Automated Optimization of the ISAC-II Heavy Ion Linac

by

Spencer Ray Kiy
B.Sc., University of British Columbia, 2010

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astonomy

Automated Optimization of the ISAC-II Heavy Ion Linac

by

Spencer Ray Kiy
B.Sc., University of British Columbia, 2010

Supervisory Committee

Dr. Oliver Kester, Co-Supervisor
(Department of Physics and Astronomy & TRIUMF)

Dr. Dean Karlen, Co-Supervisor
(Department of Physics and Astronomy & TRIUMF)

ABSTRACT

The ISAC (Isotope Separator and ACcelerator) II superconducting heavy ion linac at TRIUMF enables delivery of radioactive ion beams for various nuclear physics experiments at TRIUMF. Since the installation and commissioning of the linac in two phases (2006, 2010), it has been operated predominantly by manual adjustments of control system parameters by operators or experts. These empirical adjustments of beam optics elements, as well as amplitude and phase of accelerating cavities, constitutes a significant amount of overhead, reducing the efficiency of the facility. To implement automated optimization of the linac, a model is first built in `TRANSOPTR`, a beam envelope code widely used at TRIUMF. The control system parameters are then closely studied so as to properly calibrate the model to the real system. Finally, the model is used to automatically calculate control system settings during optimization of the linac.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

I would like to thank:

my supervisors, Oliver Kester and Dean Karlen for their patience and support throughout the writing of this thesis. I am very grateful to have benefited from your significant knowledge and experience.

TRIUMF for giving me the rare opportunity to pursue a graduate degree while continuing as a full time employee.

many others including my TRIUMF supervisor Friedhelm Ames for supporting and guiding me in my career, Marco Marchetto and Bob Laxdal for their help in digging up old field maps of solenoids, and their general expertise regarding the ISAC-II linac. Rick Baartman for including me in Beam Physics group meetings where I have learned so much about `TRANSOPTR` and beam optics. Olivier Shelbaya for the Opera modeling of the DSB buncher field map, digitization of Trace 3-D envelope simulations, help and guidance on the use of `TRANSOPTR`, and many helpful discussions. Tiffany Angus for my first introduction to the ISAC-II linac and training as an operator many years ago as well as many insightful discussions. Everyone in the RF group at TRIUMF for their assistance in setting up for measurements, general knowledge, and many useful conversations.

my in-laws, Rob and Gio for raising such a wonderful daughter in Alex and for all your help watching our girls while I snuck away on weekend days to get this thesis written.

my mom, Gerri for all the wonderful things a parent does for their kids, including watching and playing with our girls while I worked on my thesis. I'm so grateful for everything that you and dad did for me.

my wife, Alex and daughters Elliott and Maren. Alex we welcomed both our girls into the world during this thesis project and you have been incredibly supportive and patient throughout. I will forever be convinced that you did more extra work towards the completion of this thesis than I did. I can never express in words how lucky I am to have you in my life.

DEDICATION

This thesis is dedicated to my father, Dirk Kiy. You were a kind, patient, human with a passion for lifelong learning and teamwork. I have no doubt that some of what you have taught me helped me to write this thesis. You are dearly missed.

Chapter 1

Introduction

Particle accelerators use the electromagnetic force to take ensembles of charged particles, typically referred to as particle or ion 'beams', and apply time-varying electric fields in such a way so as to increase the kinetic energy of the particles. Particle accelerators all over the world provide beams for a wide variety of purposes including industry, radiation therapy, and science. This includes experiments ranging from Nuclear and Particle Physics, Astrophysics, Material Science, Life Sciences and more.

The ISAC-II (Isotope Separator and ACcelerator) superconducting heavy ion linac (LINear ACcelerator) [1] at TRIUMF accelerates stable or radioactivate ion beams produced in ISAC-I up to as fast as 18% the speed of light. The ISAC-II linac provides ion beams to one of 4 separate experimental stations in the ISAC-II experimental hall, primarily conducting experiments studying the structure of rare nuclei or rate of nuclear reactions of interest in astrophysics.

1.1 The ISAC-II Heavy Ion Linac at TRIUMF

The ISAC-II linac at TRIUMF is primarily optimized (or 'tuned') by manual adjustments of control system parameters (focal strength of quadrupoles, timing of accelerating cavities). The adjustments are typically made to one device at a time while observing the effect on some property of the beam (size, intensity, energy) downstream of the linac. To our dismay, humans can only process a limited number of inputs/outputs simultaneously, meaning often the beam property being observed and optimized (e.g. beam intensity) is done so without simultaneously considering

[1] from here on colloquially referred to as the 'ISAC-II linac'

other outputs of interest (robustness of the tune; divergence, beam size, etc).

This presents one primary issue:

- *Operation of the linac requires a significant amount of time (overhead) from highly experienced personnel, negatively impacting the total deliverable hours of rare ion beams to experiments.*

  While the electromagnetic force has been well understood for some time and was of course used in the design of the linac, our knowledge of electromagnetism is not being fully utilized in regular operation. As a result, the relationship between cause (adjusting system parameters) and effect (changes to the beam properties) is treated as a 'black-box' rather a than well understood recipe. The operation of the machine is empirical rather than model-based.

Above is a general statement that encapsulates the core of the problem to be addressed, however it doesn't quite articulate the particular physical problems to be tackled. These can be summarized as follows:

1. *The ISAC-II linac has not been modelled in a code that enables real-time simulations of the machine*

2. *The effect of the beamline elements providing transverse focusing (quadrupoles, solenoids) must be more carefully benchmarked to the simulated beam enevelope*

3. *The parameters in the control system used to adjust the phase and amplitude of accelerating cavities have not been well calibrated to the physical properties of the cavities*

4. *An algorithm for sequential optimization of the linac has not yet been implemented, primarily due to the fact it has as prerequisites the above three points*

With the ARIEL (Advanced Rare IsotopE Laboratory) [1] facility coming online over the next five years, the available hours for delivery of RIBs (radioactive ion beam) to experiments are expected to triple. Transitioning from one to three simultaneous RIBs while minimizing the increase in personnel presents a significant challenge and motivates many angles of attack [2] [3] [4] at TRIUMF to improve automated and model-based tuning methods. This thesis constitutes one of these projects.

To address the identified problems, the proposed new approach is as follows:

1. *The ISAC-II linac will be modelled in the beam envelope code* `TRANSOPTR`*, and compared to previous simulations in the code Trace 3-D*

2. *The simulated beam envelope in* `TRANSOPTR` *will be compared to beam size measurements at available diagnostics*

3. *The control system variables for phase and amplitude of each cavity will be calibrated to measured data*

    a *Independent of the beam, the effect of moving the phase setpoint in the control system on the output signal going to the cavity amplifiers will be measured*

    b *Beam-based data will be collected, recording the energy provided by each cavity while varying the phase*

    c *Measurements will be used to calculate both the phase offset of each cavity, as well as the peak electric field in each cavity for a given amplitude in the control system*

4. *An algorithm for sequential optimization of the linac within the* `TRANSOPTR` *model will be implemented*

Finally, with the new approach implemented, a number of model-based adjustments to the ISAC-II linac configuration will be applied to the machine, and the results of these adjustments discussed.

As will be shown in this thesis, the resulting implementation and measurements have provided an improved understanding of the machine. In a broad sense, this will reduce overhead, improve institutional knowledge, allow easier identification of system faults, and expand the number of staff members able to operate the linac.

1.2 Claims

To verify my new approach to the optimization of the ISAC-II linac, I make *four* claims about the new paradigm that demonstrate the problems identified in the previous section have been addressed.

1. The transverse beam size at available diagnostics has qualitatively good agreement with `TRANSOPTR` simulations.

   This implies that the effects of the beamline optics and RF devices on the transverse beam envelope are understood

2. Using automated calculation of linac parameters, the model predicts cavity phase to within five degrees and output energy to within 2% at the output of each operational cavity.

   This demonstates that the peak accelerating field and phase offset (with respect to a common reference) are sufficiently well known such that beam will remain within the acceptance of the next cavity

3. Automated calculation of linac parameters results in a beam envelope where more than 90% of the beam is successfully transported to the end of the linac.

   This shows that the calculated settings are sufficiently accurate to keep the transverse beam envelope within the necessary apertures and to keep the beam entering each cavity within the necessary acceptance in longitudinal phase space

4. Automated calculation of linac parameters predicts the final beam energy to within 1% (for a configuration require all operational cavities).

   This indicates that the calculated settings, along with the chosen accelerating phase of -25 degrees, are sufficient to accurately set the final beam energy

The overarching consequence of these four claims is that an improved understanding of the machine has been attained, enabling a significant reduction in facility overhead.

1.3 Outline of Chapters

Chapter 1 introduces the ISAC-II linac and briefly describes the new developments to be introduced, why these are needed, and the expected impact.

Chapter 2 dives into the details of the previous methods for operation of the ISAC-II linac, the main problems to be addressed, and why they are relavant.

Chapter 3 contructs a comprehensive picture of the new approach for operation of the ISAC-II linac, laying out the details of its implementation, provides an

introduction to the envelope code TRANSOPTR, and explains the algorithm used for automated optimization of the linac.

Chapter 4 catalogues the experimental setup utilized and the methodology behind the data collection. First described is the data collection for calibration of the model. This is followed by an explanation of the four tests done to demonstrate improved control over the linac.

Chapter 5 provides analysis of the collected data. This includes comparing TRANSOPTR to Trace 3-D, analyzing the accuracy of the TRANSOPTR model to observations of beam properties, and comparing the newly implemented approach to what has been utilized at other labs.

Chapter 6 summarizes the claims made from chapter 1 and how they were demonstrated. Opportunities for future improvement and their potential impact on linac performance are discussed.

Chapter 2

Empirical Optimization of the ISAC-II Linac

2.1 Description of the ISAC-II Linac

2.1.1 Overview

The ISAC-II linac is designed to accept heavy ion beams from the ISAC-I facility [5] at a velocity of $\beta = 0.057$ with a mass-to-charge ratio [1] between 2 and 7. The linac can be conceptually separated into three sections: DSB (Drift-tube linac to Superconducting Linac Beamline), SCL (Superconducting Linac), and SEBT (Superconducting Linac to Experiment Beam Transport).

These beamlines are composed of a myriad of different devices that serve various purposes. This includes magnetic quadrupoles, dipoles and steerers, superconducting solenoids, RF cavities, and diagnostics to measure beam position and intensity.

[1] mass-to-charge ratio in accelerator physics refers to the mass (in atomic mass units) of the isotope composing the ion beam divided by the charge of the isotope in an integer number of +e

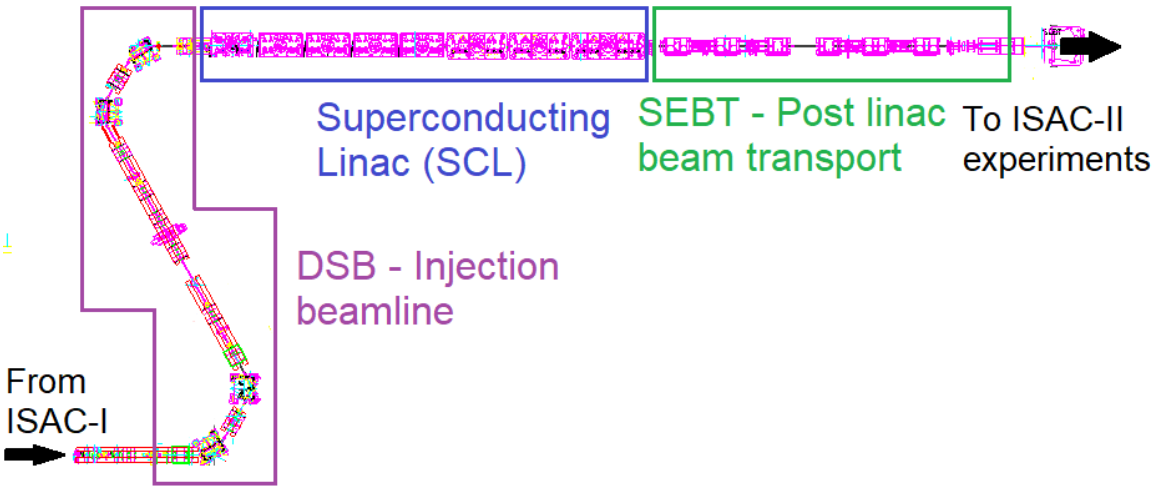

Figure 2.1: The ISAC-II superconducting heavy ion linac shown here, contains three primary sections: the injection beamline 'DSB', the linac itself 'SCL' and 'SEBT,' the post-linac transport and diagnostic section

2.1.2 DSB - Injection Beamline

The DSB transfer line begins at a diagnostics box downstream of the ISAC-I drift tube linac (DTL). The DSB beamline continues for just under 28 m until it meets the start of the SCL. This section consists of twenty magnetic quadrupoles, four magnetic dipoles, and one room-temperature bunching RF cavity, as well as many diagnostics for beam profile and intensity measurements. The DSB beamline can be further subdivided into four tune sections in the following sequence:

1. A doubly achromatic[2], focus-to-focus, -119 degree bend section (6 quads, 2 dipoles)

2. A four quad focus-to-focus transport that ends at the centre of the DSB buncher cavity.

3. A four quad focus-to-focus transport that ends at the start of the final bend section.

[2]In a doubly achromatic bend, two criterea are met: at the exit of the bend there is no correlation between energy and horizontal position, AND the derivative of the energy-position correlation is 0.

4. A doubly achromatic, focus-to-focus, +119 degree bend section (6 quads, 2 dipoles)

2.1.3 SCL - The Superconducting Linac

Consisting of fourty superconducting coaxial quarter wave cavities and eight superconducting solenoids distributed over eight cryomodules, the SCL is where the beam is accelerated to as high as 18% the speed of light over a distance of 19.3 metres. There are three unique cryomodule designs and three unique cavity designs (see figure 2.2), distributed as follows:

- The first two cryomodules, referred to as SCB1 and SCB2 each contain four 'flat' type cavities that operate @ 106.08 MHz
- The next three cryomodules, referred to as SCB3, SCB4, and SCB5 each contain four 'round' type cavities that operate @ 106.08 MHz
- The sixth and seventh cryomodules, SCC1 and SCC2, each contain six 'donut' type cavities that operate @ 141.44 MHz
- The last cryomodule, SCC3, contains eight 'donut' type cavities that operate @ 141.44 MHz

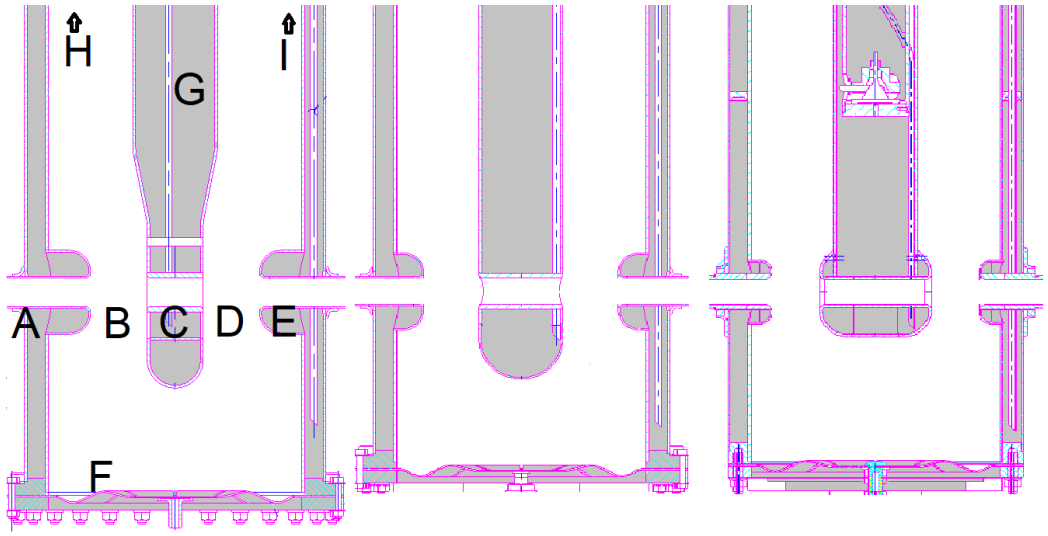

Figure 2.2: Shown here are the three cavity types used in the SC linac. From left to right they are named 'flat', 'round', and 'donut' in reference to the shape of the inner conductor. Components of the flat cavity are labelled: A is the entrance beam port, B is the first acceleration gap, C is the drift tube in the inner conductor through which the beam travels as RF changes sign, D is the second acceleration gap, E is the exit beam port, F is the tuner plate, G is the inner conductor, and H/I (not shown) are the RF coupler and pickup ports

2.1.4 SEBT - Post-linac Beam Transport

The SEBT beamline beginning at the exit of the SCL is composed of 12 quads and 12 steerers. The primary diagnostics for tuning of the beamline are located in this section, which are the three flight time monitors (FTMs) spanning a total distance of just over 12 meters.

2.1.5 Beamline optics

The devices along a beamline which exert forces on the charged particles composing a particle beam are often divided conceptually between two categories: transverse (acting perpendicular to the direction of motion) and longitudinal (acting primarily along the direction of motion). The Lorentz force law

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}) \tag{2.1}$$

shows clearly that while an electric field can exert forces in any direction (the

direction of the field), magnetic fields may only act perpendicular to the ion's velocity.

Transverse optics found in the ISAC-II linac include magnetic quadrupoles and superconducting solenoids. Quadrupoles are designed so as to have a magnetic field in the x direction of increasing intensity as you move further from the beam axis in the y direction (and conversely, swapping x/y). This means there is an approximately linear increase in the field, such that particles further from the beam axis receive a stronger correction back towards the beam axis compared to near-axis particles.

$$B_x(y) = Gy \tag{2.2}$$

$$B_y(x) = Gx \tag{2.3}$$

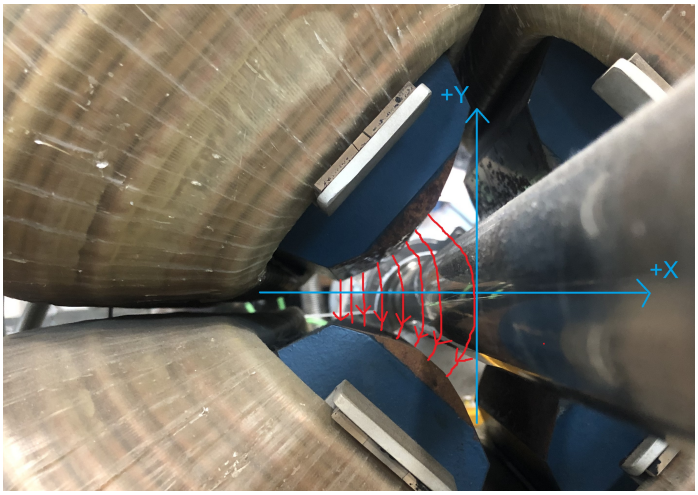

Figure 2.3: Danfysik 1982 L5 Magnetic Quadrupole, with magnetic field lines overlaid in red for clarity.

Where the constant G is the field gradient of the quadrupole. A given quadrupole will provide a focusing force in one transverse plane and a defocusing force in the other. Multiple quadrupoles in succession however can provided a net focusing effect (if desired), known as the strong focusing principle.

The ISAC-II linac contains eight superconducting 9 Telsa solenoids that provide transverse focusing.

Four magnetic dipoles in the DSB injection beamline serve to turn the beam around corners

Magnetic steerers throughout the linac in theory should not be needed for a perfectly aligned system. They exist due to the fact that the alignment process is in fact

imperfect and small adjustments need to be made to retain the beam on the desired axis and minimize the steering effect of focusing devices.

2.1.6 RF Cavities

Accelerating RF cavities make use of electric fields so as to act on the beam in the longitudinal direction. This is most often used to increase the kinetic energy of the beam, although RF cavities are also used to provide focusing in the longitudinal direction.

There are a wide diversity of RF cavities to be found in the field of accelerator physics. They are typically roughly categorized by their shape (pillbox, coaxial, elliptical), resonant frequency and excitation mode, number of accelerating gaps, design velocity (of the particles to be accelerated), and accelerating gradient.

At a basic conceptual level, an RF cavity is a metal [3] box within which an electromagnetic field can be excited and contained. The simplest example, a one-gap pillbox cavity is shown in Figure 2.4. The bunch of ions transit the cavity as the electric field is large and co-linear with the beam's total momentum. Not shown is the other half of the RF cycle, when the magnetic field would rise in the opposite direction then fall off, followed by electric field rising in the reverse direction and then falling off, then the next bunch of ions approaches and the cycle begins again. It is important for acceleration that the bunches or particles are either far away from the cavity or shielded while the field changes sign[4].

[3] Typically Copper for a normal-conducting cavity, or Niobium for a super-conducting cavity

[4] Multi-gap cavities use drift tubes which shield the particles from the electric and magnetic fields for the portion of the RF cycle when the electric field is anti-parallel to the beam's momentum)

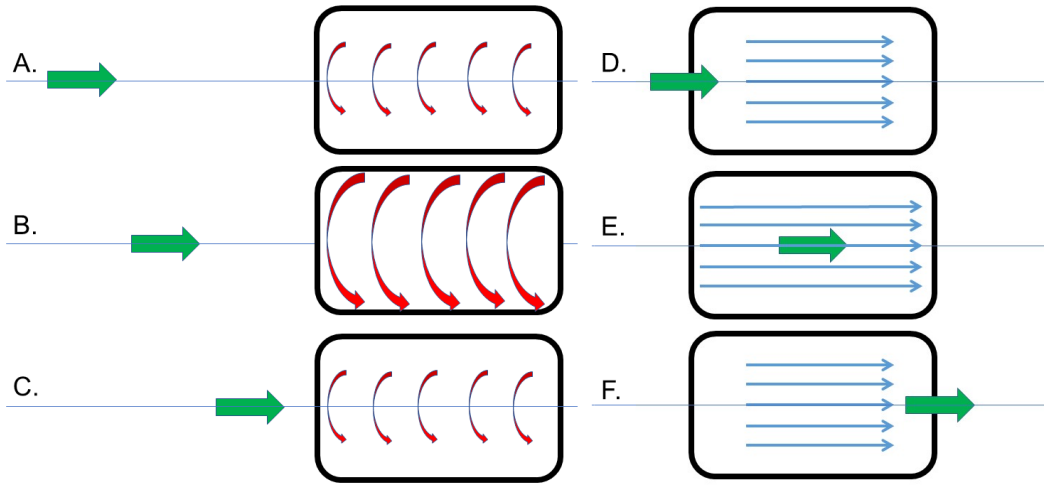

Figure 2.4: A bunch of ions (shown as a green arrow) approach the cavity (depicted as a black rectangle) as the axial magnetic field (red) is large but decreasing in magnitude (B and C). The ions enter the cavity as the electric field in the cavity is rising (D) and reach the centre of the cavity when the field is near it's maximum (E). The ions exit the cavity as the electric field is again decreasing (F)

From here, cavity designs get more complex but the principle remains the same - an electomagnetic wave is confined within a cavity and the beam transverses this cavity in such a way so as to extract energy via the electric field. For an on-axis electric field profile $\mathcal{E}_z(s)$ where s represents the distance along the beamline, the energy gained ΔE can be calculated by integrating through the field as it varies in time t

$$\mathcal{E}_z(s,t) = \mathcal{E}_z(s)cos(\omega t + \phi - \theta) \tag{2.4}$$

$$\Delta E = q \int_{-l/2}^{l/2} \mathcal{E}_z(s,t)ds \tag{2.5}$$

Where ω is the angular frequency of the RF, l the length of the cavity, ϕ the adjustable phase of the cavity, θ a fixed offset of the cavity phase, and q the ion's charge.

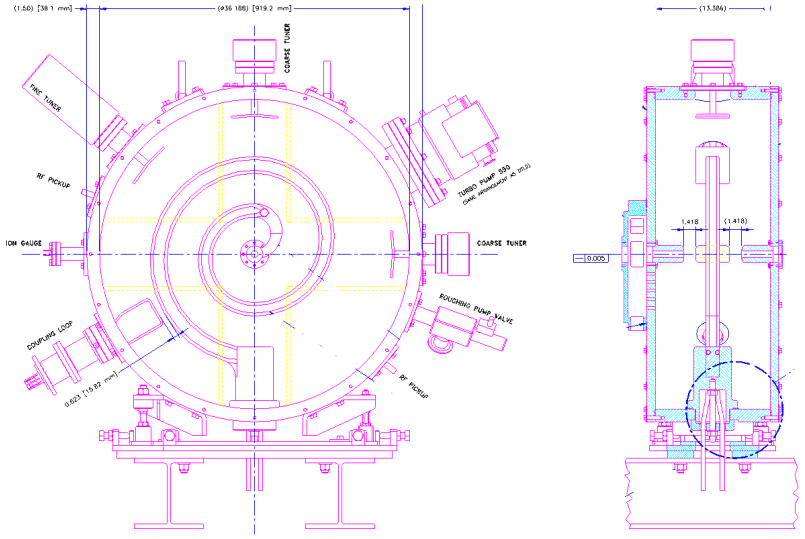

Figure 2.5: Schematic of the spiral two-gap DSB buncher. (left) Front view as-seen by the approaching ion beam with the drift-tube in the centre of the spiral, (right) side view showing the two RF gaps on either side of the drift tube at the centre

In the ISAC-II linac, the room temperature bunching cavity and all fourty superconducting accelerating cavities are two-gap cavities. For the beam to be accelerated in each gap of a multi-gap cavity, RF phase needs to advance by either π or 2π depending on the mode the cavity operates in.[5]

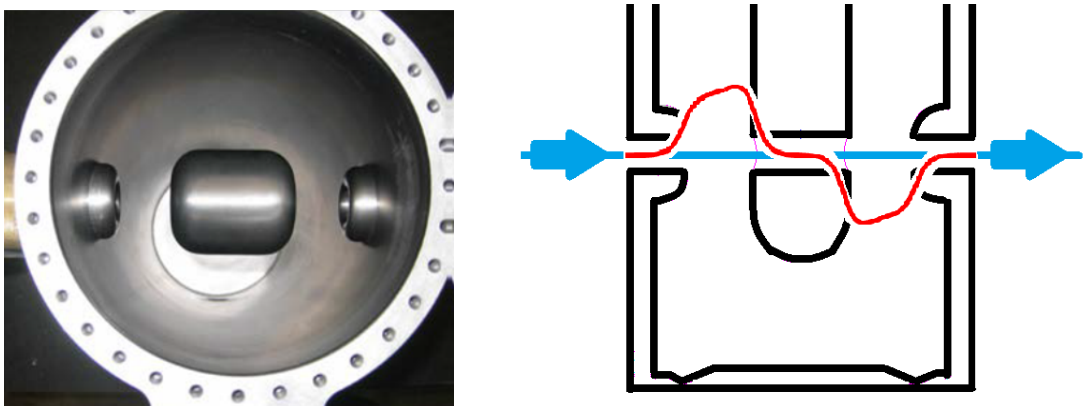

Figure 2.6: Above left is a picture of a round type cavity. Beam ports can be seen on the outer wall on either side, with the inner conductor in the centre of the picture. To the right is a side-view schematic of the same cavity type with the beam axis overlaid in blue, and the E_z field map (at one instant in time) overlaid in red

[5]For our purposes all cavites operate in π mode, meaning adjacent gaps are out of phase with each other by π

Between gaps, the beam enters a region where the electric field is quite small regardless where the cavity is within its RF cycle. For both the DSB buncher and ISAC-II cavities, this is within a 'drift tube' in the central conductor of the cavity. The field does not penetrate into this drift tube and thus the ions are not decelerated while the RF phase advances by π.

The fourty superconducting cavities found in ISAC-II are made of bulk Niobium, a Type-II superconductor with a critical temperature of 9.2 K. These cavities are cooled with LHe at 4.2 K. The cavity design is referred to as a co-axial quarter wave cavity. Twenty of the cavities operate at 106.08 MHz while the other twenty operate at 141.44 Mhz. An example of the E_z profile overlaid on top of the cavity schematic is shown in Figure 2.6 while a cavity schematic with overlaid electric fields is shown in Figure 2.7.

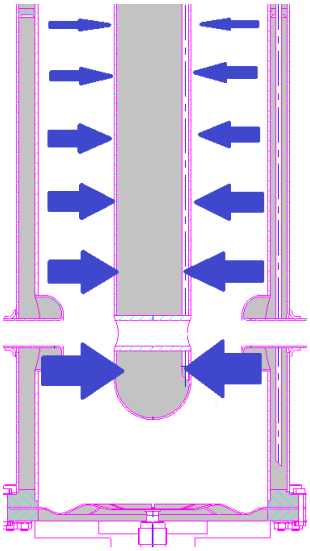

Figure 2.7: ISAC-II 'Round' cavity type, showing relative location and strength of the electric fields.

2.2 Digression for a Few Definitions

It is useful at this point to clarify some definitions commonly used when referring to properties of an accelerating cavity:

- *Electric field*, $\vec{\mathcal{E}}(x, y, s, t)$ - this is the electric field at any point and time within the cavity.
- *Peak surface field*, $\mathcal{E}_p$ - the peak surface field of an RF cavity is the highest magnitude electric field found anywhere within the cavity. Typically these occur near corners and this quantity is of interest mainly to SCRF engineers as there are theoretical limits to $\mathcal{E}_p$ due to field emission.

$$\mathcal{E}_p = max \left| \vec{\mathcal{E}}(x, y, s, t) \right| \tag{2.6}$$

- *Longitudinal component of the on-axis electric field*, $\mathcal{E}_z(s, t)$ - this is the component of the electric field of interest in the acceleration of the beam. Often colloquially referred to as the 'on-axis field'.

$$\mathcal{E}_z(s, t) = \hat{z} \cdot \vec{\mathcal{E}}(x = 0, y = 0, s, t) \tag{2.7}$$

- *Field map* or *Time-independent on-axis field*, $\mathcal{E}_z(s)$ - this is the on-axis field at one instant in time, specifically the time t^* at which the field is at a maximum magnitude.

$$\mathcal{E}_z(s) = \mathcal{E}_z(s, t^*) \tag{2.8}$$

- *Peak accelerating field*, $\mathcal{E}_z$ - this is the highest magnitude electric field found along the beam axis within the cavity. The peak accelerating field is the quantity/variable used to scale the normalized on-axis field map within the envelope code `TRANSOPTR`.

$$\mathcal{E}_z = max\,|\mathcal{E}_z(s)| \tag{2.9}$$

- *Normalized Time-independent on-axis field*, $\mathcal{E}_{norm}(s)$ - this is the field map $\mathcal{E}_z(s)$ relative to its maximum.

$$\mathcal{E}_{norm}(s) = \frac{\mathcal{E}_z(s)}{\mathcal{E}_z} \tag{2.10}$$

- *Average electric field*, $\mathcal{E}_0$ is the average electric field over the cavity length L of the time-independent on-axis field $\mathcal{E}_z(s)$.

$$\mathcal{E}_0 = \frac{1}{L}\int_{-L/2}^{L/2} |\mathcal{E}_z(s)|ds \tag{2.11}$$

- *Effective voltage*, V_{eff} - The effective voltage of a cavity is a representation of how much energy it transfers to an ion. It is equal to the kinetic energy gained by the ion divided by their charge. Effective voltage of a single cavity can vary depending on its phase, the A/q of the beam, applied RF power, and the incoming beam velocity - so it can be misleading to state a value for V_{eff} without including the neccessarly supplementary information. If unspecified, it is this author's experience that most often, an underspecified V_{eff} refers to the incoming beam velocity matching the design velocity of the cavity, maximum accelerating phase ($\phi = 0°$), maximum operable RF without breakdown, and the maximum design A/q for the accelerator.

$$V_{eff} = \frac{\Delta E}{q} \tag{2.12}$$

- *Effective (or cavity) length*, L - This most often refers to the distance between the interior of the two cavity walls, although this isn't always the case. This is a useful figure as it is typically economically beneficial to maximize the effective voltage V_{eff} while minimizing the length of the cavity L.

- *Accelerating Gradient*, $\mathcal{E}_a$ or G - This is the most commonly quoted figure of merit for an accelerating cavity and is simply the effective voltage divided by the cavity's length. As with V_{eff} it is important to qualify other parameters of interest when quoting a gradient.

$$G = \frac{V_{eff}}{L} \tag{2.13}$$

  For clarity I will define $\mathcal{E}_a$ and V_{eff} to be at incoming beam velocity matching the design velocity of the cavity, maximum accelerating phase ($\phi = 0°$), and A/q = 6)

- *Control system amplitude*, A_{cs} - This refers to the value[6] in the control system which represents the peak accelerating field scaled to some arbitrary units. This quantity is relevant, as an algorithm modelling a cavity based on current RF cavity performance needs to first read A_{cs} from the control system as a boundary condition.

$$\mathcal{E}_z = \mathcal{E}_z(A_{CS}) \approx A_{CS}C \tag{2.14}$$

  where C varies from one cavity to the next.

- *Accelerating (or relative) phase of a cavity*, ϕ_{rel} - In accelerator physics, it is conventional to express the energy gain as a function of the cosine of an accelerating phase (see 2.17) such that particles see maximum acceleration when the phase is 0. The convention further describes that accelerator phases between 0 and -180 are *bunching* (electric field increasing as the ions cross the cavity)

[6] A signal from a pickup in the magnetic field region of each cavity is read into the control system via a power meter and ADC

while phases between 0 and +180 are *de-bunching* (electric field decreasing as the ions cross the cavity)

- *Cavity (control system) phase*, ϕ_{cs} - To achieve acceleration, RF cavities in a linac must be on the same 'clock' and the physicist must have a means to adjust the timing of individual cavities relative to this clock such that they exert the desired forces on the arriving bunches of ions. This is done by means of a master oscillator which acts as the main clock for all RF including the RF that creates the beam bunches from an initial continous beam of ions. For independent cavity linacs such as the ISAC-II linac, each cavity has a 'phase shifter' which is a variable (knob) in the control system that can be adjusted to change the timing of the cavity with respect to the beam.

$$\phi_{rel} = \omega t + \phi_{cs} - \theta \tag{2.15}$$

Where θ is another calibration parameter, a phase offset that can vary from one cavity to the next and is related to the differing lengths of RF cables.

- *Synchronous phase*, ϕ_s - Most multi-gap cavities are designed to operate at a specific voltage. When crossing each accelerating gap i, the design/reference particle may be some non-zero relative RF phase $\phi_{rel,i}$ away from receiving maximum acceleration in that gap to provide longitudinal stability. The synchronous phase is the phase such that the relative phase in each cell of a multi-gap cavity is the same.

Mathematically, for an n-gap cavity, the synchronous phase satisfies:

$$\phi_s = \phi_{rel,i} \qquad \text{for all } i \in 1..n \tag{2.16}$$

2.3 Transit time factors - to integrate now or to integrate later

A widely utilized concept in the field of accelerator physics is that of the so-called 'transit-time factor,' often abbreviated TTF. A particle transiting the accelerating

gap sees a varying electric field and thus, the actual energy gained [7] is always less than one would expect from just naively integrating through the time-independent peak electric field profile $\mathcal{E}_z$.

Cavities that consist of multiple gaps typically have a further reduction in efficiency related to the timing of the particle at each gap. For a two gap cavity, the cavity phase could be adjusted such that a particle experiences the maximum possible acceleration in the first gap, but arrives at the second gap too early or too late and thus does not experience maximum acceleration through the second gap. Multi-gap cavities are built for a design beta β_s, which is a velocity at which (in a simple approximation) the beam transits the distance between the centre of each gap in a half RF period.

With the assumption that the velocity change through the cavity is negligible and that the electric field is uniform over the acceleration gap, Wangler [6] derives the energy gain ΔW and transit-time factor, T, of a two gap cavity to be

$$\Delta E = q\mathcal{E}_0 T L cos(\phi_{rel}) \tag{2.17}$$

$$T = \frac{sin(\pi g/\beta\lambda)}{\pi g/\beta\lambda} sin(\frac{\pi\beta_s}{2\beta}) \tag{2.18}$$

Where q is the charge, ϕ_{rel} the relative phase of the cavity, L the cavity length, g the length of the accelerating gaps, β the incoming particle velocity over c, λ the RF wavelength, and β_s the design velocity of the cavity. $\mathcal{E}_0$ is the average accelerating field over the cavity length L.

For the ISAC-II linac, TTF and the analytic expression for energy gain 2.17 have typically been used to calculate cavity performance. For improved accuracy, a CST microwave studio simulation of the 3-dimensional cavity fields was performed, and particle tracking simulations through these fields done to get a more accurate measure of TTF as a function of β (V. Zvyagintsev, Pers. comm.). The analytic 2-gap ttf expression was then fit to this data, allowing the gap length g and design velocity β_s to vary. The three TTF curves for the ISAC-II cavities are shown in Figure 2.8.

[7] For a gap with a perfectly symmetric $\mathcal{E}_z$ around the geometric centre of the gap and in the limit of $\Delta v << v_i$, maximum acceleration occurs if particle is at the centre of the gap when the electric field stength is at a maximum

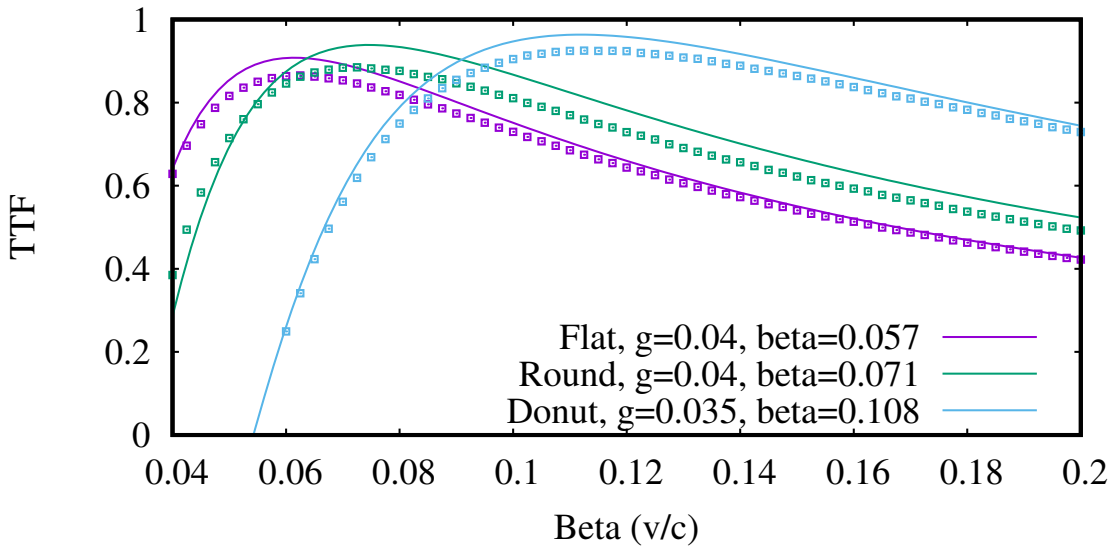

Figure 2.8: Analytic form of the TTF curves for ISAC-II cavities in solid lines compared to `TRANSOPTR` (squares)

Essentially, what this boils down to is that models using the TTF approach are integrating through the cavity fields using particle tracking codes offline so that the model can be simplified down to the analytic expressions 2.17 and 2.18. However, due to the assumptions made in the approximation, it encounters issues as the A/Q of the beam is changed. As shown in Figure 2.9, the TTF curves are not actually the same as the A/Q is varied.

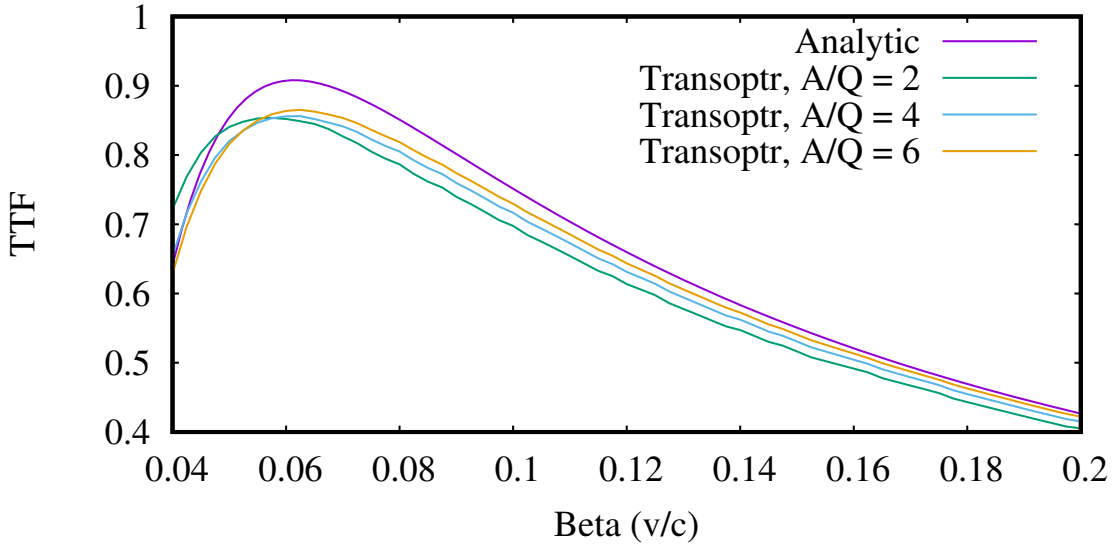

Figure 2.9: `TRANSOPTR` generated TTF curves for ISAC-II flat cavity at varying A/Q's

Another method is possible - integrate through the fields in real-time while modelling the energy gain through the linac. Basically, leave the differential equation 2.19 unsolved, and numerically solve it when running your envelope code.

$$\frac{dE}{ds} = qA_iC_i\frac{\mathcal{E}_z(s)}{\mathcal{E}_z}cos(\omega t + \phi_i - \theta_i) \tag{2.19}$$

Where for cavity i, A is the control system amplitude, C is the scaling constant between the control system and physical field within the cavity, ϕ_i is the control system phase setting, and θ_i is the fixed phase offset relative to the reference signal associated with that cavity.

2.4 Considerations for Linac Operation

Historically, when preparing for delivery to a given experiment, there are two primary variables that result in a need for re-configuring of the linac (or 're-tuning'):

1. RF cavity performance
2. Requested beam properties

RF cavity performance

The fourty cavities in the linac were specified for 1.08 MV of *effective voltage* [5], however due to various reasons the actual performance varies from cavity to cavity. As of writing, two cavities are not operational - one due to an issue with operation of the tuner, the other due to a damaged RF transmission line. The remaining 38 cavities vary in their performance [7] for various reasons including initial purity of the cavity material, quality of the assembly, conditioning status, and contamination of the cavity due to an influx of particulates. The SCRF group at TRIUMF is responsible for monitoring and maintaining cavity performance. Typically at the start of each experiment or at a minimum the beginning of each year, the SCRF group condition cavities and identify their stable operating setpoints, which are taken as inputs for linac operation.

Since the initial installation of the ISAC-II linac, cavity amplitudes have been operated at the operational setpoints identified by the SCRF group, with the exception of the amplitude of the last cavity which is reduced to match the requested energy. This was done to get the most acceleration possible out of the fewest cavities, and reduced the number of variables involved in the linac optimization process. However, in principle with a calibrated model of the linac the model can be permitted to vary the cavity amplitudes when optimizing for a desired output.

Requested beam properties

The ISAC-II linac accelerates beam for delivery to one of three experimental stations for various nuclear structure and nuclear astrophysics experiments [8] [9] [10]. The requested beam energy is identified depending on the planned experiment and can vary from as low as 1.53 MeV/u (no acceleration required, all RF off) to as high as 16.5 MeV/u for beams with sufficiently low A/Q.

Typically, an upcoming experiment specifies the following requirements: isotope, intensity, energy, and purity. The Beam Delivery Group at TRIUMF discuss these requirements and machine capabilities with the users to further identify the charge states(s) [8] best suited to meet the request.

[8] While the charge state at the RIB sources at TRIUMF are typically singly charged (1+), a number of options are available to sacrifice intensity in order to boost the charge state which is often desirable to overcome energy or purity concerns. This can be done in an ECRIS or EBIS-type charge breeder or by inserting a thin foil into the beam at higher energies to strip off additional electrons.

2.5 Beam Envelope Simulations

2.5.1 Brief Introduction to Envelope Optics

With an introduction to the linac complete, it is now timely to peform a brief review of envelope optics before digging into the specifics of the code used for the ISAC-II linac simulations. Readers already experienced in beam physics and envelope optics may prefer to skip ahead to Section 2.5.2.

Coordinate System

A beam transport system and accelerators are designed for an ideal particle that follows what is known as the reference trajectory - the intended path the physicist paves out for the particle. It is convenient to define a co-moving non-intertial coordinate system (known as a Frenet-Serret frame) that follows the ideal particle through the reference trajectory.. The three axes in this reference frame are typically referred to x, y, and z. The letter s is reserved for a distance along the reference trajectory while z refers to the distance along s that a given particle is from the reference particle.

Ensembles of particles

Reality is rarely so simple and not all particles are cut out to follow the reference trajectory. Two factors play a role in this - engineering and physics.

Ions are first neutral atoms at some non-zero temperature, typically either floating around a plasma chamber or within a hot thin tube of metal [9]. Their non-zero temperature means these atoms have some spread in momentum in all three spatial dimensions. When ionized, they then see a large potential well towards the ground electrode and all recieve a large increase in momentum along the reference trajectory. However, they still retain their small momentum differences in x, y, and z. Further to this, because the plasma chamber or ionizing tube has a non-zero spatial extent in x and y, there is also a small but non-zero spread in the positions of the ions in x and y.

Even were some marvelous feat of engineering to cause the position and momentum spread of the ions from the source to approach zero, physics literally pushes back, as the positively charged ions will exert a repulsive Coulomb force on one other.

[9] Often called the 'ionizer' or 'ionizing tube'

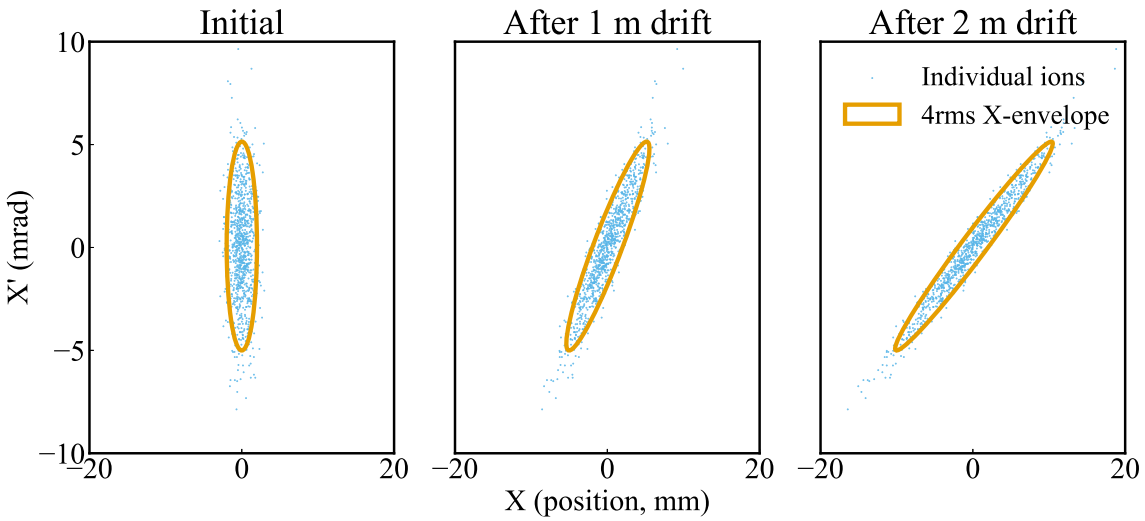

Figure 2.10: Demonstration of the evolution of a 1000 particle ensemble and the 4rms ellipse representation as the beam transits through a drift space of 2 m

And so, we must accept the need to describe collections of particles that have some momentum and position spread around the reference trajectory. This means we have a 6-dimensional phase space (x, x', y, y', z, z') [10]. Luckily, a well designed system will typically have no correlation between the x, y, and z coordinates and we can approach the matter initially by considering a 2-dimensional phase space [11].

In Figure 2.10 we consider a collection of 1000 particles and plot them in the $X - X'$ phase space with no initial correlation and following a gaussian distribution [12]. As this beam drifts down the beamline, X' remains constant while positions change depending on X'.

$$x'_1 = x'_0 \tag{2.20}$$

$$x_1 = x_0 + L \cdot x'_0 \tag{2.21}$$

Where L is the length of the drift space. Writing the position and angle of each particle as a vector (x, x'), this transformation can be written with the simple matrix

[10] Note here that z/z' are not always used as the 5th and 6th coordinates but are more intuitive for this introduction. They are more often instead written as $\Delta W/\Delta\phi$ or $\Delta E/\Delta t$

[11] Keeping in mind that this is just a projecting of the 6-dimensional phase space projected onto the $X - X'$ plane

[12] In reality, the distribution in X is uniform within the circular aperture of the ionizing tube. This means projected onto the x-axis the distribution in X should follow a Wigner semicircle distribution with a hard edge, as no ions can be created or transported at $|x| > r_{tube}$

formula

$$\begin{bmatrix} x_1 \\ x_1' \end{bmatrix} = \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x_0 \\ x_0' \end{bmatrix} \tag{2.22}$$

Where the matrix $\left(\begin{smallmatrix} 1 & L \\ 0 & 1 \end{smallmatrix}\right)$ is referred to as the *transfer matrix*, M. This is helpful, but if we want fast simulations we don't want to repeat this type of calculation for thousands of particles. Now we note that for distributions of multiple variables we can calculate a covariance matrix $\boldsymbol{\sigma}$, which is composed of the second moments of the beam distribution in the (x, x') phase space.

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_x^2 & \sigma_{xx'} \\ \sigma_{x'x} & \sigma_{x'}^2 \end{bmatrix} \tag{2.23}$$

where

$$\sigma_x^2 = \frac{1}{n}\sum_{i=1}^{n}(X_i - \mu_x)^2 \tag{2.24}$$

$$\sigma_{xx'} = \frac{1}{n}\sum_{i=1}^{n}(X_i - \mu_x)(X_i' - \mu_{x'}) \tag{2.25}$$

The covariance matrix $\boldsymbol{\sigma}$ describes the ellipse shown in Figure 2.10 and transforms [11] as:

$$\boldsymbol{\sigma_1} = \boldsymbol{M} \cdot \boldsymbol{\sigma_0} \cdot \boldsymbol{M^T} \tag{2.26}$$

Thus, any known transfer matrix $\boldsymbol{M}$ can be applied to the covariance matrix, and this is primarily how envelope codes work - they track the covariance or sigma matrix $\boldsymbol{\sigma}$ through the beamlines, applying a stored matrix for each element encountered along the way. This allows envelope codes to execute in times orders of magnitudes faster than particle tracking codes.

Courant–Snyder Parameters

The sigma matrix description of the ellipse in phase space can be written [12] as:

$$\sigma_{22}x^2 - 2\sigma_{12}xP_x + \sigma_{11}P_x^2 = \epsilon^2 \tag{2.27}$$

However, it is also common to transform the ellipse description to the so called

Courant–Snyder (often referred to as Twiss) parameters. This is convenient, as σ_{11}, σ_{22}, and σ_{12} in Equation 2.27 themselves are dependent on the beam emittance ϵ. This means that for two different emittances, fitting a transfer matrix to achieve a certain σ_{11} will result in different values for σ_{22}. In contrast, the Courant–Snyder parameters α, β, and γ as shown below are independent of the emittance, such that they describe a beam envelope through some optical elements that will scale with the square root of the beam emittance.

$$\gamma x^2 + 2\alpha x P_x + \beta P_x^2 = \epsilon \tag{2.28}$$

Some of the relationships between the second moments of the beam distribution, the 4rms ellipse, and the Courant–Snyder parameters are shown graphically in in Figure 2.11.

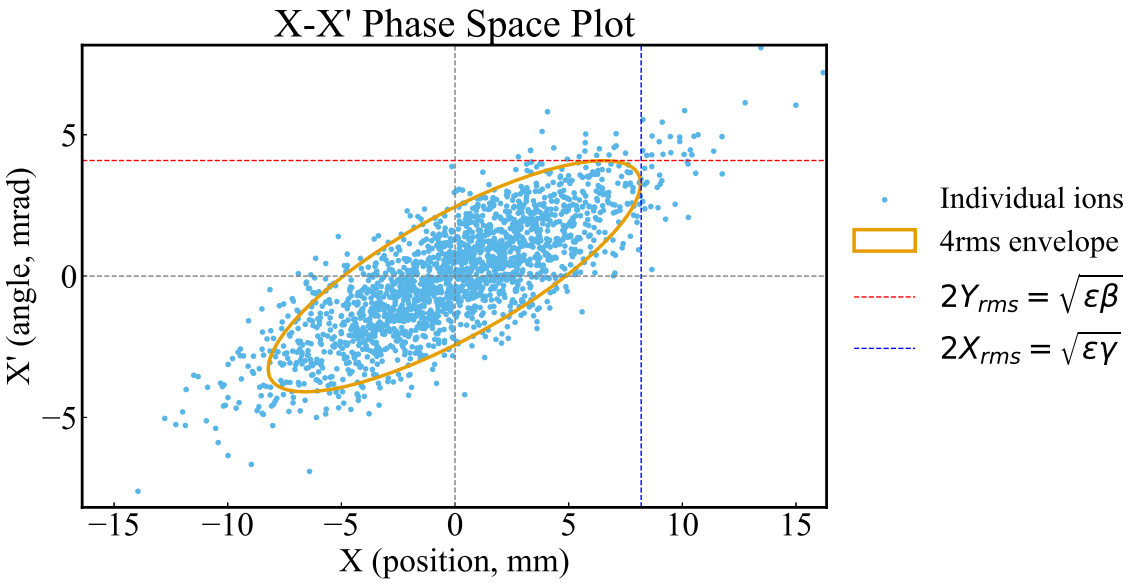

Figure 2.11: Representation of the relationship between the many ions (2000 in this simple example) composing a beam and the 4rms envelopes.

2.5.2 Beam Envelope Simulations in Trace 3-D

The ISAC-II linac was designed primarily using one particle tracking code and one beam envelope code [5].

The particle tracking code used was LANA (Linear Accelerators Numerical Analysis) developed at the Institute for Nuclear Research (INR) in Russia [13]. It was

used at TRIUMF as early as 1997 [14]. While particle tracking (or 'ray-tracing') codes have many applications in accelerator physics, they are outside the scope of this thesis.

The beam envelope code used in the design of the ISAC-II linac was Trace 3-D [15]. Developed at Los Alamos National Laboratory (LANL), Trace 3-D uses the σ-matrix formalism introduced by `TRANSPORT` [16] to track the second moments (variance) of the beam distribution (x, x', y, y', z, z') and their correlations as the beam moves through the beamline along the reference trajectory.

After being used for initial beamline design, Trace 3-D has continued to be the primary code used for calculating the necessary optics settings. For a given tune section, a single ion beam, charge state[13], and energy was used as an input to a Trace 3-D simulation, which was optimized for desired fit constraints. Trace 3-D then gave the optimal magnetic field gradients to be used for quadrupoles. An example of an optimized tune can be seen in Figure 2.12 where Trace 3-D has optimized the injection beamline for the four desired tune sub-sections noted earlier in section 2.1.2.

[13]The charge state, or charge number, of an ion is the total electric charge of the ion divided by the elementary charge e. A neutral atom that has a single electron removed thus has a charge state of 1+

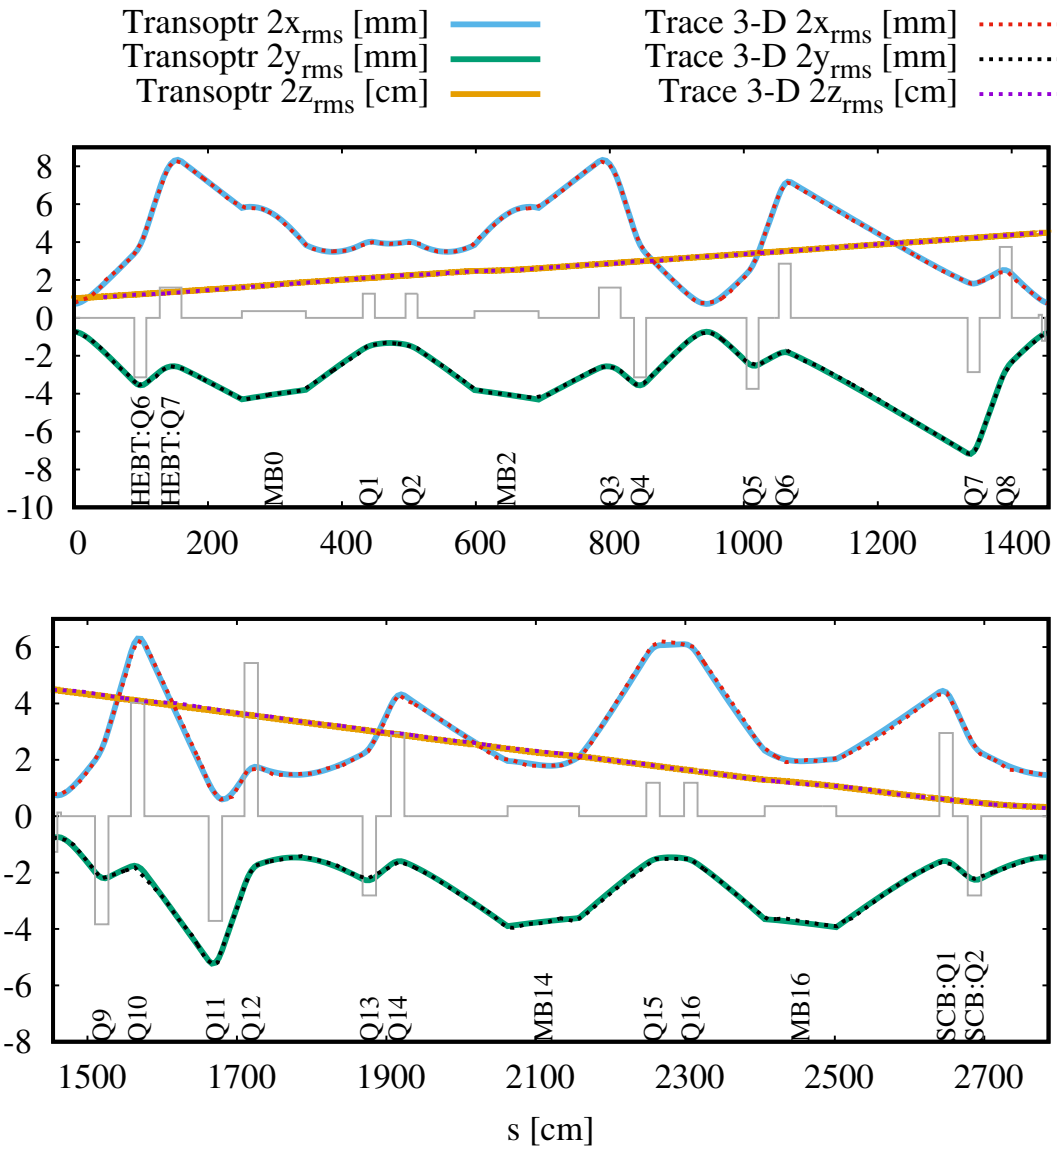

Figure 2.12: The 2rms envelope through the DSB section, comparing TRANSOPTR (solid lines) to Trace 3-D (dashed lines)

Ahead of an experiment, the field settings were scaled based on the A/Q and energy, and these were then converted into the necessary power supply current using a spreadsheet in the control room.

While Trace 3-D was a useful tool in linac operation, it did have some drawbacks. As suggested above, Trace 3-D was used only in an offline-mode, calculating desired tunes for given constraints well ahead of live beam delivery. Further, it modelled

each of the fourty accelerating cavities as an 'instantaneous kick', with the phase and amplitude of each cavity uncoupled from the physical machine.

These topics and the proposed improvements will be discussed further in Chapter 3 after completing the current description of the empirical linac optimization

2.5.3 Beam Envelope Simulations in TRANSOPTR

Also an envelope code that traces is roots to TRANSPORT, TRANSOPTR has been in use at TRIUMF since the 1980's[17] and is noteworthy in a number of ways:

- TRANSOPTR allows users access to the source code, enabling easy connection of model simulations to applications used for tuning.
- There is extensive expertise on the use of TRANSOPTR at TRIUMF, it has been used in the design of all low energy beamlines, as well as CANREB and the electron linac (e-linac).
- The code was expanded to allow for numerical integration through on-axis fields of accelerating cavities [18], enabling real-time simulations of accelerating cavities.

However, while TRANSOPTR has considerable benefits, it wasn't used in the design of the high-energy beamlines in ISAC and until recently hasn't been used for tuning applications of the high-energy beamlines. More to come on this in Chapter 3.

2.6 Empirical Linac Optimization

Herein is a description of the empirical method that has been the primary procedure for configuring the ISAC-II linac for a given requested output, built upon the concepts covered thus far.

2.6.1 Assumptions

With the empirical optimization method, the following assumptions are made:

1. The Courant–Snyder parameters of the beam arriving at the injection beamline are as desired, namely $\alpha_x = \alpha_y = 0$, $\beta_x = \beta_y = 0.2\,mm/mrad$, and $\epsilon_x = \epsilon_y \leq 5.0\,mm \cdot mrad$

2. There are no other unexpected correlations in the σ matrix, for example non-zero dispersion resulting in a non-zero σ_{16}

3. Hysteresis effects in the steel of quadrupoles are sufficiently small so as to be negligible in beam transport

4. Fringe field effects of the quadrupoles are negligible

While some or all of these assumption may not hold true, they are nonetheless assumed for the procedure that follows.

2.6.2 Preparation

All RF cavities in the linac are turned off, and the optics (quads, solenoids, dipoles) are loaded to the control system from a spreadsheet, scaled accordingly based on the Trace 3-D simulations performed offline.

2.6.3 Steering

Ideally, one would carefully move through the beamlines, ensuring steerers are adjusted to minimize the steering effects of optical elements. However this method is tedious and time consuming so in practice what happens is that beam transmission to the end of the ISAC-II linac is maximized using steerers and dipoles, while intermittently checking profile monitor to ensure that the beam remains near the reference trajectory. This is most reliably done immediately downstream of dipoles, where the dipole field is adjusted to centre the beam on the downstream profile monitor.

What this amounts to is basically an incomplete and error prone scanning of a large parameter space to maximize a scalar (the downstream beam intensity) that is a function of these many parameters.

2.6.4 Empirical matching

'Matching' in accelerator physics refers to utilizing a section of beamline (often 4 quadrupoles) to fit the sigma matrix at the output of that section to a desired output. Once this is accomplished the beam through the downstream section behaves just as expected, or at least behaves reproducibly to the last time the same fit was achieved.

Basically, one has a desired sigma matrix σ_f, a measured input matrix σ_i, and the goal is to find quadrupole settings that give rise to a transfer matrix M that produces σ_f from σ_i.

$$\boldsymbol{\sigma_f = M \cdot \sigma_i \cdot M^T} \tag{2.29}$$

This is most efficiently done by completing what is known as an emittance measurement at the initial location, which gives the inital sigma matrix σ_i. Then a computer is used to run a numerical optimization that finds the quadrupole settings that give a σ_f closest to that which is desired. In Trace 3-D, a regula falsi method is used to solve for the quadrupole settings, while other codes such as `TRANSOPTR` use different methods like the downhill simplex or simulated annealing.

In reality, the equipment required for emittance measurements is expensive and budgets have limits. No such diagnostic is available at the start of the injection beamline of the ISAC-II linac. Thus, an empirical match is done looking at one or both of two available diagnostics: profile measurements and beam transmission.

Profile measurements can be done at a number of points along the beamline, giving the $2X_{rms}$ and $2Y_{rms}$ size of the beam. An accelerator operator or physicist can adjust the quadrupole settings to achieve the desired sizes, however these measurements are just a projection of the 6-dimensional sigma matrix onto either the x or y axis, and don't provide information on the divergence of the beam.

More often used for empirically matching the tune is to provide four quadrupoles as variables, and attempt to maximize the transmission of the beam through a given section of beamline. For example the four quad matching section upstream of the ISAC-II injection beamline might be tuned to maximize the percentage of the beam that is transported *all the way to the end of the injection beamline*. This relies on the assumption that a mis-matched beam will exceed the expected envelope size somewhere along that beamline section and lose ions to the beam pipe or other apertures located along the beamline [14]. A more robust version of this is to manually tune four quadupoles through a pair matchings slits that define the beam size in both x and y. If these slits are placed at a phase advance of $\frac{n \cdot \pi}{2}$ where n is an integer, the result is that both the beam size and divergence is being defined at the two slit locations.

[14] This is basically the same manual scanning of a parameter space as done in the optimization of steerers

2.6.5 Time-of-Flight measurement

Following the empirical matching of quadrupoles at two distinct matching sections, beam is transported through the SCL to SEBT, the post-linac transport and diagnostic section. Here the three flight time monitors (FTMs) spaced along the SEBT section are used to measure the arrival time of the beam [19]. The time spectrum at each monitor is fit to a gaussian and the time between each of the three pairs of gaussian centroids (one for each monitor) gives three flight times t_{jk} over three different distances D_{jk}. These velocities give an energy per nucleon as shown in 2.30 with error 2.31. These are combined to give a weighted average beam energy per nucleon and error as shown in 2.32 and 2.33 [20].

$$E_{jk} = mc^2 \left[\frac{1}{\sqrt{1 - \left(\frac{D_{jk}}{c\Delta t_{jk}}\right)^2}} \right] \tag{2.30}$$

Where $j, k = 1, 2, 3$ and $j < k$

$$\epsilon_{jk} = \frac{\sqrt{\sigma_j^2 + \sigma_k^2}}{\Delta t_{jk}} E_{jk} \tag{2.31}$$

$$\bar{E} = \frac{\frac{E_{12}}{\epsilon_{12}^2} + \frac{E_{23}}{\epsilon_{23}^2} + \frac{E_{13}}{\epsilon_{13}^2}}{\frac{1}{\epsilon_{12}^2} + \frac{1}{\epsilon_{23}^2} + \frac{1}{\epsilon_{13}^2}} \tag{2.32}$$

$$\epsilon_{\bar{E}}^2 = \frac{1}{\frac{1}{\epsilon_{12}^2} + \frac{1}{\epsilon_{23}^2} + \frac{1}{\epsilon_{13}^2}} \tag{2.33}$$

Note however two mistakes in this approach. First, the profile at each monitor has a spread σ_j primarily because the beam itself is a collection of many particles that have some energy spread. The error in the energy measurement calculation should use the error of the mean of the arrival time at each monitor. A better estimate of the error of the mean is $\sigma_{\bar{t}} \approx \frac{\sigma_t}{\sqrt{n}}$. The typical number of data points in a single ftm profile collected in 10 seconds is 1e4, which would suggest the standard error of the mean is 1/100th of the standard deviation.

Second, Equation 2.31 is missing a factor of (approximately) two in relating the relative energy error to the time-of-flight error. In the non-relativistic limit this can be easily derived from the formula for kinetic energy and progagation of error to be:

$$\frac{\sigma_E}{E} \approx 2\frac{\sigma_t}{t} \tag{2.34}$$

Observations however do show disagreements between different pairs of FTMs that can't be explained from the error of the mean of the time profiles alone. There has to be some other sources of error contributing to this, and this will be discussed further in Section 5.2.

These calculations are typically done in a 'TOF Phasing' MATLAB program [21] that, when enabled, runs in a continuous loop to pull the FTM profiles, fit them, and calculate an average energy per nucleon and error. This program also helps assist the user with other operations in tuning as will be shown below.

2.6.6 DSB buncher

Once the ion beam has been transported to the end of the ISAC-II linac and the energy has been measured on the FTMs, setup of the DSB buncher begins. The DSB buncher is a room-temperature spiral 2-gap cavity operating at 35.36 MHz and its function is to take the beam that is diverging in time from the DTL and apply a corrective kick, slowing down the fastest ions and speeding up the slowest. The result is that all ions arrive at the first accelerating cavity of the linac with as small a spread in time as possible to allow for efficient acceleration.

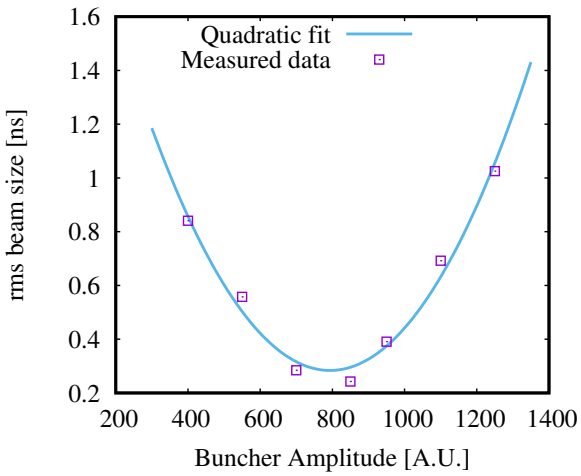

Figure 2.13: Typical quadratic fit of measured beam size vs buncher amplitude (focal strength) during optimization of the buncher

To optimize the buncher, the phase is scanned to re-produce the same energy at the FTMs as observed with the buncher off. Of course, two phases will accomplish this with one of the phases being de-bunching. The bunching phase is selected by ensuring the variance of the time distributions at the monitors is actually smaller than that observed with the buncher off.

Finally, a minimum of five data points are collected at various buncher amplitudes, recording the rms size of the time profile at each amplitude. A quadratic fit to these is

performed by the phasing program such that the amplitude producing the minimum beam size in time at the first FTM is identified, as shown in Figure 2.13 A basic optics equation is applied to shift this focus in time from the first FTM to farther upstream at the first accelerating cavity.

2.6.7 Phasing Cavities

With the setup of the DSB buncher complete the repetitive task of phasing cavities begins. With the amplitude of each cavity fixed by RF experts, each cavity is reduced to a single variable, the phase. The approach, briefly summarized, is roughly as follows:

1. Record the arrival time of the beam at the first FTM
2. Turn on the cavity in question
3. Scan the cavity phase to find the approximate accelerating phase
4. Record the arrival time of the beam at 5 different phases, typically at increments of 30 degrees
5. Plot ΔTOF vs ϕ_{cs} and fit to a cosine curve [15] using a downhill simplex method as show below in Figure 2.14
6. Select and set the -25 degree phase, which provides some longitudinal focussing while still harnessing more than 90% of the accelerating power of the cavity for acceleration
7. Measure the beam energy with the new cavity set.

[15] Note also that although the change in TOF is approximately proportional to the cosine of the phase under typical conditions, it isn't correct and a better fit can be achieved using the correct equation

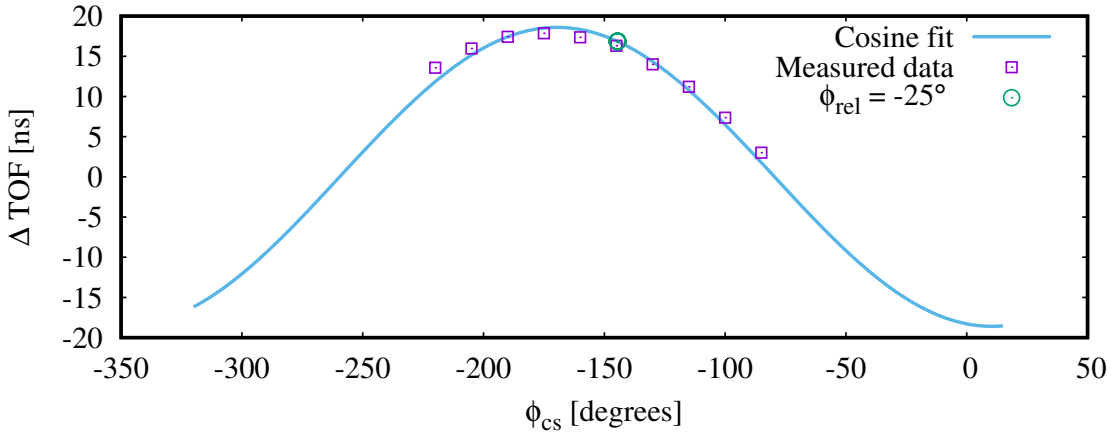

Figure 2.14: Typical data set (here with 10 points rather than 5) collected when phasing a cavity to find the desired -25° phase

These steps are repeated for each cavity, pausing to load and optimize the transverse optics as needed when beam intensity drops prohibitively low. The cavity phasing stops when the energy requested by the experiment has been exceeded. At this time the amplitude of the last cavity is reduced until the measured energy matches the requested.

2.6.8 Transverse Optics

The DSB injection beamline is set in such a way so as to ensure a circular beam just upstream of the first accelerating cavity. Setting the cavities at a longitudinally bunching phase producing a transverse RF defocusing effect. This paired with the length of the linac necessitate transverse focusing which is provided by superconducting solenoids. However, setting the strength of these solenoids is a suprisingly complex manner due to the many possible configurations of the cavities. Not only do the operational amplitudes of each cavity change on occasion - there are also cavity failures and not all cavities are typically used for every experiment. Even if the cavity amplitudes were all accurately known and phases fixed at -25 degrees, there are still 40! different unique configurations of the linac. Now in practice for the most part it is preferred to turn on cavities from the start of the linac until the desired energy is reached, leaving 40 configurations if all cavities are operational.

Offline, in Trace 3-D, the settings of the eight solenoids were optimised to bring the beam to a focal point at the diagnostic box between each cryomodule (M. Marchetto,

Pers. comm.). This was done for 41 different configurations: all cavities off, only 1st cavity on, etc. Then, the optimal solenoid field as a function of the boolean cavity on/of status was fit to find coefficients A_{ij} and D_j in the equation:

$$B_{ref,j} = \sum_{i=1}^{40} A_{ij} C_i + D_j \tag{2.35}$$

Where $j = 1..8$ for each of the 8 solenoids. and C_i is the cavity status as 0 or 1. During setup/tuning of the linac, this reference field $B_{ref,j}$ was calculated based on the on/off status of cavities and then scaled based on A/Q and energy compared to the initial Trace 3-D simulations and manually loaded into EPICS.

2.7 What's missing, and what more can be done?

While the ISAC-II linac has now been operated routinely for over a decade, it typically requires an operator with a significant amount of skill, training and experience [16] approximately eight to twelve hours to setup the linac with a new beam. Expert system owners have completed the task in four hours.

A number of opportunities for improvement have been hinted at and may have jumped out of the page at the reader already. It is worth here reflecting on the identified problems brought forth in Chapter 1.

1. *The SCL has not been modelled in a code that enables real-time simulations of the machine*

2. *The effect of the beamline elements providing transverse focusing (quadrupoles, solenoids) must be more carefully compared to the simulated beam enevelope*

3. *The parameters in the control system used to adjust the phase and amplitude of accelerating cavities have not been well calibrated to the physical properties of the cavities*

4. *An algorithm for sequential optimization of the linac has not yet been implemented, primarily due to the fact it has as prerequisites the above three points*

[16] While difficult to accurately quantify, a survey of the last four operators to tune the linac required on average a B.Sc. in Physics and a strong intuition for problem solving, 40+ hours of training time with an expert, and 120+ hours of hands-on experience tuning the machine before becoming fully proficient.

This chapter has now provided an introduction to the ISAC-II linac and beam optics, as well as a summary of the historical method which has been used over the past decade to operate the machine. With this behind us we move into Chapter 3 where I propose my improvment: to move to a model-based optimization of the linac utilizing as fully as possible the known physics behind the machine and more flexible simulation tools.

Chapter 3

Model-based Optimization of the ISAC-II Linac

To excercise sufficiently precise control in automated setup of the linac, we need a sufficiently accurate model of the machine paired with a simulation code that is flexible and executes quickly enough for real time optimization.

The first and easiest choice in selection of a simulation code is that of a many-particle or few-particle simulation. Multiparticle simulations typically trace individual ions through a map of electromagnetic fields and require long exectution times. In comparison, tracking the beam envelope as done in `TRANSOPTR` or even just a single reference particle is orders of magnitude quicker.

For optimization of the beam size both in the transverse and longitudinal directions, we have no choice but to use an envelope code which tracks the second moments of the beam distribution.

For calculation of cavity phases and energy gain, we are primarily looking at the first moment (mean) of the beam and do have another choice to make. Often used in accelerator physics is the analytic TTF descripion and drift-kick model of the energy gain introduced in Section 2.3, which applies an instantaneous kick at the centre of each cavity, as done in Trace 3-D. [1] The alternative is to use a numerical solver to integrate through the on-axis field, as done in `TRANSOPTR`.

Table 3.1 summarizes the two approaches for transversing the length of a single cavity from the starting point a, through point b at the centre of the cavity, to point c at the end of the cavity.

[1] While convenient for an analytic description of the energy gain, this is of course physically impossible, requiring an infinite force

	Integration at run-time approach	**Analytic approach**
Example code	`TRANSOPTR`	Trace 3-D
Electric field	$\mathcal{E}_z(s,t) = \mathcal{E}_z(s)cos(\omega t + \phi_{cs} - \theta)$	$\mathcal{E}_0$
Energy gain	$\Delta E = q\int_a^c \mathcal{E}_z(s,t)ds$	$\Delta E = q\mathcal{E}_0 TLcos(\phi_{rel})$
Transit time factor	N/A	$T = \frac{sin(\pi g/\beta\lambda)}{\pi g/\beta\lambda} sin(\frac{\pi\beta_s}{2\beta})$
Elapsed time	$\Delta t = \int_a^c \frac{1}{v(s)}ds$	$\Delta t = \frac{d_{ab}}{v_i} + \frac{d_{bc}}{v_f}$

Table 3.1: Comparison of two different approaches to tracking energy gain and elapsed time through a cavity

Where the average accelerating field $\mathcal{E}_0$ is calculated as shown in Equation 2.11. As `TRANSOPTR` is required for the calculation and fitting of the transverse beam envelope anyway, I will evaluate both the integration and the simpler drift-kick models for their expected accuracy in calculation of cavity phases. Figure 3.1 below shows a graphical representation of these two approaches. Note that because the relative velocity change in each accelerating gap is small, there is very little deviation between the time of flight between the two models

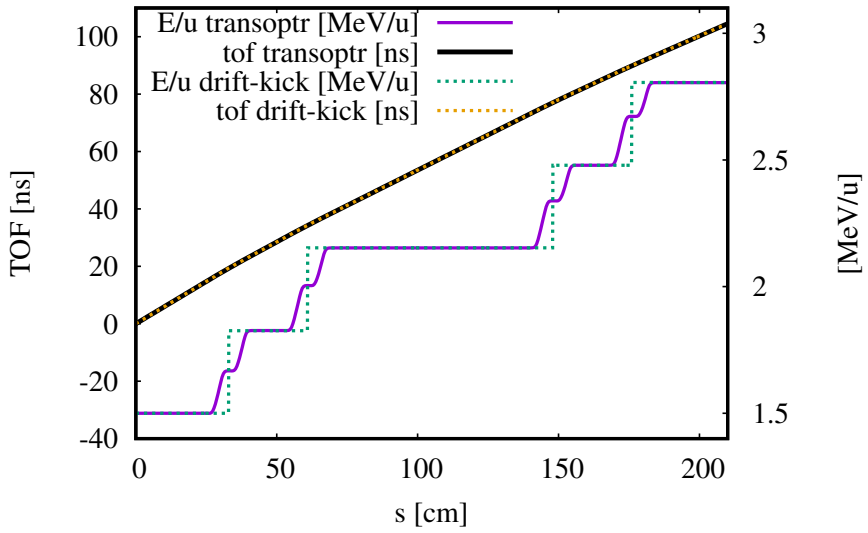

Figure 3.1: Comparison of the drift-kick and TRANSOPTR models for the first cryomodule in the ISAC-II linac

3.1 Modelling with TRANSOPTR

As an envelope code, TRANSOPTR tracks the second moments of the beam distribution and their correlations through the various optical elements of the beamline. Each particle in a beam distribution can be specified by a position vector $\boldsymbol{X}_i$ decribing its location in 6-dimensional phase space. The internal coordinates in TRANSOPTR are

$$\boldsymbol{X} = \left(x, \frac{P_x}{P_0}, y, \frac{P_y}{P_0}, \beta c \Delta t, \frac{\Delta E}{\beta c P_0} \right)^T \tag{3.1}$$

Where x and y are the transverse distances, P_x and P_y are the conjugate momenta, βc is the beam velocity, Δt is the arrival time of the particle at a given position s with respect to the reference particle, ΔE is the energy with respect to the reference particle, and P_0 is the total momentum of the reference particle. However, TRANSOPTR (and envelope codes in general) does not track each particle. Instead, the beam distribution is handled statistically. Internally, the beam distribution is represented within a 6x6 sigma matrix

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} & \sigma_{14} & \sigma_{15} & \sigma_{16} \\ \sigma_{21} & \sigma_{22} & \sigma_{23} & \sigma_{24} & \sigma_{25} & \sigma_{26} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} & \sigma_{34} & \sigma_{35} & \sigma_{36} \\ \sigma_{41} & \sigma_{42} & \sigma_{43} & \sigma_{44} & \sigma_{45} & \sigma_{46} \\ \sigma_{51} & \sigma_{52} & \sigma_{53} & \sigma_{54} & \sigma_{55} & \sigma_{56} \\ \sigma_{61} & \sigma_{62} & \sigma_{63} & \sigma_{64} & \sigma_{65} & \sigma_{66} \end{bmatrix} \tag{3.2}$$

which is the covariance matrix of the beam where the elements σ_{ij} along the diagonal are the central second moments (variance) of the beam distribution in each dimension, and the off-diagonal elements are the covariances between each pair of dimensions.

$$\sigma_{ij} = \langle x_i x_j \rangle = \frac{1}{n} \sum_{k=1}^{n} (X_{i_k} - \mu_{x_i})(X_{j_k} - \mu_{x_j}) \tag{3.3}$$

From a given initial distribution TRANSOPTR then runs in one of two primary modes. Either it tracks the $\boldsymbol{\sigma}$ matrix through the optical elements with simple matrix multiplication:

$$\boldsymbol{\sigma_2} = \boldsymbol{M} \boldsymbol{\sigma_1} \boldsymbol{M^T} \tag{3.4}$$

With the transfer matrix $\boldsymbol{M}$ for each optical element specified within TRANSOPTR. In the other mode, TRANSOPTR performs Runge-Kutte integration through the differential equation:

$$\frac{d\boldsymbol{\sigma}}{ds} = \boldsymbol{F}\boldsymbol{\sigma} + \boldsymbol{\sigma}\boldsymbol{F}^T \tag{3.5}$$

where $\boldsymbol{F}$ is the infinitesimal transfer matrix from s to $s + ds$

$$\boldsymbol{M} = \boldsymbol{I} + \boldsymbol{F}ds \tag{3.6}$$

The infinitesimal transfer matrix $\boldsymbol{F}$ is in turn calculated by extracting the equations of motion from the Hamiltonian [18]:

$$\boldsymbol{F} = \begin{bmatrix} \frac{\partial^2 H}{\partial P_x \partial x} & \frac{\partial^2 H}{\partial P_x^2} & \frac{\partial^2 H}{\partial P_x \partial y} & \frac{\partial^2 H}{\partial P_x \partial P_y} & \frac{\partial^2 H}{\partial P_x \partial z} & \frac{\partial^2 H}{\partial P_x \partial P_z} \\ -\frac{\partial^2 H}{\partial x^2} & -\frac{\partial^2 H}{\partial x \partial P_x} & -\frac{\partial^2 H}{\partial x \partial y} & -\frac{\partial^2 H}{\partial x \partial P_y} & -\frac{\partial^2 H}{\partial x \partial z} & -\frac{\partial^2 H}{\partial x \partial P_z} \\ \frac{\partial^2 H}{\partial P_y \partial x} & \frac{\partial^2 H}{\partial P_y \partial P_x} & \frac{\partial^2 H}{\partial P_y \partial y} & \frac{\partial^2 H}{\partial P_y^2} & \frac{\partial^2 H}{\partial P_y \partial z} & \frac{\partial^2 H}{\partial P_y \partial P_z} \\ -\frac{\partial^2 H}{\partial y \partial x} & -\frac{\partial^2 H}{\partial y \partial P_x} & -\frac{\partial^2 H}{\partial y^2} & -\frac{\partial^2 H}{\partial y \partial P_y} & -\frac{\partial^2 H}{\partial y \partial z} & -\frac{\partial^2 H}{\partial y \partial P_z} \\ \frac{\partial^2 H}{\partial P_z \partial x} & \frac{\partial^2 H}{\partial P_z \partial P_x} & \frac{\partial^2 H}{\partial P_z \partial y} & \frac{\partial^2 H}{\partial P_z \partial P_y} & \frac{\partial^2 H}{\partial P_z \partial z} & \frac{\partial^2 H}{\partial P_z^2} \\ -\frac{\partial^2 H}{\partial z \partial x} & -\frac{\partial^2 H}{\partial z \partial P_x} & -\frac{\partial^2 H}{\partial z \partial y} & -\frac{\partial^2 H}{\partial z \partial P_y} & -\frac{\partial^2 H}{\partial z^2} & -\frac{\partial^2 H}{\partial z \partial P_z} \end{bmatrix} \tag{3.7}$$

In this mode, TRANSOPTR also tracks the first moments of the beam distribution. This is primarily of use for tracking the beam energy and time of flight through the model, and is also done through Runge-Kutte integration:

$$\frac{dE}{ds} = q\mathcal{E}_z(s)cos(\omega t + \Phi) \tag{3.8}$$

$$\frac{dt}{ds} = \frac{1}{v(s)} \tag{3.9}$$

The input variable Φ that TRANSOPTR takes in is (in terms of variables already introduced) a combination of the variable phase shifter ϕ and the fixed phase offset θ, $\Phi = \phi - \theta$. The Hamiltonian and resulting $\boldsymbol{F}$ matrix for the axially symmetric linear accelerator have been derived by Baartman [18] and are used for the work in this thesis.

TRANSOPTR is fast, open source, and provides the functionality needed for linac beam-based model calibration and model-based beam tuning planned for this thesis. The next section now focuses on how to accurately and re-producibly calculate the

transfer matrices used by TRANSOPTR from the control system paramters (quadrupole current, cavity ampltiude and phase, etc.).

3.2 Model Calibration

The next step towards the new operational paradigm is a continued survey of the previously used Trace 3-D model, establish the new model in TRANSOPTR, and benchmark the two with each other. A number of physical properties representing the machine in the model are investigated within this section.

3.2.1 Quadrupoles - effective length, B-I, fringe field effects

Magnetic quadrupoles are in essence four electromagnets all fed by the same power supply. The poles and yoke of a magnetic quad are constructed out of steel, a ferromagnetic material. Around each of the poles is wound many turns of a coil which carries a DC current. The coil around each of the poles induces a very strong magnetic field at the tip of each pole due to the relatively high magnetic permeability of steel [2]

The 32 magnetic quadrupoles comprising the ISAC-II linac can be subdivided into five unique types, identified by their design and year of manufacturing.

- Danfysik 1982 L1
- Danfysik 1987 L1
- Danfysik 2003 L1
- Danfysik 1982 L2
- Danfysik 1987 L2

The manufacturer Danfysik delivers quads for which they specify the so called 'effective length.' This is the longitudinal extent L of an equivalent idealized 'square field' quad with the same peak pole tip field as the actual quad. A comparison of the two can be seen below in Figure 3.2. The effective lengths quoted by Danfysik for the

[2] It is energetically favourable for the magnetic moments of the iron atoms in the steel to align with the applied magnetic field.

L1 and L2 types are 180 and 325 mm respectively, and these were the lengths used in the historical Trace 3-D linac model.

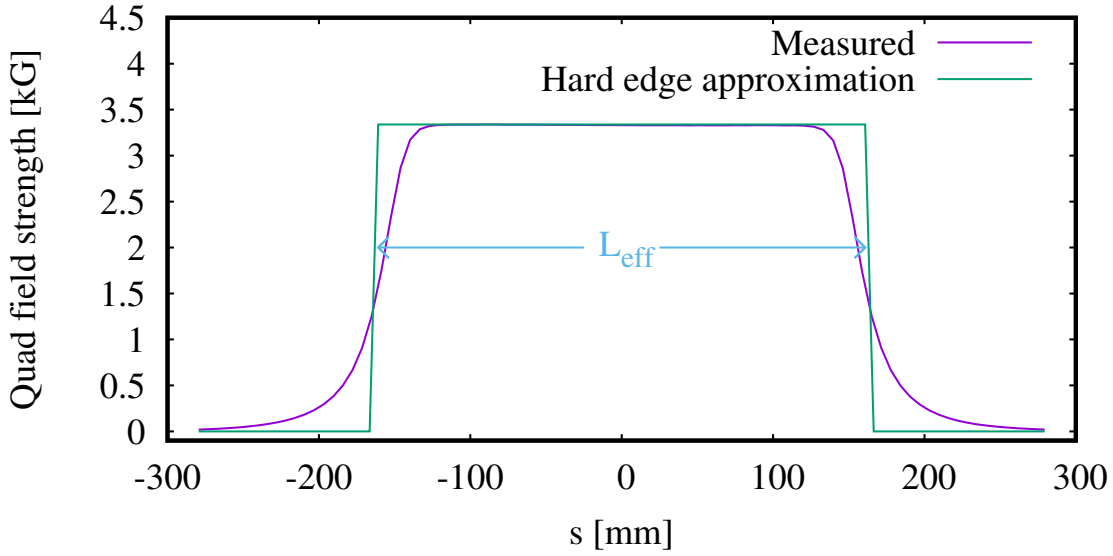

Figure 3.2: Square field approximation of a quad of effective length 325 mm compared with actual measured field

However, a number of variables can introduce errors in the construction of each quad. Most notably, the induced magnetic field is sensitive to the batch of steel from which it was constructed. This is why it is preferable to sub-divide quad types by manufacturing year in addition to their geometry. Every quad that is installed in the ISAC-II linac was measured in a test stand prior to installation during the years 2000 through 2005. These measurements typically consisted of:

1. Measurement of the induced magnetic field (in the longitudinal centre of the quad) vs the power supply current (commonly referred to as 'B-I curves')

2. Magnet tests at maximum current, checking resistance of the coil circuit, voltage, temperatures of cooling water, etc.

3. Magnet surveys, often referred to as a 'field map'. This involves recording the measured magnetic field components at a collection of points[3] within the poles and typically extends longitudinally past the edge of the poles until the field has fallen off to essentially 0.

[3]At TRIUMF these were most often 'cylindrical' scans done at a radius of 0.75" from the optical axis, and at angular steps of 15 degrees. A small selection of quads also had their radial linearity measured.

The Beam Physics group at TRIUMF processed the data from these measurements to calculate a B-I fit [22] as well as a measured effective length and fringe field coefficients [23]. However, not all of this information was included in the Trace 3-D model.

There is also the complicating issue of magnetic hysteresis. In short, after a quad is driven to a certain current and subsequently turned off, the steel in the quad will maintain some remnant magnetization. This in effect introduces an error into the magnetic field resulting from a given current, as the field depends on both the driving current *and* the recent history of the magnet. One way to mitigate this effect is to de-magnetize the steel by alternating the polarity of the current driving the magnet while slowly reducing the current amplitude until near zero is reached.

However, the power supplies for the quads in the ISAC-II linac are all uni-polar, meaning they can't be easily de-magnetized. Further to this, the initial field maps done appear to have been done starting from a de-magnetized state. What this means is the B-I relationship fit to this data and used for the model has both a systematic error and a pseudo-random error. This is depicted in Figures 3.3 and 3.4. In 3.3, the measured data points are shown both while ramping the magnet up and back down, and a fit of the data to a fifth order polynomial overlaid with the data. A linear term is subtracted off the field to better show the hysteresis effect.

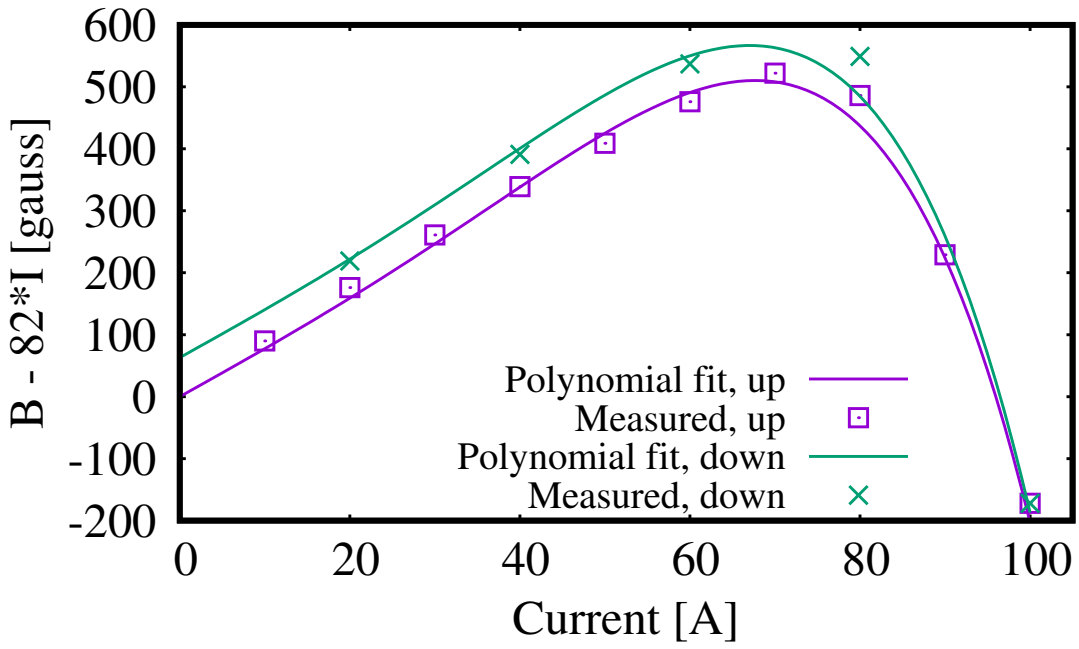

Figure 3.3: Measured B-I data compared to a 5th order (odd terms only) polynomial fit with a linear dependence of $82 \cdot I$ subtracted

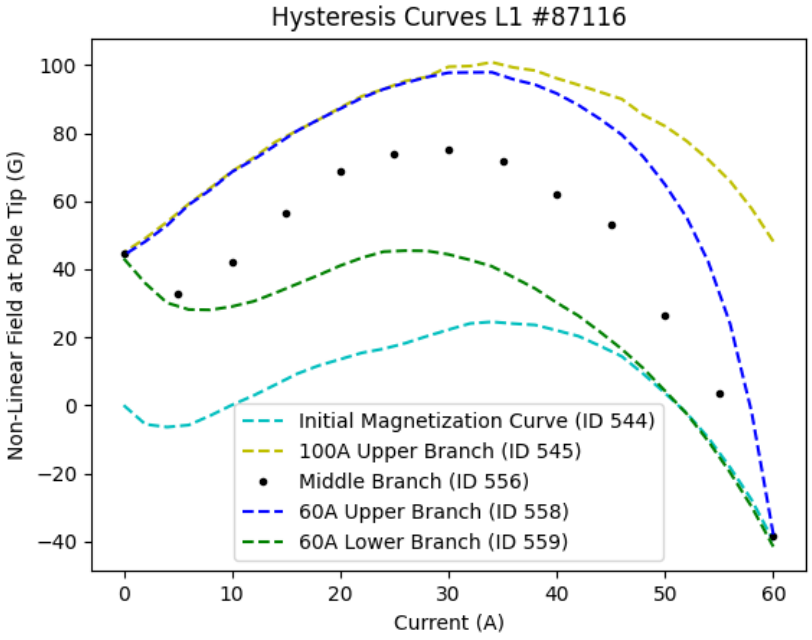

Figure 3.4: Outer hysteresis curves shown in light blue and yellow, uni-polar curve in green and dark blue. Figure courtesy of J. Nasser from the work done in [24].

In Figure 3.4, the initial magnetization curve up to 100 A are plotted, as well as the measured behaviour that arises from then installing the quad with a unipolar 60 A power supply, as typically done in ISAC. The B-I relationship used in the historical model would follow the light blue curve, however the actual behaviour online would always be within the two 60 A curves. The systematic error is the disagreement between the middle branch shown in black and the light blue initial magnetization curve, and the pseudo-random error is the vertical height between the two 60 A curves - depending on the recent history of the magnet, the actual field could be anywhere within those two curves. Figure 3.5 shows the relative significance of these systematic and random errors on the model accuracy.

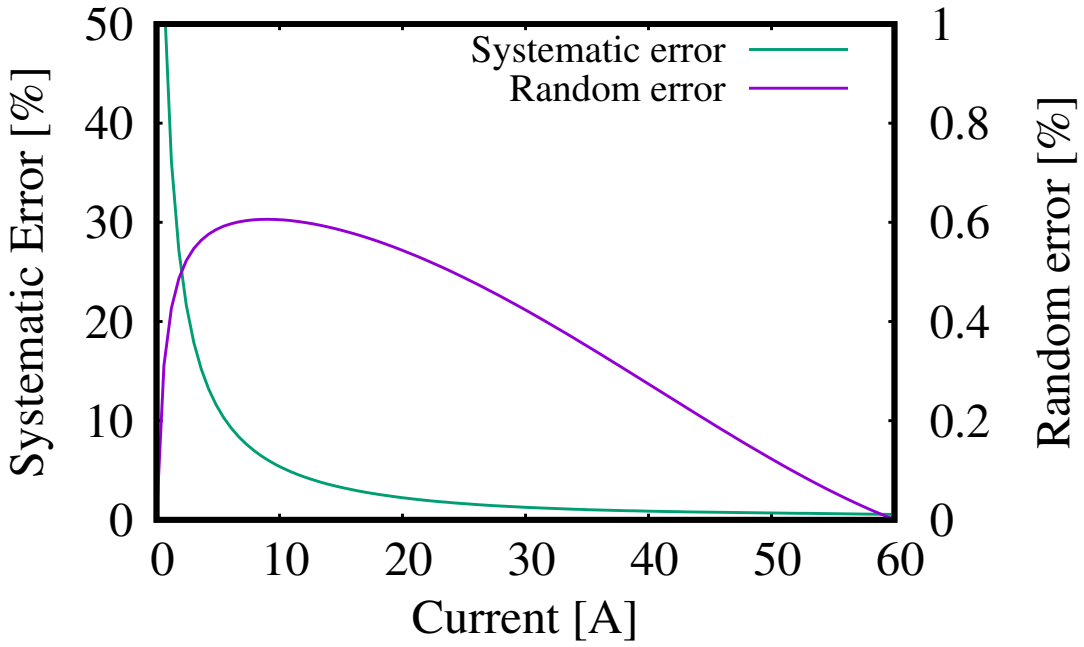

Figure 3.5: Relative disagreement between model and reality resulting from systematic and random errors in unipolar quads

Note that systematic error is larger than the random error for the entire operational range of the magnet. Further, in practice the quads were over-specified and most of them are operated in the current range of 5 to 20 amps, over which the systematic error is unacceptably large.

There is ongoing work [24] at TRIUMF to address both of these errors and attain a more precise and reproducible field to aid with model-based tuning, however for the purposes of this thesis, these sources of error are simply noted.

3.2.2 Superconducting Solenoids - Field maps, B-I

Historically, the Trace 3-D model of the ISAC-II linac represented the eight superconducting solenoids as 'hard-edged' solenoids with a constant axial magnetic field B_z and an effective length L_{eff}. In contrast, `TRANSOPTR` accepts the longitudinal magnetic field profile $B_z(s)$ and treats it similarly to an accelerating cavity in that the $\boldsymbol{F}$ matrix is calculated from the Hamiltonian, which depends on B_z profile. `TRANSOPTR` integrates through this with adaptive Runge-Kutte integration.

The B_z field maps were originally measured in 2004 and 2005 by the vendor, Accel, that produced the solenoids for TRIUMF. These field maps were normalized to a maximum amplitude of 1, and used as the `TRANSOPTR` input field maps. The normalized maps are shown below in Figure 3.6.

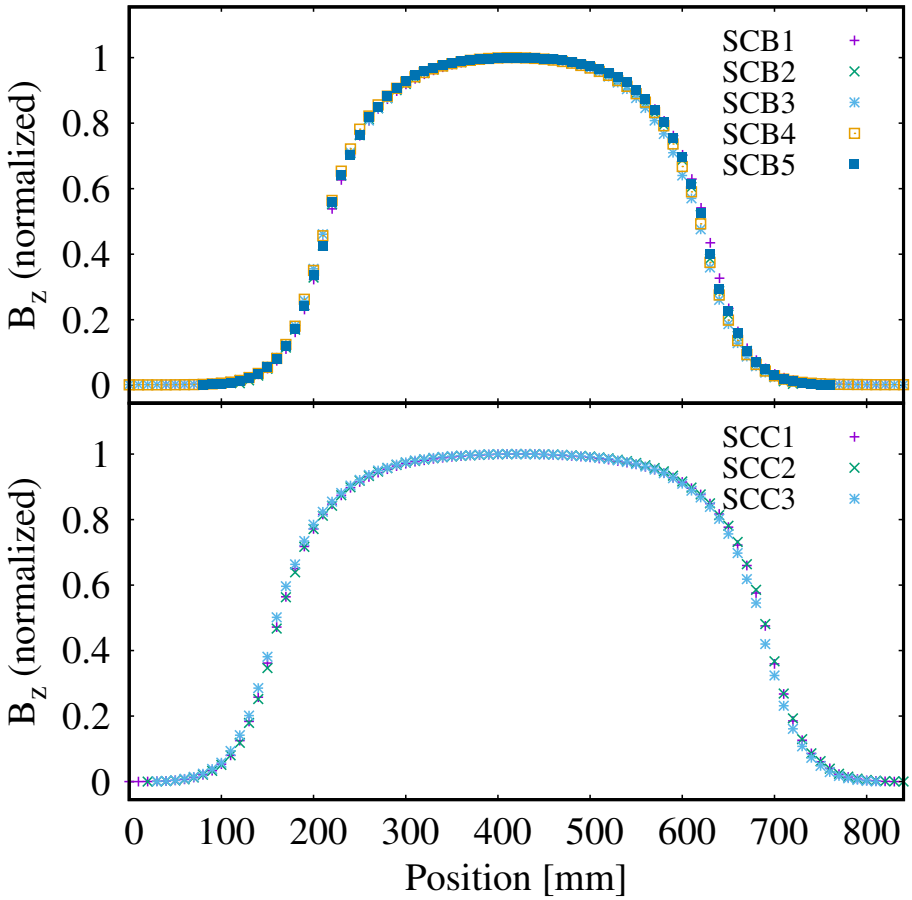

Figure 3.6: Normalized B_z field maps of the eight SC solenoids, data collected by the manufacturer

The B-I relationship between the supply driving current to the solenoids and the

maximum field at the solenoid centre is also vital. Previously with the Trace 3-D model, an average linear relationship between the current and field was used, with one fit for the shorter solenoid type (SCB) and one fit for the longer solenoid type (SCC). This has been replaced with the individual B-I relationship for each of the eight solenoids supplied by the vendor.

It should be noted that while the relationship is expected to be very linear, there is only one B-I data point provided for each solenoid, the current required for a 9 Tesla field at the solenoid centre. A future improvement could be to collect a full B-I data set from 0 to 100 A for each solenoid.

3.2.3 Cavity Field maps

Using CST Microwave Studios, 3-D simulations of the time-variant electromagnetic fields within the three superconducting cavity types have previously been completed by V. Zvyagintsev. The on-axis field map $\mathcal{E}_z(s)$ is extracted from this and provided to `TRANSOPTR` as the field input[4]. The normalized field maps are shown below in Figure 3.7. One can see that as the design velocity, β, of the cavity increases, the region with vanishingly small electric field at the centre of the cavity (within the drift tube) gets longer to allow enough time for the RF cycle to still complete a half-period prior to the beam entering the second gap.

[4]It should be noted here that the established subroutine in `TRANSOPTR` is for an axially symmetric field such as that in a multi-cell elliptical cavity. The error associated with modelling a cavity that isn't axially symmetric with this routine has not yet been investigated for the ISAC-II linac

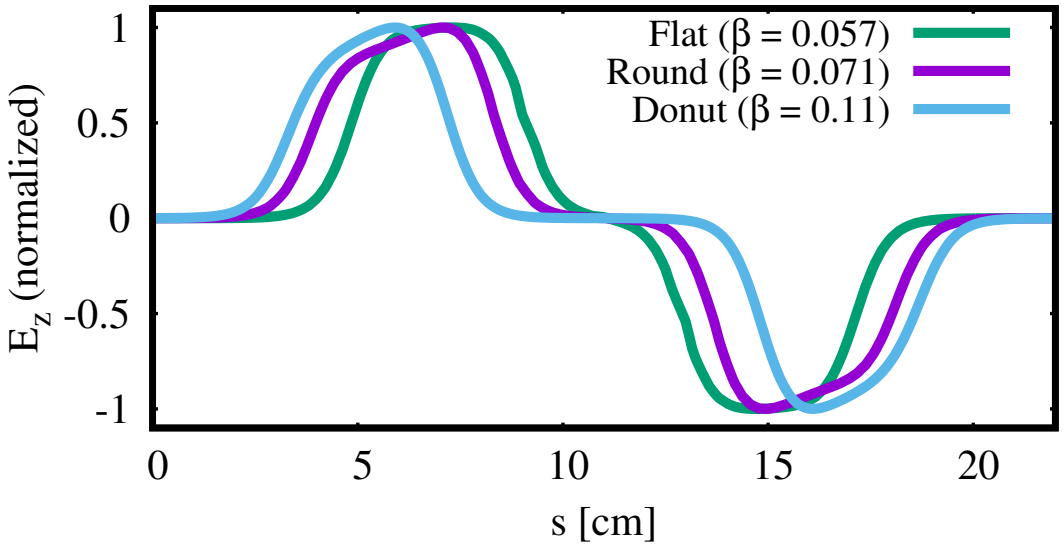

Figure 3.7: Normalized E_z field maps of the three cavity types in the ISAC-II linac

CST Microwave Studios is widely-used commercial software that uses the finite element method (FEM) to solve Maxwell's equations for given cavity geometries and initial conditions.

This differs than the Trace 3-D approach which applied an instaneous kick to the beam to increase the energy. This was basically equivalent to a E_z that was a delta function with a finite integral:

$$\frac{dE}{ds} = \delta(s) q \mathcal{E}_0 T L cos(\phi_{rel}) \tag{3.10}$$

and

$$\Delta E = q \mathcal{E}_0 T L cos(\phi_{rel}) \tag{3.11}$$

Regarding the calibration of the cavity field maps, the last step which is properly calibrating the amplitude of the field map, is done with the beam itself and is covered later in Section 3.2.6.

3.2.4 RF phase shifter

The RF cavities providing acceleration to the beam must be timed correctly. The beam must be bunched and must be crossing the acceleration gaps at a time when an accelerating electric field is established, which is pointing along the direction of travel. Further, the timing of the cavities may need to be adjusted as beam properties

change.

The RF cavities in the ISAC accelerator chain are all driven from the same single master oscillator, and at various points along the distribution of this single reference signal, the frequency may be multiplied to a higher order harmonic, or shifted by a variable amount, known as a 'phase shifter'. The phase shifter is the adjustable setting ϕ_i from Equation 2.19. A more detailed discussiong of the ISAC-II RF distribution will follow later in Section 3.2.6.

Based on past discussions with beam physicists and low-level RF (LLRF) experts and the nature of the circuits that convert a digital signal (setting in the control system) into an analog reference signal driving each cavity, it was known that the behaviour is not perfectly linear, nor well investigated.

There are fourty four phase shifters in the ISAC-II linac, and the effect of these was calibrated using a Hewlett Packard 8508A Vector Voltmeter. The calibration procedure was as follows:

1. Connect the vector voltmeter to two signals:
   - Reference signal arriving at the local RF console
   - Signal coming out of phase shifter board for each cavity
2. Setup the voltmeter to output the phase difference between the two signals
3. Step the EPICS/LLRF setpoint from -220 to +220 in steps of 10, recording the voltmeter measurement at each step.
4. Repeat for each of the 44 phase shifters in the ISAC-II linac.
5. Implement script to convert between control system setpoint and actual degrees in model.

To clarify, the ϕ_i from equation 2.19 is really a 'corrected' control system phase, which is some non-linear function of the control system phase ψ_i, $\phi_i = \phi_i(\psi_i)$. To keep the discussion going forward clear, I will primarily mention only the (corrected) control system phase ϕ_i, and will not further dwell on the detail that this must be converted to a (uncorrected) control system phase ψ_i before actually being delivered to the control system. The data from these calibrations are presented and discussed in Section 4.2.1.

3.2.5 RF Amplitudes & Power Meter

When running in closed-loop for beam delivery, the LLRF system locks cavities onto a desired amplitude. An RF pickup in the magnetic field region of the cavity picks up the RF signal and it is sent through a power meter and ADC, which converts the signal to a digital value in the control system. However, the reading of the power meters used only scale linearly with the cavity voltage over a certain limited range in power. This effect has been measured by the SCRF group at TRIUMF (Z. Yao, Pers. comm.) and is shown below in Figure 3.8. The relationship between the accelerating field $\mathcal{E}_a$ and the control system amplitude A_{cs} is shown for three different cavities. As can be seen, the scaling parameter C_i of the linear behaviour between $\mathcal{E}_a$ and A_{cs} is different for each of the three cavities. Further, it is observed that CAV2 and CAV4 deviate from the linear relationship at amplitudes above ~ 1800.

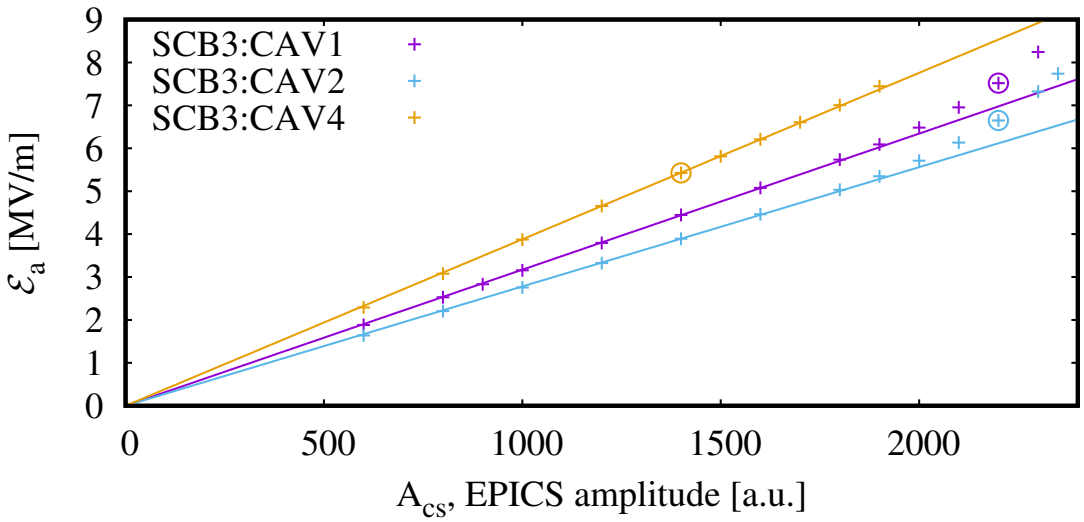

Figure 3.8: Accelerating field vs amplitude setpoint, operational setpoints are circled. Data courtesy of Z. Yao

For the purposes of this thesis, the control system amplitudes have been assumed to have linear relationships with the electric fields, with some analysis of the according systematic error given later in Section 5.2. For future refinement of the model, two options are available:

1. Add or adjust attenuators to the RF cables upstream of the power meter to decrease the signal strength such that it is completely within the linear range.

To continue to make use of past beam measurements for calibration this will require a conversion to past control system data.

2. Include the non-linear relationship between $\mathcal{E}_a$ and A_{cs} within the model and automated routine, similar to how the phase shifter non-linearity is handled. Either fit the data from the SCRF group's measurements or do analagous beam-based measurements at varying amplitudes.

3.2.6 Beam-based Model Calibration

The methods described for calibrating the `TRANSOPTR` model to the linac thus far have required measurements independent of the beam. The final step of calibration requires input using the beam itself.

Ideally, all cavities would be 'in-sync' if no phase shift is applied to any of them. However, the length of cables carrying the LLRF reference signal from the master oscillator to the various subsystems, finally to the power amplifiers and the cavities themselves all vary. So in reality there is some constant delay offset θ_i between each different cavity in addition to the variable phase shift ϕ_i, as depicted below in Figure 3.9.

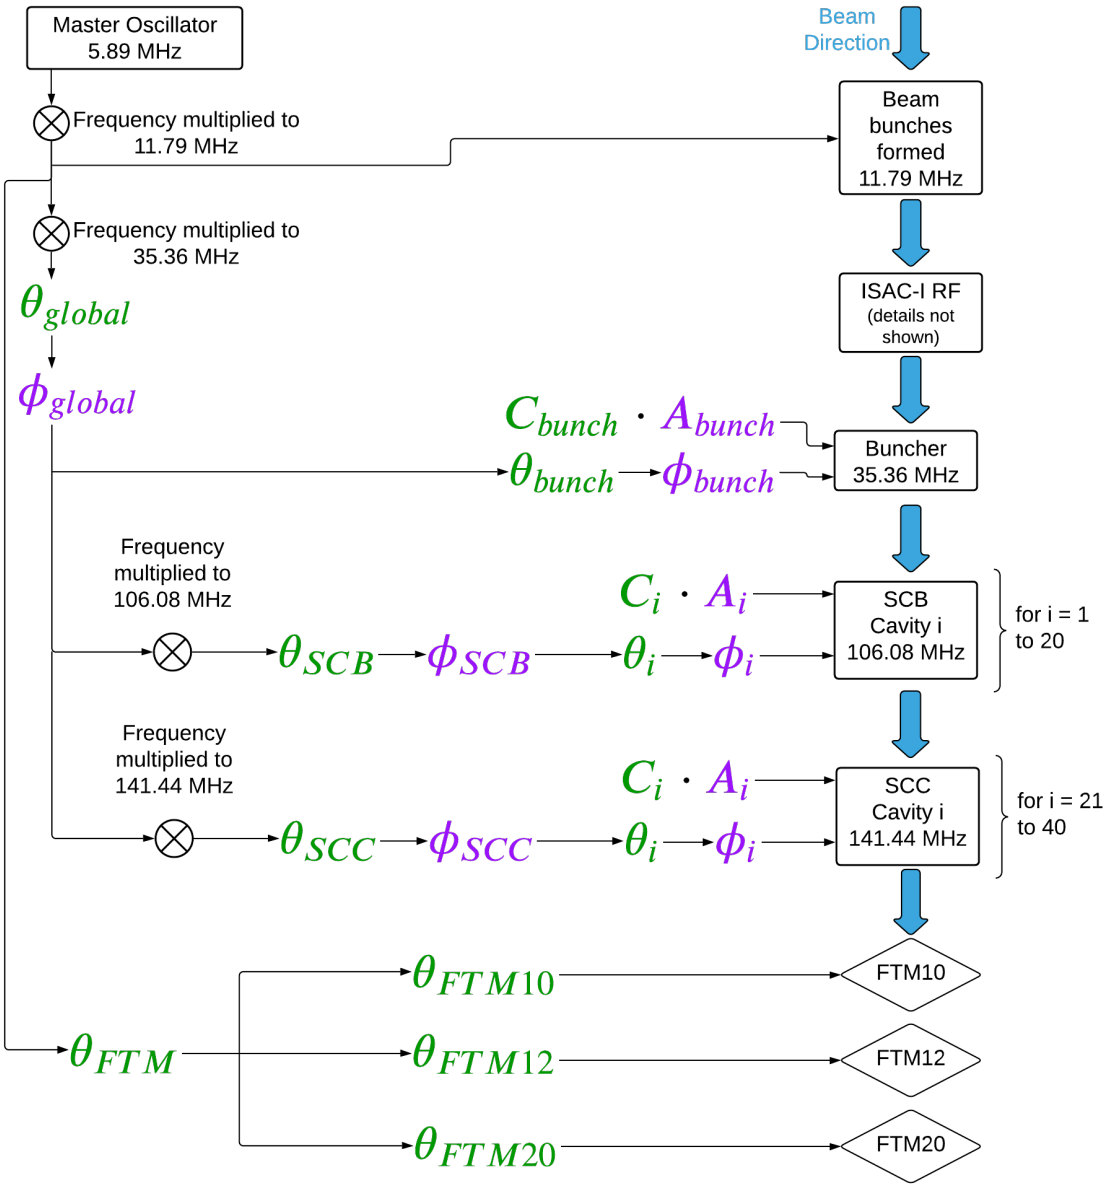

Figure 3.9: ISAC-II RF Distribution

In this figure, unknown static parameters are written in *green*, while variable control system parameters are written in *purple*. In addition to the phase offset and cavity amplitude calibration for each of the 40 cavities, there are also two unknown parameters for the upstream buncher, θ_{bunch} and C_{bunch}, as well as an unknown offset θ and variable control system phase ϕ for each of the three global phases: 'Global', 'SCB', and 'SCC'. These global phase shifters allow adjustments of all downstream signals by the same amount, which will be leveraged later.

In the actual calibration and running of the model, redundancy is removed by setting the reference point of the ISAC-II linac model *downstream* of the global phase, and by setting:

$$\theta_{bunch} = \theta_{SCB} = \theta_{SCC} = 0 \tag{3.12}$$

The proposed model-based optimization takes in a collection of data sets, each a snapshot of the linac tune for a given beam. Basically, for each cavity along the chain we have an output energy E_i that is a function of the control system variables A_i and ϕ_i, and depends on the known tune parameters mass (a), charge (q), input energy (W, to avoid confusion with the output energy E), and flight time (t) as well as the unknown parameters C and θ. We want to fit those unknown parameters to the measured data. This method is independent of the model (the model is just the function that gives the output energy E_i). So it can be applied both to the `TRANSOPTR` and drift-kick models, as well as any other model that provides the same inputs and outputs. For calibration of the `TRANSOPTR` model, the equation used is:

$$\Delta E = q \int A_i C_i \mathcal{E}_{norm}(s) cos(\omega t + \phi_i - \theta_i) ds \tag{3.13}$$

While for the drift-kick model, Wangler's analytic Equation 2.17 mentioned earlier can be re-written in terms of the control system phase and amplitude for a clearer comparison.

$$\Delta E = q A_i C_i T cos(\omega t + \phi_i - \theta_i) \int |\mathcal{E}_{norm}(s)| ds \tag{3.14}$$

Where the integral of the normalized field map needs to be evaluated only once for each cavity type.

The proposed data collection for calibration is as follows:

1. Begin with all RF cavities in ISAC-II linac off.
2. Measure the beam energy/nucleon on the FTMs.
3. Turn on the next available cavity in the linac.
4. Vary the cavity phase, measuring the output energy/nucleon on FTMs at each step.
5. Set the desired operating cavity phase using the legacy MATLAB phasing program (the method used to determine the operational phase isn't important)
6. Repeat steps 2-5 for each cavity.

The proposed algorithm for calibration of the model is as follows:

1. Calculate the energy per nucleon W_k and time-of-flight t_k arriving at the entrance of the next cavity
2. For each data point (ϕ_i, A_i) collected for the cavity, calculate the measured output energy/nucleon E_i
3. Use a non-linear least square algorithm to solve for the unknown parameters C and θ to best fit the measured data
4. Input the new parameters C and θ into the model
5. Repeat steps 1-4 for each cavity in the linac, skipping any cavity for which no data points are available.

Graphically, this algorithm is shown below in Figure 3.10

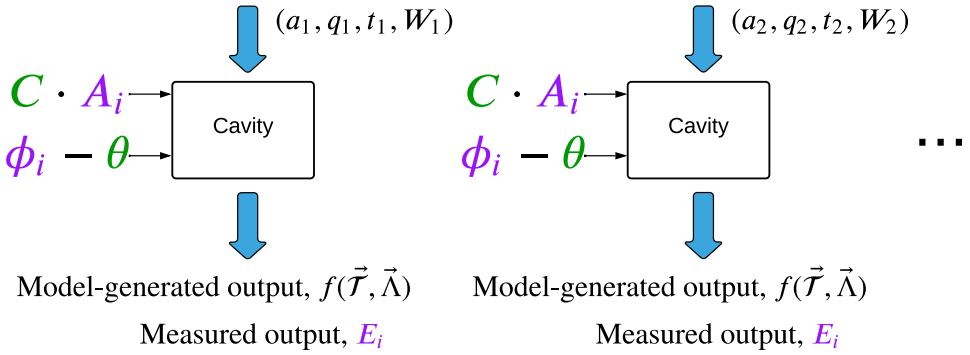

Figure 3.10: Cavity calibration algorithm

Mathematically, we have a tune with properties (a, q, t, W), where a is the mass of the beam in amu, q is the beam charge state, t is the elapsed flight time entering the next cavity, and W is the energy entering the next cavity. At each cavity we collect

$$n \text{ data points } A_i, \phi_i, E_i, \text{ where } i = 1, ..., m$$

Where A_i is the control system amplitude, ϕ_i is the control system phase, and E_i is the measured output energy per nucleon at that phase/amplitude setting. So we end up with n data points for the fitting of the unkown parameters C, θ at each cavity. The non-linear least squares algorithm is then minimizing:

$$S = \sum_{k=1}^{n} (E_i - f(\vec{\mathcal{T}}, \vec{\Lambda}))^2 \tag{3.15}$$

where the data points are

$$\vec{\mathcal{T}} = (a, q, t_i, W_i, A_i, \phi_i) \tag{3.16}$$

and the unknown parameters are

$$\vec{\Lambda} = (C, \theta) \tag{3.17}$$

The function $f(\vec{\mathcal{T}}, \vec{\Lambda})$ that the data is being fit to is the energy per nucleon coming out of the cavity calculated by the model. So in the `TRANSOPTR` model case, it is numerically calculated by Runge-Kutte integration through the on-axis field (Equation 3.13), while in the drift-kick case it is the result of the analytic expression in equation 3.14.

3.3 Implementation of Automated Optimization

With the method for calibrating the model to the machine in hand, we can now turn towards the algorithm for automatic optimization. At its core, it is just the inverse process of the calibration described in the previous section: having completed beam-based model calibration (tuning), we can now have confidence in carrying out model-based beam tuning.

We have some known beam properties arriving at the linac: (a, q, t, W), we now know the calibration parameters for all linac cavities (C, θ), and we use the model to vary the cavity phases ϕ to achieve some desired output energy per nucleon E_{final}.

Measuring beam properties at linac entrance

While not a main feature of this thesis work, a desirable long term approach to tuning of the ISAC-II linac is to also automatically match in the transverse tune. As described in section 2.6.4, this is best done with an emittance measurement and re-matching to the desired beam properties at the linac entrance.

As an emittance meter is not currently available upstream of the linac, it is currently proposed to utilize emittance re-construction, which employs the scanning of a quadrupole to rotate the beam in phase space at a downstream profile monitor. By measuring the beam size at this profile monitor for many different *known* transport matrices between the quad and profile monitor, the sigma matrix descibing the beam's distribution in phase space upstream of the quad can be re-constructed to best fit the measured data.

Model Inputs

The user of the automated optimization routine must specify only the beam mass, charge, incoming beam energy, and requested output energy.

3.3.1 The automated routine

The first step of the automated routine is to load the user specified beam properties into the model: mass, charge, and input energy (a, q, E).

Re-zeroing the clock

Next, we run into an interesting hurdle: there is no guaruntee that the beam is arriving at the start of the ISAC-II model at the same time as it was in the calibration tunes, and we do not have a time-domain diagnostic upstream of the DSB buncher with which to measure the arrival time. Instead, the proposal is to use the bunching cavity itself as the time-domain diagnostic:

1. The bunching cavity is optimized by the user to bring the beam to a focus in time at the first accelerating cavity, as historically has been done and described back in section 2.6.6, yielding some buncher phase $\phi_{empirical}$.

2. Next, the model is used to calculate the expected buncher phase ϕ_{model} and due to the unknown arrival time at the start of the model, there will likely be some disagreement, namely $\phi_{empirical} \neq \phi_{model}$

3. The global phase shifter (ϕ_{global} in Figure 3.9), which adjusts the timing of all cavities can then be adjusted to be the difference $\phi_{empirical} - \phi_{model}$ such that the buncher can be set to the predicted model value.

$$\begin{aligned} \phi_{bunch} &= \phi_{empirical} \\ \phi_{global} &= 0 \end{aligned} \tag{3.18}$$

becomes

$$\begin{aligned} \phi_{bunch} &= \phi_{model} \\ \phi_{global} &= \phi_{empirical} - \phi_{model} \end{aligned} \tag{3.19}$$

So that timing of the buncher is identical to its previous setting, but because the global phase has been shifted we essentially have re-synced all the cavities to be at the expected starting phase in their RF cycles at the start of the model.

It is expected that this re-zeroing will also be essential at the first accelerating cavity. This is due to the significant drift space between the bunching cavity and the first accelerating cavity. This large distance means that a small error in the measured energy can result in a relatively large phase offset at the first accelerating cavity.

$$\Delta\phi \approx \frac{360}{T_{RF}} \frac{(\Delta\beta)d}{\beta c} \tag{3.20}$$

Where T_{RF} is the RF period, β is the beam velocity, d is the distance between the start of the model and the cavity, $\Delta\beta$ is the velocity error as a percentage, and $\Delta\phi$ is the resulting phase offset in degrees. For the first linac cavity (SCB1:CAV1), we get:

$$\Delta\phi \approx \frac{360}{9.43ns} \frac{(\Delta\beta)13.9m}{0.057c} \approx 309\Delta\beta \tag{3.21}$$

So a 1% shift in velocity gives a shift in phase of over 300 degrees. To ensure the phase is accurate to better than 5 degrees, the velocity would need to be known to much better than a tenth of a percent, or an energy/nucleon to 0.5 keV/u or better, out of a total of 1.53 MeV/u. This exceeds the accuracy to which we can measure the energy, so SCB1:CAV1 is used to reset the model clock as well, in an analagous way to the DSB buncher:

1. SCB1:CAV1 is manually phased by collecting TOF vs phase data points, as historically has been done and described back in section 2.6.7, yielding some cavity phase $\phi_{empirical}$.

2. Next, the model is used to calculate the expected buncher phase ϕ_{model} and due to the unknown arrival time at the start of the model, there will likely be some disagreement, namely $\phi_{empirical} \neq \phi_{model}$

3. The SCB and SCC phase shifters (ϕ_{SCB}, ϕ_{SCC} in Figure 3.9), which adjust the timing of all SCB and SCC cavities can then be adjusted to be the difference $\phi_{empirical} - \phi_{model}$ such that SCB1:CAV1 can be set to the predicted model value.

$$\begin{aligned} \phi_1 &= \phi_{empirical} \\ \phi_{SCB} = \phi_{SCC} &= 0 \end{aligned} \tag{3.22}$$

becomes

$$\begin{aligned} \phi_1 &= \phi_{model} \\ \phi_{SCB} &= \phi_{empirical} - \phi_{model} \\ \phi_{SCC} &= \frac{4}{3}(\phi_{empirical} - \phi_{model}) \end{aligned} \tag{3.23}$$

Note that the factor of $\frac{4}{3}$ comes from the difference in RF frequencies between the SCB (106.08 MHz) and SCC (141.44 MHz) cavities.

The main algorithm

1. Read the amplitude of the next cavity from the control system, input into the model.
2. Run the model, allowing it to vary the cavity phase to maximize the output energy[5], thus finding the 0-degree phase of the cavity.
3. Shift the phase by some amount (typically -25 degrees) to achieve the desired bunching effect.
4. (Optional) If intermittent checks of linac transmission and energy are desired, run the tranverse tune algorithm (below).
5. Check if the model output energy has exceeded the requested energy. If it has, break out of the main algorithm to the final energy optimization algorithm. If not, repeat steps 1-4 for the next availabile cavity in the linac.

Tranverse tune algorithm

1. With current state of the modelled cavity amplitudes and phases, run the model, allowing it to vary the focal strength of solenoid 1 to match the desired beam waist at the location of the next diagnostics box.
2. Fix solenoid 1 in the model to the optimized strength identified in the step above
3. Repeat steps 1 & 2 for all eight solenoids in the linac.

[5]As mentioned in an earlier chapter, TRANSOPTR uses either a downhill simplex or simulated annealing algorithm to solve the least squares problem

4. Based on the assumption that the desired match is achieved at the linac exit, scale the downstream transport quads from the design tune to the new beam and energy/nucleon.

Final energy optimization algorithm

1. With current state of the modelled cavity amplitudes and phases, run the model, allowing it to vary the amplitude of the last cavity in use to match the request output beam energy of the linac.

2. Carry out the transverse tune algorithm to find the optimal settings for focusing devices.

3. Load model setpoints to the control system.

4. Measure the energy per nucleon and print out summary for user, indicate if expert intervention is required.

Algorithm for Recovery from a Cavity Failure

During delivery to experiments, a status of the linac is routinely saved during initial setup. This can be used along with the model to re-calculate linac settings regardless of whether or not the delivered linac tune was established via the automated linac algorithm.

1. A daemon running constantly monitors the status of all cavities and compares to the last tune saved for delivery.

2. If a change is detected that lasts for an extended period of time, the algorithm to calculate new settings begins automatically, but waits for operator intervention before setting to the control system.

3. Starting from the start of the linac, read the amplitude of the next cavity from the control system, input into the model.

4. Run the model, allowing it to vary the cavity phase to maximize the output energy, thus finding the 0-degree phase of the cavity.

5. Shift the phase by the desired relative phase saved with the saved initial linac tune data.

6. Continuing repeating steps 3-5 each subsequent cavity, until the failed cavity is reached.

7. It is assumed at this point in the ISAC-II linac that the model and actual energy are in agreement, but an accomodation is made to allow for the possibility that cavity phases were manually tuned.

8. If modelled phases up to the failed cavity agree within some threshold to those used in the saved tune, continue the main linac tuning algorithm (omitting the failed cavity)

9. If modelled phases disagree beyond the threshold, calculate a 'model-machine' phase disagreement, and add this phase to downstream cavities.

10. When the model output energy has exceeded the requested energy, break out of the main algorithm to the fine-tuning algorithm.

3.4 What is new about this approach?

The tuning of a linac is not new - however some of the details of the implementation described in this chapter warrant highlighting:

- The linac model is calibrated to beam-based measurements of phase-energy data collected for each available cavity.

- The model is not restricted to the analytic drift-kick model, but can via `TRANSOPTR` do real-time numerical integration through the cavity fields for more accurate results, avoiding the use of a transit time factor.

- The envelope code `TRANSOPTR` is used to read live cavity status from the control system and re-optimize the transverse optics for the desired fit in real-time.

This new method for optimization of the linac builds upon previous developments both within TRIUMF and externally. This includes literature for implementations of automated phasing and recovery at other labs [25], the establishment of a high level application (HLA) framework at TRIUMF [2], and investigations into the properties of optical elements [24]. A more detailed discussion of analagous implementations at other labs is given in Appendix A.

This chapter has presented the proposed new approach for the calibration of the linac model and the subsequent use of that model for automated phasing of linac cavities. In summary, this chapter has:

1. Reviewed and updated the model to match reality as close as possible first *without* use of the ion beam.

2. Described the method for using the ion beam and the linac to calibrate the phase offset θ and amplitude scaling parameter C for each of the fourty cavities.

3. Described the algorithms to be used for model-based tuning of the ISAC-II linac.

The algorithms described enable a highly desirable functionality that is important for regular operation of the linac - the ability to automatically rephase cavities in the event of a cavity failure.

Equipped with this new approach, the next chapter will describe in more detail the experimental setup, and present the collected data both for the calibration and automatic phasing experiment.

Chapter 4

Calibration and Experimental Data

This chapter begins with a summary of the experimental setup, followed by presenting data from the calibration of the linac model, and finally the results from tests of the automated phasing routine.

4.1 Experimental Setup

All data collection and experiments were performed at the ISAC-II facility at TRIUMF [26][27] using the existing diagnostics and measurement capabilities of the facility. A short summary is provided below.

4.1.1 Diagnostics

Beam Intensity Measurements

A Faraday Cup (FC) is essentially a cup shaped block of copper[1] that can be either inserted into the beam axis to both fully stop the beam, or retracted to a position fully out the beam axis. A FC is oriented such that the opening of the cup is facing towards the incoming ion beam. A wire grid biased at approximately +100 volts is located across the cup opening, to supress secondary electrons that are emitted when high-energy ions impinge on the copper (or tantalum) surface. This supression of secondary electrons is essential, as a FC is used to measure the beam current at this location.

[1]Faraday cups in the ISAC-II linac and downstream are actually copper lined with Tantalum, to raise the Coloumb barrier and reduce undesired nuclear reactions at these higher energies.

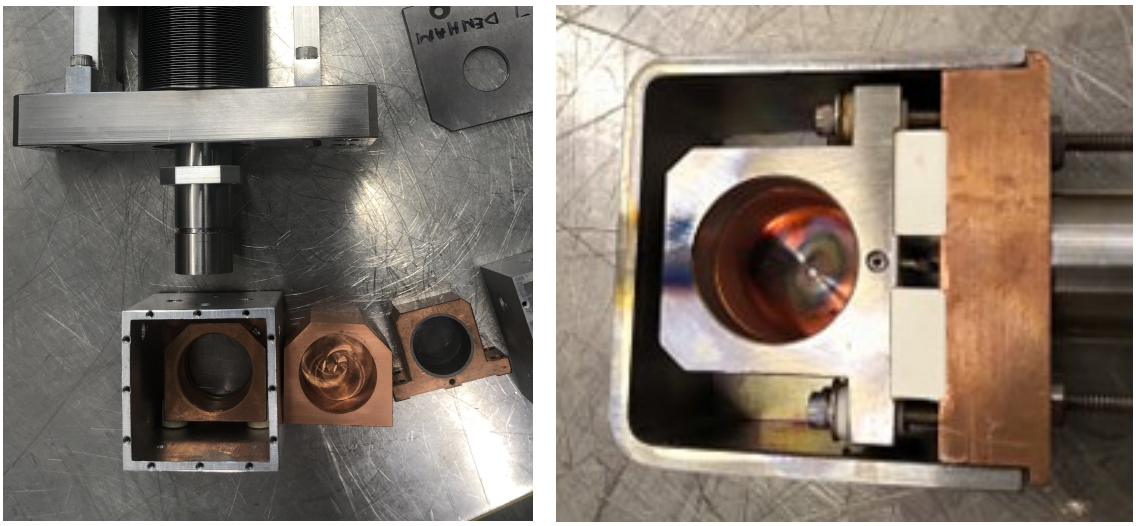

Figure 4.1: Typical ISAC Faraday cups: Various components (above left) and an assembelled cup with face plate removed (above right)

As the ion beam is positively charged and ions are typically coming at some fixed rate, it represents an electrical current. FCs are insulated from the beamline itself, and the block of metal is connected to a circuit with an ammeter. In this way, with sufficiently sensitive instruments, the beam current is measured. If the charge state of the beam is known, and the beam does not contain any contaminants, this current can be converted to a number of ions per second.

Typical beam currents measured in the ISAC-II linac are in the range of 100 pA to 10 nA, and the FCs can accurately measure to as low as around 10 pA. There are 16 FCs distributed along the ISAC-II linac beamlines.

Transverse Profile Measurements

Rotating Profile Monitors (RPMs) are the most commonly used device for measuring the transverse beam profile in ISAC. They consist of two 0.5 mm diameter tungsten wires, mounted at the end of a long arm that pivots around a fixed point. The tungsten wires are at an angle of 101 degrees with respect to each other, each at 50.5 degrees from the pivoting arm they are mounted on. The arm is driven from the exterior through a total angle of 22 degrees. During this time the two tungsten wires sweep through the beam, and an ammeter connected in series with the wires records the current at various points. One wire gives the horizontal profile and the other giving the vertical profile. A potentiometer connected to the external end of the pivoting arm provides a voltage reading which is converted into a corresponding

position at the beam axis end based on the geometry of the devices.

A picture of a typical RPM can be seen below in Figure 4.2. The head of the pivoting arm is circled in light blue (Tungsten wires aren't mounted in this picture), pivot point is circled in green, and the motor that drives the motion is circled in purple. There are 9 RPMs distributed along the ISAC-II linac beamlines.

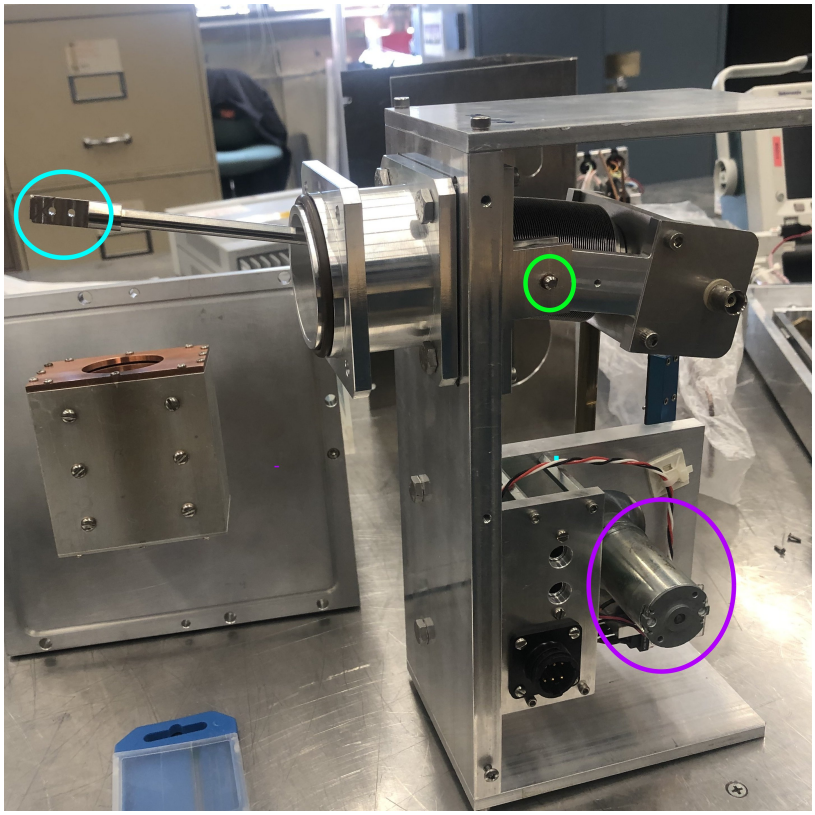

Figure 4.2: A picture of a typical ISAC Rotating Profile Monitor

Linear Profile Monitors (LPMs) are similar to RPMs in that they provide transverse beam profiles, but do so in a different manner. Shown in Figure 4.3, they are essentially a block of copper with two small slits at 90 degrees with respect to each other. The LPM is mounted at 45 degrees between the horizontal and vertical axes, such that one slit is parallel to the horizontal axis and the other parallel to the vertical axis. The slits are 0.5 mm wide by 32 mm long, and are intended to only allow a small slice of the beam profile through in a given dimension. The beam that goes through the slit is measured by a Faraday cup immediately downstream. A potentiometer connected to the LPM records the voltage as the LPM scans through the beam axis, which is then converted to a distance based on a known calibration. There are 9

LPMs distributed along the ISAC-II linac beamlines.

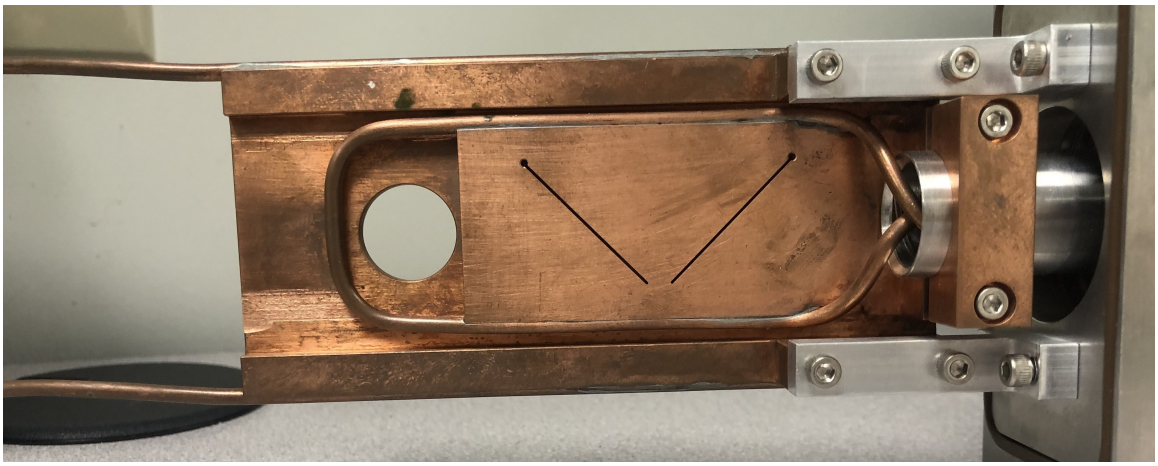

Figure 4.3: A typical ISAC Linear Profile Monitor

Longitudinal Profile Measurements

The time structure and energy per nucleon of the beam are measured on three flight time monitors (FTMs). FTMs are an assembly consisting of a 50 μm diameter biased tungsten wire that intersects the beam and emits secondary electrons. These electrons are detected by a micro-channel plate (MCP) detector with a time resolution of < 100 ps, giving a resulting energy/nucleon resolution of 0.1% [28].

4.1.2 Control System and Software

Virtually all equipment related to the operation of the beamlines in the ISAC-II linac is controlled via EPICS (Experimental Physics and Industrial Control System)[2]. Originally a collaboration between Los Alamos National Laboratory and Argonne National Laboratory, EPICS is now a free open source control system widely used across many different accelerator facilities around the world.

The control system consists of a large number of IOCs (input-output controllers), which are basically just stand-alone computers that are each responsible to communicate with a sub-section of the equipment comprising the linac. Equipment includes power supplies, solenoid valves (for mechanical movement of devices via compressed air), vacuum pumps, diagnostics (current readings, vacuum measurements).

EDM (Extensible Display Manager) is an EPICS extension used to create user-interfaces for operators, physicists, and other staff to monitor and control/communicate

[2]https://epics.anl.gov/index.php

with EPICS.

A phasing program written in Matlab[21] has historically been used to phase the linac. It uses a Matlab Channel Access (MCA) toolbox[29] developed at the Stanford Linear Accelerator Center (SLAC) that allows Matlab to read from and write to EPICS devices.

4.1.3 Data acquisition

Data was collected from the RIB control room at TRIUMF. This was done from a control room console, utilizing the MATLAB phasing program to carry out the linac calibration tunes. Concurrent with this, a python script was used to capture additional data, pulling the FTM profiles and reading FC currents at various points during the phasing process. Data was saved in ASCII files on a TRIUMF desktop and backed up to a different server.

4.2 Calibration Data

Presented in this section is the collected data for the calibration of the model. After summarizing the phase shifter data for handling the non-linear behaviour of the control system phase, a data set for a full linac tune is presented that can be used for calibration. These include calibration data on:

4.2.1 Phase Shifters

The behaviour of the linac phase shifters is summarized below. Data was collected as described earlier in Section 3.2.4. While a total of 44 phase shifters were measured, only four are presented here to give the reader a brief summary of the calibration - measurements of other cavities not shown varied in similar ways as those presented.

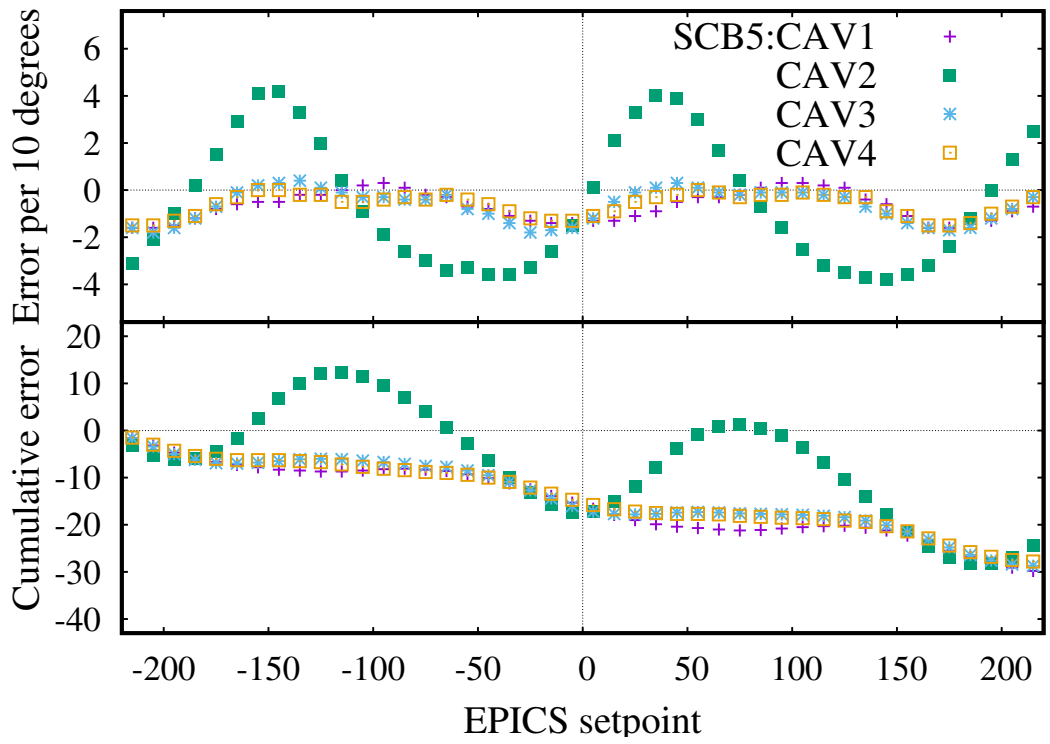

Figure 4.4: SCB5 cavities 1-4. Measured error in phase per 10 degree step (above) and cumulative error from starting point of -220 degrees (below) vs EPICS setpoint

Shown above in Figure 4.4, the phase error is the disagreement between a step in EPICS, $\Delta\phi_{epics}$, and the measured resulting step in the vector voltmeter, $\Delta\phi_{vm}$. The top plot shows the error for each individual 10 degree step, while the bottom plot shows the cumulative error in stepping from -220 degrees in the control system to the other measured phases.

$$Error = \Delta\phi_{epics} - \Delta\phi_{vm} \tag{4.1}$$

The resulting observations are as follows:

1. The relationship between the control system setpoint and the actual effect varies between different cavities.

2. The effect is not linear, the error changes depending on the range the setpoint is moving over.

3. The setpoint range over which the actual phase completes a full cycle ranges between around 374 and 384.

4. The ISAC-II global phase appears to adjust the phase relative to the reference signal in the opposite direction of all other phase shifters.

The measured data has been input into a python package[3], which allows conversion between the uncorrected control system phase and the correct control system phase.

4.2.2 Calibration tune - Oxygen-18, 4+ charge state

A beam of Oxygen-18, at a charge state of 4+ and energy per nucleon of 1.545 +/- 0.005 MeV/u was tuned through the linac and used for calibration measurements as outlined earlier in Section 3.2.6 from May 19th, 2021 through to May 21st, 2021.

The initial energy per nucleon was measured using the FTMs and is shown below in Figure 4.5.

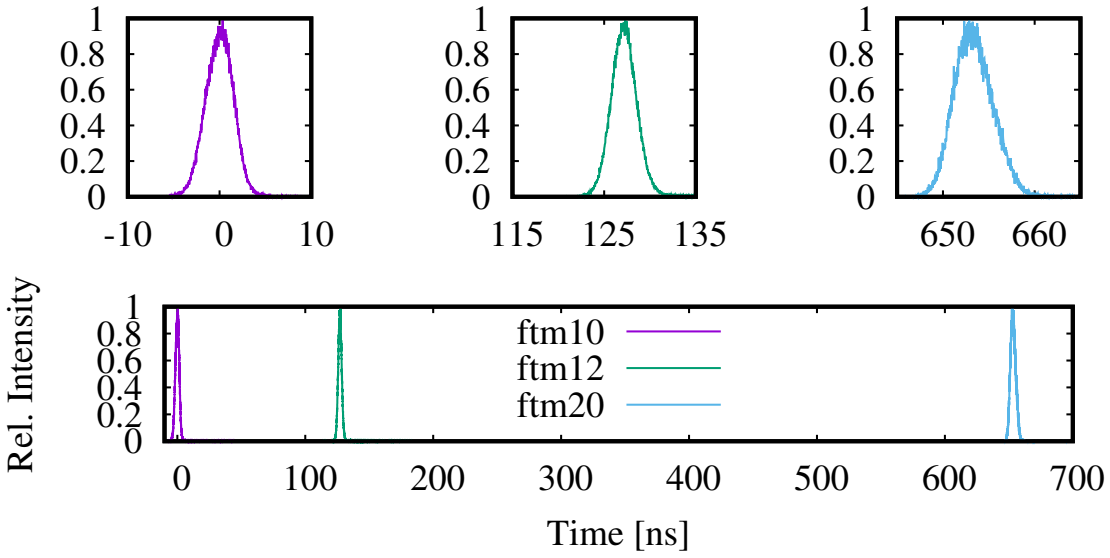

Figure 4.5: Initial flight time measurement of $18O^{4+}$ with DSB buncher and ISAC-II linac off.

DSB buncher calibration and longitudinal emittance reconstruction

The DSB buncher was varied to ten different amplitudes to both calibrated it's amplitude scaling parameter as well as to attempt to reconstruct the longitudinal emittance arriving from the upstream ISAC-I accelerators. A summary of the data collected can be seen below in Table 4.1, while the `TRANSOPTR`-generated fit to the collected data can be seen in Figure 4.6

[3] https://gitlab.triumf.ca/hla/acc-utilities/phase_shifter

A_{EPICS} [a.u.]	Z_{ftm10} [cm]	Z_{ftm12} [cm]	Z_{ftm20} [cm]	E/u (measured) [MeV/u]
776.6	0.75	0.78	1.04	1.540
790.0	0.66	0.69	1.02	1.540
742.2	0.75	0.72	0.9	1.541
400.0	2.9	2.65	3.74	1.543
550.0	1.92	1.74	2.14	1.542
700.0	0.98	0.93	1.1	1.541
850.0	0.84	0.92	1.52	1.540
950.0	1.35	1.41	2.46	1.540
1100.0	2.38	2.39	3.97	1.539
1250.0	3.53	3.43	5.35	1.538

Table 4.1: Collected beam size data for the DSB buncher from the $18O^{4+}$ calibration tune established in May 2021

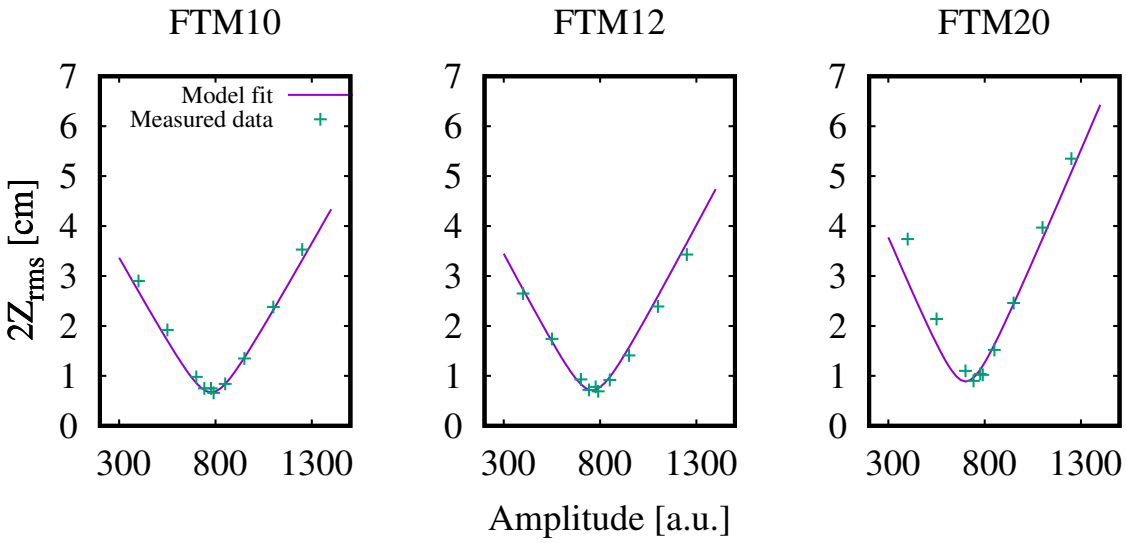

Figure 4.6: Comparison of the measured $2Z_{rms}$ beam size with the TRANSOPTR fit to the data for the three FTMs.

Linac calibration tune, 38 cavities

The final energy per nucleon reached with all 38 cavities on and optimized was 8.93 +/- 0.01 MeV/u, and is shown below in Figure 4.7. The energy achieved after each cavity is shown in Table 4.2. As an example of the result arising from the calibration procedure (recall Section 3.2.6), the fit parameters are also given. The linac energy gain profile can be seen graphically in Figure 4.8.

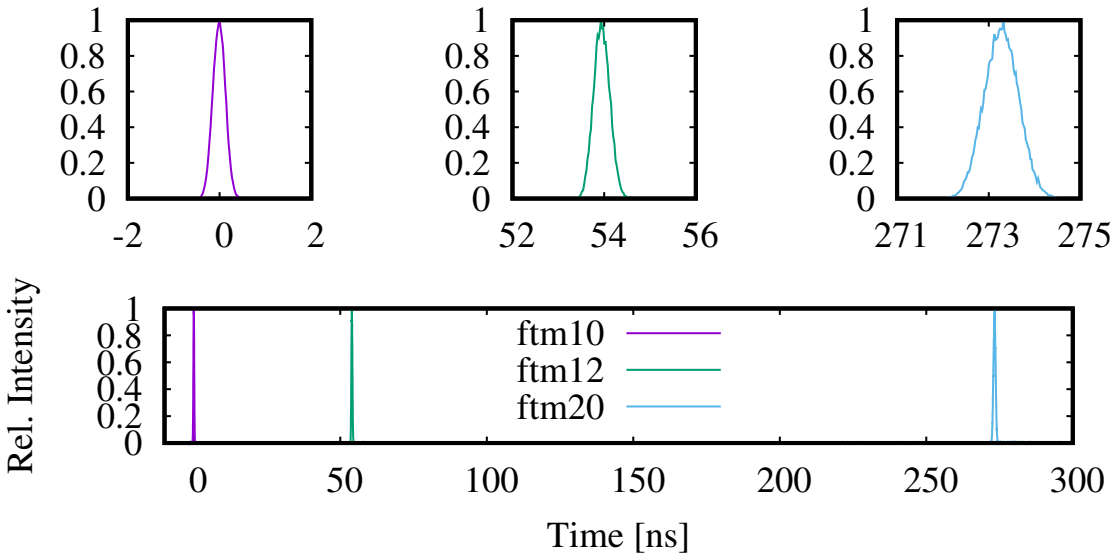

Figure 4.7: Final flight time measurement of $18O^{4+}$ with DSB buncher and 38 ISAC-II cavities on.

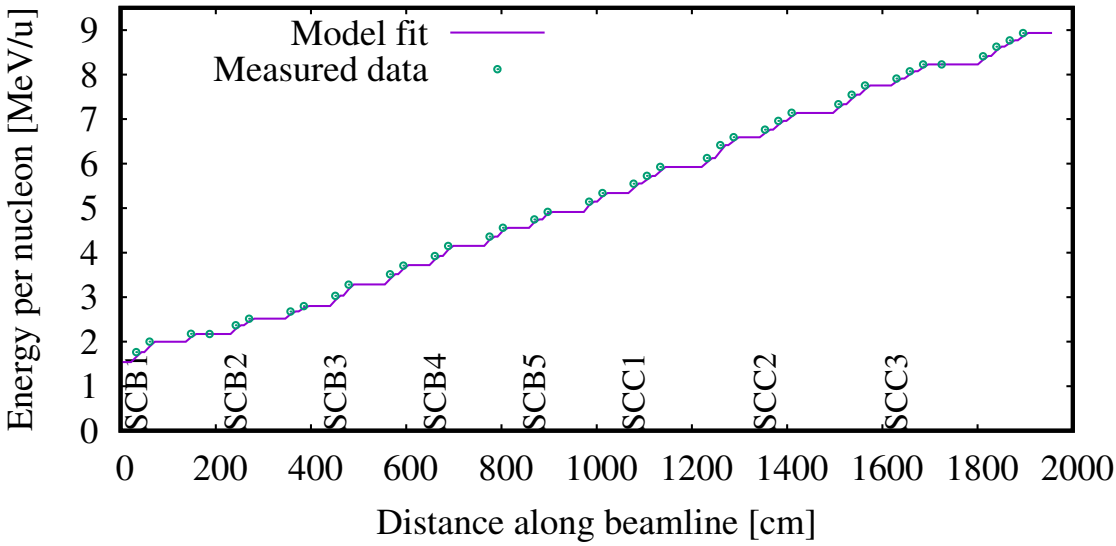

Figure 4.8: Linac energy profile of $18O^{4+}$ with 38 ISAC-II cavities on showing model fit to calibration data.

Cavity	$\boldsymbol{A}$	**C** $[\cdot 10^{-3}]$	$\boldsymbol{\phi}$ $[deg]$	$\boldsymbol{\theta}$ $[deg]$	**E/u [MeV/u] (model fit)**	**E/u [MeV/u] (measured)**	$\boldsymbol{\sigma_C}$ $[\cdot 10^{-3}]$	$\boldsymbol{\sigma_\theta}$ $[deg]$
SCB1:CAV1	2000	7.50	121.63	-187.0	1.765	1.766	0.22	0.6
SCB1:CAV2	1200	12.62	17.87	-145.2	1.998	2.000	0.39	0.6
SCB1:CAV3	1600	7.90	-64.97	-124.3	2.172	2.176	0.2	0.6
SCB1:CAV4	0	-	0.0	-	-	-	-	-
SCB2:CAV1	2000	6.76	-52.67	-67.1	2.368	2.368	0.14	0.6
SCB2:CAV2	1100	9.10	-110.7	96.0	2.517	2.516	0.14	1.0
SCB2:CAV3	1100	10.41	-108.03	30.1	2.680	2.678	0.11	0.8
SCB2:CAV4	1000	8.70	157.39	167.4	2.800	2.800	0.18	1.2
SCB3:CAV1	2200	6.81	41.99	57.2	3.034	3.030	0.07	0.6
SCB3:CAV2	2200	7.47	-119.13	-178.3	3.285	3.282	0.14	0.5
SCB3:CAV3	2200	6.30	49.5	117.2	3.517	3.514	0.06	1.0
SCB3:CAV4	1400	9.39	-103.51	-106.7	3.719	3.708	0.29	0.8
SCB4:CAV1	1900	7.25	-145.23	-1.6	3.927	3.924	0.22	0.8
SCB4:CAV2	2000	7.54	50.87	156.4	4.151	4.151	0.08	0.6
SCB4:CAV3	1900	7.95	156.67	153.6	4.360	4.360	0.08	0.5
SCB4:CAV4	2100	7.10	-118.94	-121.0	4.558	4.558	0.07	0.7
SCB5:CAV1	1800	7.58	125.5	-18.5	4.745	4.746	0.10	0.5
SCB5:CAV2	1800	8.47	-93.17	145.0	4.913	4.915	0.08	0.5
SCB5:CAV3	2100	8.00	-147.0	89.6	5.147	5.143	0.06	0.6
SCB5:CAV4	1800	7.27	-29.41	-131.4	5.338	5.338	0.07	0.7
SCC1:CAV1	1900	8.02	99.39	50.3	5.554	5.550	0.08	0.6
SCC1:CAV2	2300	5.56	138.86	-2.1	5.724	5.724	0.06	0.7
SCC1:CAV3	1900	7.88	168.45	-56.9	5.924	5.926	0.08	0.6
SCC1:CAV4	1900	7.60	-18.22	53.9	6.124	6.125	0.07	0.8
SCC1:CAV5	2100	9.44	180.58	-176.5	6.415	6.417	0.07	0.5
SCC1:CAV6	1100	10.99	189.67	132.5	6.592	6.593	0.12	0.9
SCC2:CAV1	1500	8.10	4.83	66.9	6.764	6.762	0.08	0.8
SCC2:CAV2	1600	8.63	173.89	-175.1	6.959	6.959	0.09	0.7
SCC2:CAV3	1600	8.01	-79.1	-91.6	7.139	7.142	0.08	0.7
SCC2:CAV4	1800	7.84	163.11	-167.2	7.335	7.335	0.08	0.6
SCC2:CAV5	1800	8.27	129.52	139.4	7.550	7.548	0.07	0.6
SCC2:CAV6	2000	7.29	-116.24	-124.5	7.755	7.755	0.07	0.6
SCC3:CAV1	1200	9.12	138.63	-50.2	7.911	7.910	0.14	0.7
SCC3:CAV2	1800	6.29	121.91	-78.0	8.075	8.073	0.10	0.7
SCC3:CAV3	1200	8.80	44.71	-164.2	8.230	8.228	0.13	0.9
SCC3:CAV4	0.0	-	0.0	-	-	-	-	-
SCC3:CAV5	1600	7.92	49.48	-11.1	8.418	8.413	0.10	0.7
SCC3:CAV6	1800	8.11	-68.45	-126.0	8.628	8.627	0.08	0.6
SCC3:CAV7	1200	8.30	115.29	43.4	8.770	8.768	0.15	0.7
SCC3:CAV8	2300	4.93	81.32	17.7	8.935	8.933	0.07	0.7

Table 4.2: Collected data and resulting fit parameters resulting from the $18O^{4+}$ calibration tune established in May 2021

More detailed analysis is given later in Section 5.2, but there are a few pertinent comments to make about this data now. First, it is clear immediately that the amplitude scaling parameter C varies significantly over the 38 available cavities, ranging from a low of $4.93 \cdot 10^{-3}$ up to a high of $1.26 \cdot 10^{-2}$. If not already self-evident, this highlights the need for beam-based calibration: the calibration of cavities was not indended to vary, but rather arose from differences in the way the cavity fields are read, transmitted, and converted to digital signals. This could include: the geometry of each individual cavity and possible small difference differing from the ideal designs, positions of the RF pickups, length and attenuation of each individual RF cable carrying the signal back to the control consoles, and the individual power meter and ADC for each cavity. Similarly, the phase offset θ varies significantly primarily due to the differing cable lengths between the common RF refence signal and the RF coupler for each individual cavity. The final observation to note is that the output energy in the model resulting in the fits agrees well throughout the linac, differing by at most 0.011 MeV/u and agreeing to within the uncertainty of the measurement for the majority of cavities.

4.3 Automated linac adjustments

4.3.1 Transverse beam size

While the main focus of this work is on the optimization of the cavity phases for a given desired energy, it was also crucial to ensure the linac model was capable of accurately modeling the transverse beam size. RF cavities typically have a defocusing effect on the beam in the transverse plane[4], and in addition to this the necessary focusing strength of elements changes as the beam energy changes.

To benchmark the quality of the model in the transverse dimensions, a beam of $20Ne^{4+}$ was drifted through the ISAC-II linac in October of 2020. A number of data sets were collected with profile monitors to allow for re-construction of the intitial beam properties into the linac, and these were subsequently used as inputs into the model and compared to measured beam sizes at other diagnostics further down the machine.

Sigma matrix reconstruction

[4]When the cavities are providing a longitudinal focusing effect

To re-construct the beam properties entering the ISAC-II linac, an assumption is made that there is no coupling between the horizontal, vertical, or longitudinal dimensions[5]. Mathematically this is equivalent to stating that the beam sigma matrix from Equation 3.2 can be reduced to three separate 2 by 2 matrices as shown below in Equation 4.2.

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{11} & \sigma_{12} & 0 & 0 & 0 & 0 \\ \sigma_{21} & \sigma_{22} & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{33} & \sigma_{34} & 0 & 0 \\ 0 & 0 & \sigma_{43} & \sigma_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{55} & \sigma_{56} \\ 0 & 0 & 0 & 0 & \sigma_{65} & \sigma_{66} \end{bmatrix} \tag{4.2}$$

Utilizing a quadrupole and drift space with known transfer matrices ending at a profile monitor, the beam size at the monitor has a simple relationship (recall Eq. 2.26) to the initial beam properties upstream of the quadrupole:

$$\begin{bmatrix} \sigma'_{11} & \sigma'_{12} \\ \sigma'_{21} & \sigma'_{22} \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \cdot \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix} \cdot \begin{bmatrix} M_{11} & M_{21} \\ M_{12} & M_{22} \end{bmatrix} \tag{4.3}$$

Which when expanded to show the relationship between the final beam size σ'_{11} and the initial sigma matrix, we get

$$\sigma'_{11} = \sigma_{11}M_{11}^2 + 2\sigma_{12}M_{11}M_{12} + \sigma_{22}M_{12}^2 \tag{4.4}$$

Where the variables are the transfer matrix elements which depend on the focal strength of the quadrupole and the drift lengths on either side of it. We then have a set of measurements $I_i, \sigma'_{11,i}$, where I_i is the current driving the quadrupole's focal strength. We use a minimization routine to find the parameters σ_{11}, σ_{12}, and σ_{22} which best fit the observed measurements.

Examples of the data received from the profile monitor is shown below in Figure 4.9. This shows how the profiles in both transverse dimensions can change as the quadrupole current is changed. From this data the mean and variance of the beam distribution is calculated.

[5]This is a common assumption made whether re-constructing the beam or during direct emittance measurements. This assumption can be violated in numerous ways, including when there is x-y coupling as is possible exiting a solenoid or x-z coupling like dispersion which can happen around non-achromatic bend sections

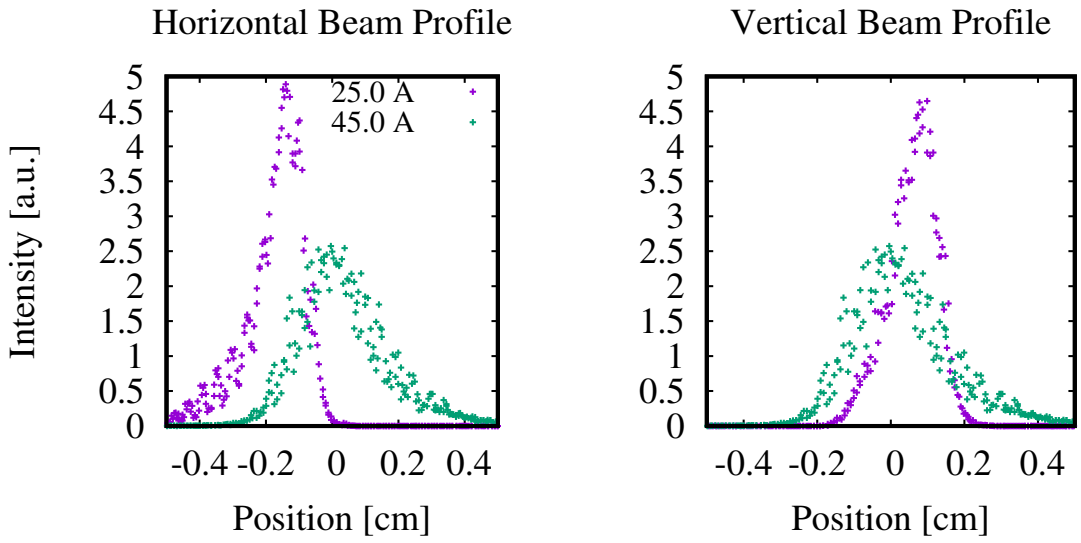

Figure 4.9: Typical beam profiles collected using DSB:RPM4 before processing to 2rms beam size.

Eleven data points were collected, and `TRANSOPTR`'s internal optimizer was used to find the initial beam properties that best fit the data. The result of this fit is depicted below in Figure 4.10.

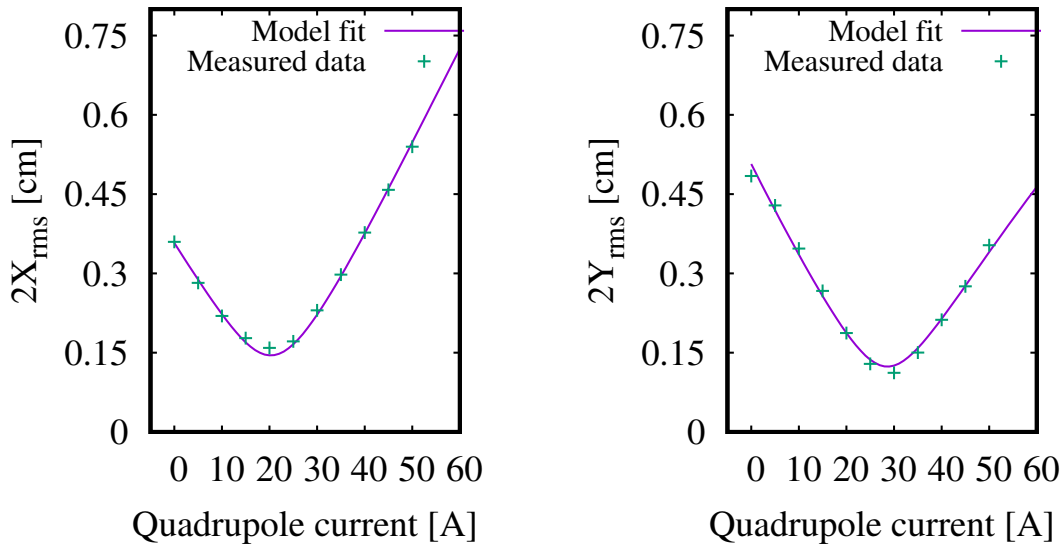

Figure 4.10: Measured and fit beam size after re-construction of sigma matrix upstream of DSB:RPM4

A graphical representation of this method is shown below in Figure 4.11. The

initial beam is the same in all eleven cases, and the quadrupole DSB:Q3 (shown as a grey rectangle near the start of the plot) has differing effects on the beam envelope as the current is varied. This method can also be plotted in phase space as shown in Figures 4.12 and 4.13, and demonstrates the transformation from a divergent to focused beam.

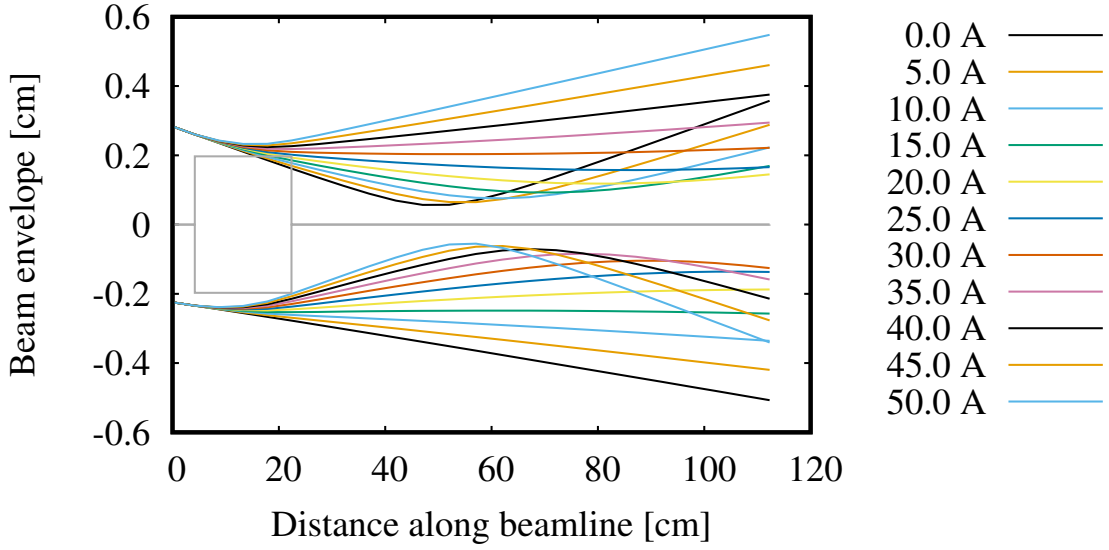

Figure 4.11: Beam envelope (y envelopes plotted as negative for clarity) for DSB section used for emittance reconstruction.

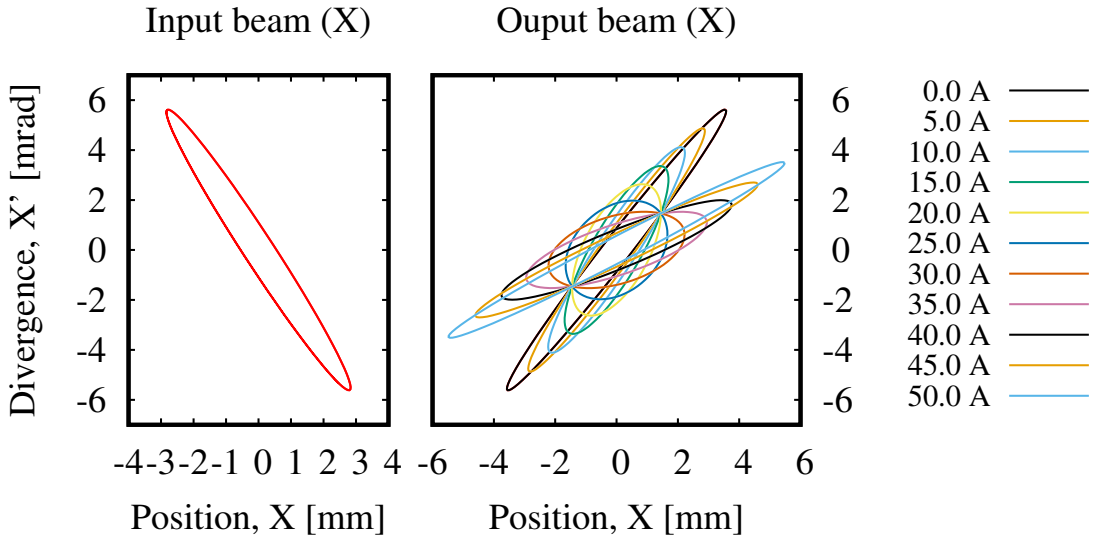

Figure 4.12: Initial ellipse in X phase space on left, various output ellipses on right for varying quad currents

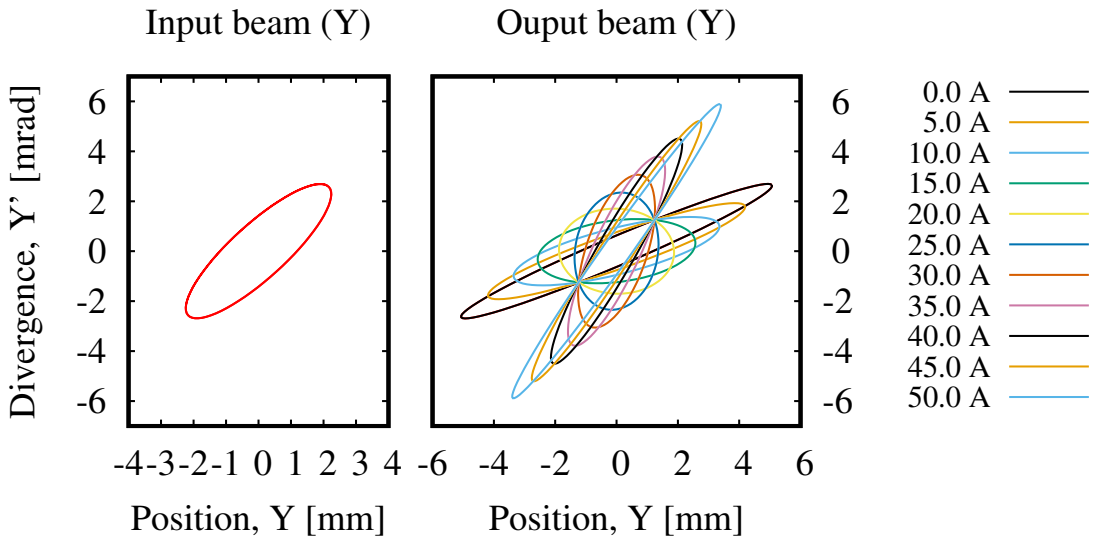

Figure 4.13: Initial ellipse in Y phase space on left, various output ellipses on right for varying quad currents

Beam size comparison

With the initial beam sigma matrix established, profiles along the length of the beamline were measured and beam size was compared to the simulated envelope that the model predicts. The result is shown below in Figure 4.14.

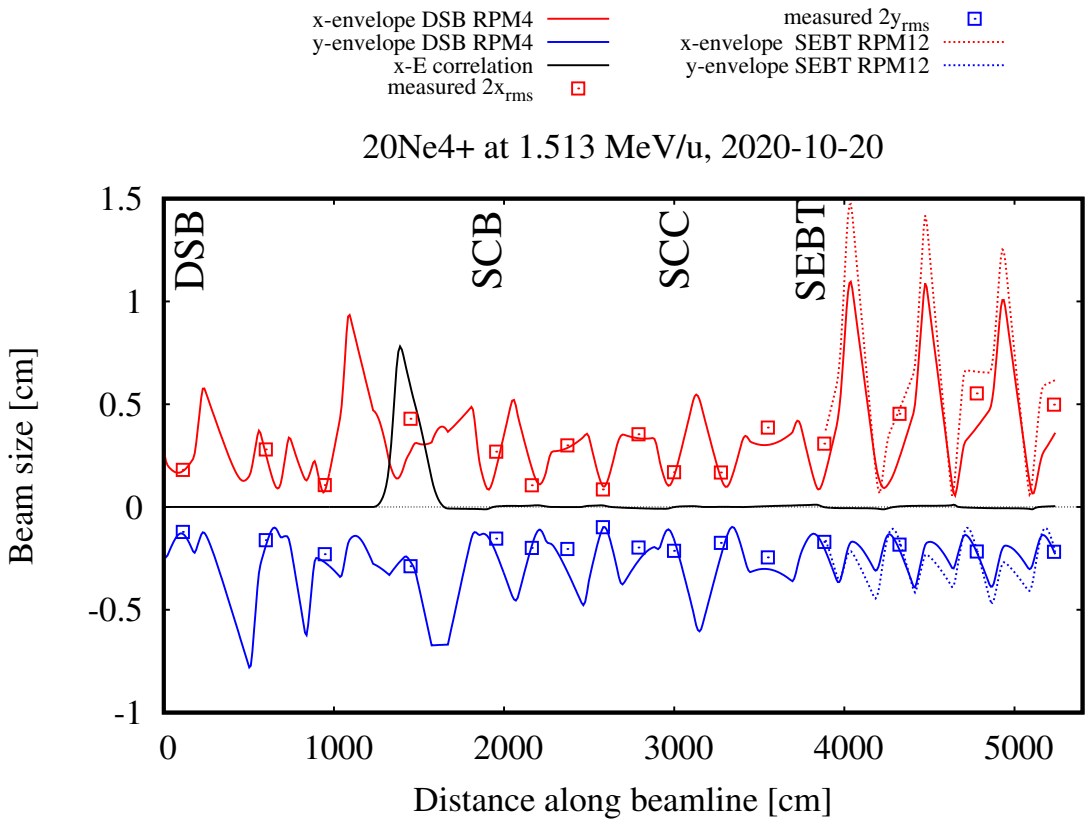

Figure 4.14: Beam envelope generated by `TRANSOPTR` model, using re-constructed initial sigma matrix, comparing to measured beam sizes along the linac

Qualitatively the agreement between the model and observation appears quite good for all of the vertical (Y) profile measurements. Horizontal (X) profiles agree quite well up to the end of the linac, approximately at 4300 cm on the x-axis of the plot. One exception to this is the measurement within a dispersive section upstream of the linac around 1500 cm. The strong dispersion in the horizontal plane suggests that the beam size here is incorrect due to an error in the modelled energy spread of the beam.

The disagreement in X for the last 3 monitors is significant and warrants further investigation. A second emittance re-construction was completed with a profile monitor downstream of the linac for comparison, and the envelope based on this measurement is shown in dashed lines. Most notably, the horizontal emittance in the measurement after the linac is larger than that measured upstream of the linac. The fit beam properties for the two measurements are shown in Table 4.3 as well as graphically in Figure 4.15.

Property	Measured at DSB:RPM4	Measured at SEBT:RPM12
$2X_{rms}$ $[mm]$	2.59	4.90
$2X'_{rms}$ $[mrad]$	2.79	2.25
r_{12}	0.903	0.910
$4\epsilon_{x,rms}$ $[\mu m]$	3.11	4.57
$2Y_{rms}$ $[mm]$	2.23	1.84
$2Y'_{rms}$ $[mrad]$	2.44	2.25
r_{34}	0.811	0.820
$4\epsilon_{y,rms}$ $[\mu m]$	3.19	3.37

Table 4.3: Re-constructed beam properties at SEBT:RPM12 from two separate measurements of 20Ne4+ on October 20th, 2020.

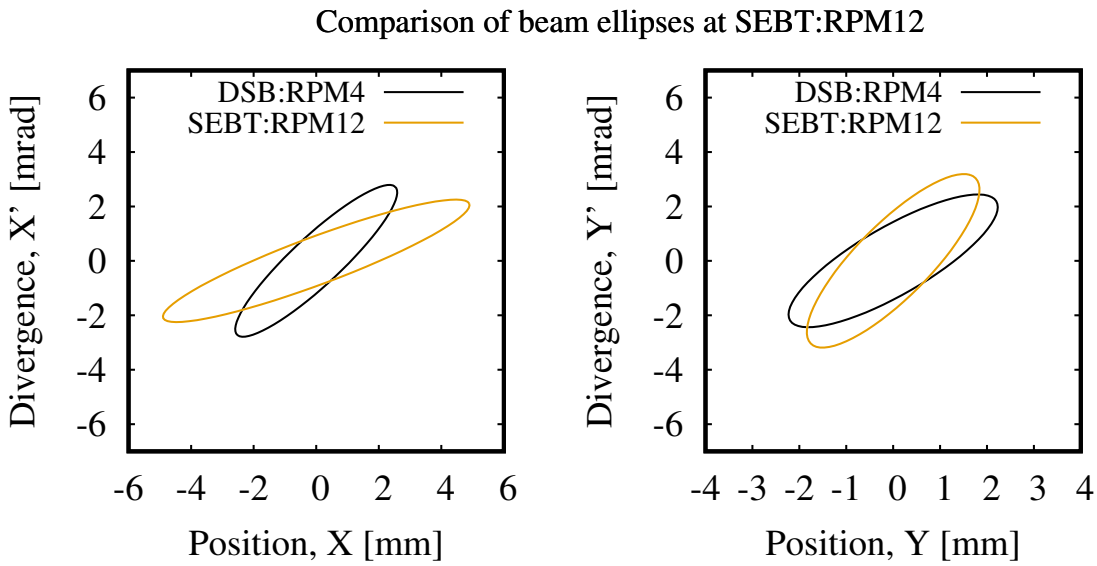

Figure 4.15: Re-constructed beam ellipses at SEBT:RPM12 from two separate measurements of 20Ne4+ on October 20th, 2020.

A few theories as to the cause of the discrepency in X in the SEBT section present themselves. First, it is possible that the beam exiting the linac entering SEBT is misaligned through the first four transport quads SEBT:Q9 through Q12. This could cause emittance growth due to higher order aberrations. The centroid of the measured profile just upstream of this section is consistently abnormally off of the beam axis, in this measurement approximately 7.5 mm. This is approximately one third of the quadrupole aperture and if the beam between the profile monitor and SEBT:Q9

continues to diverge from the beam axis, it could be that the beam is approaching the aperture radius of Q9.

The second possible cause would be changing behaviour of an upstream device. It was noted during data collection that the MEBT chopper behaviour appeared to have changed. This device is far upstream of the ISAC-II linac, and it is a high voltage plate paired with RF that kicks unwanted pulses out of the beam by applying an electric field in X. If some pulses were kicked but not sufficiently far so as to stop them on a slit, it would have the effect of increasing the horizontal beam emittance. If this happened between the measurement used to re-construct the beam emittance and the measurement of the profiles in the SEBT transport section this could possibly help explain the disagreement.

4.3.2 Full linac autophase

Using the calibration data described previously in Section 4.2.2, the linac model was run to predict the desired phase setpoint and output energy for each of the available 35 cavities in July of 2021. This was done using a beam of Argon-40 at a charge state of 11+ and an injected beam energy per nucleon of 1.485 +/- 0.008 MeV/u.

To evaluate the accuracy of the model-based calculations, the energy vs phase behaviour of each cavity was measured to compare the observed cavity phase to the predicted from the model. The method was as follows:

1. Turn on the next available cavity and collect phase vs energy data.
2. Restore the cavity phase to that predicted by the model and measure the beam energy.
3. Fit the data from Step 1 to establish the disagreement between his emperically established phase and the model-predicted phase.
4. Proceed to the next cavity and repeat.

The results of these tests are shown below in Figures 4.16 and 4.17. The former shows the measured total energy after each additional cavity compared to the model while the latter shows energy and phase error after each cavity.

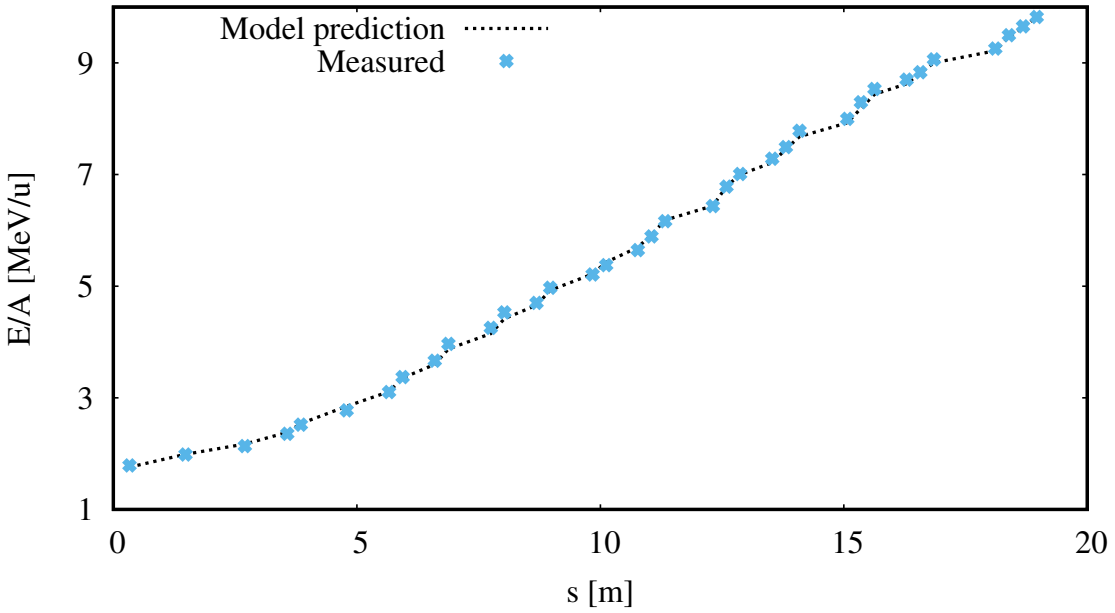

Figure 4.16: Comparison between measurement and model for full linac autophase test.

Transport efficiency (transmission) through the ISAC-II linac for the test averaged 95% over the 35 cavities measured. Qualitatively the overall output energy shown in Figure 4.16 matches well with the model, and broadly speaking this paired with the high transport efficiency indicates that the new approach is effective at automating the cavity phasing process.

The phase and energy errors shown in Figure 4.17 present an opportunity for a more detailed discussion on the accuracy of the new method. As noted previously [30], the E/A error out of the very first cavity available was 0.03 MeV/u. This is a large error of 10% relative to the total energy/nucleon gained through this first cavity of 0.3 MeV/u. With no other errors in the model, this discrepency in the output energy of the first cavity paired with the comparitively low velocity at the front-end of the linac and the large drift space before the next cavity would lead to a phase error of -20 degrees at the next cavity. The observed phase error of the next cavity is -9 degrees.

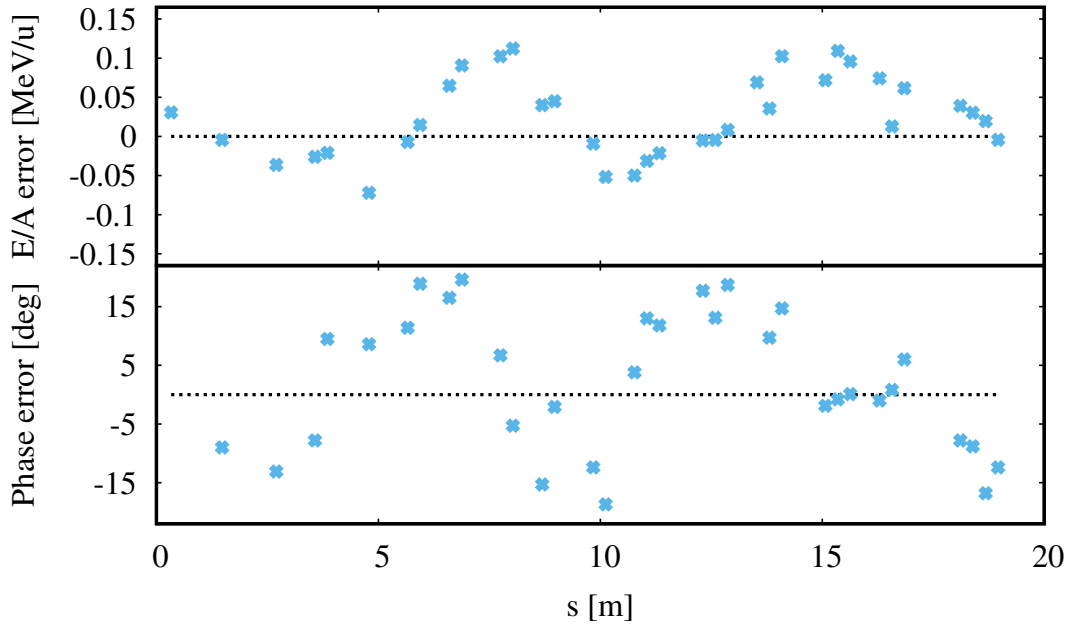

Figure 4.17: Error between measurement and model for full linac autophase test.

Note also that the output energy error of the first cavity could arise from a number of sources. It could arise from a 10% error of in calibration of the cavity amplitude, or from as little as a four degree error in the phase offset of the first cavity. A combination of these two posible errors is also possible.

Phase offset error identified during test

The calibration data used for the test was done utilizing the first linac cavity as the 're-zeroing' point, rather than the upstream DSB buncher. However during the test it was found that although cavities up to cavity 20 were close to the predicted phase, cavity 21 was very far off. After investigation it was determined that a mistake was made during the initial calibration - it is not possible to 're-zero' the 141.44 MHz cavities (cavities 21-40) using the first linac cavity at 106.08 MHz, as these are non-integer multiples of each other. Rather, the 35.36 MHz DSB buncher upstream or a time diagnostic spanning a range of at least one period of a 35.36 MHz signal must be used, as the linac cavities are both harmonics of this 35.36 MHz signal. The problem is shown graphically below in Figure 4.18 - if the model was calibrated using two random peaks both of the 106.08 MHz cavities and 141.44 MHz cavities, they would then have a calculated offset that shifts the signals so that both peak at $\phi = 0$ as shown in the Figure. However during the test or any automatic phasing, the linac

RF is re-zeroed with the arrival of the beam using a *random* peak of the 106.08 MHz signal, which could be any of the three peaks of the orange trace shown. It is clear that if either of the peaks not at $\phi = 0$ are selected, the 141.44 MHz signal will now be 'out of sync' by either 120 or 240 degrees (at 141.44 MHz).

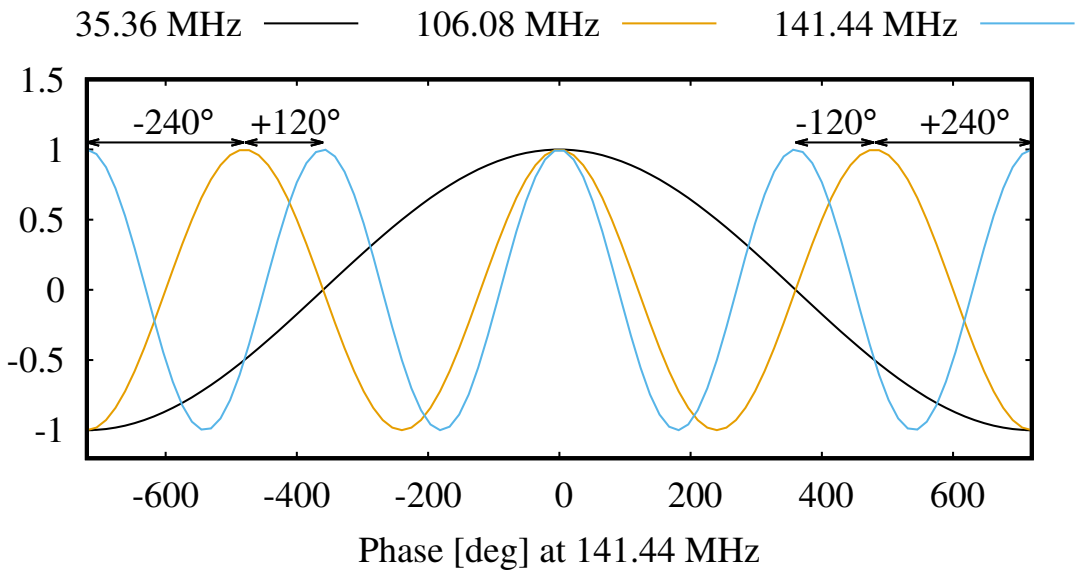

Figure 4.18: Comparison between the 35.36, 106.08, and 141.44 MHz harmonics demonstrating the problem with re-zeroing using a 106.08 MHz signal

Due to the limited opportunities available with the ISAC-II linac and the overhead involved with preparing for the test it was decided to adjust the global phase of the second half of the linac by whichever multiple of 120 degrees that brought the cavity phase closest to the model calculated value. This in practice was identical to re-calibrating the linac and restarting the test, but saved a critical amount of time during the beam test. The phase that was adjusted shifts all 20 of the 141.44 MHz cavities by the same amount. Following this adjustment the test continued as planned.

Following this test the proposed method outlined earlier in Section 3.3.1 was updated to reflect the need to re-zero with a 35.36 MHz device to avoid this discrepency during regular use of this new method.

4.3.3 Re-phasing (cavity failures)

A dedicated separate beamtime for testing of the new method functioning as a failure-recovery tool was deemed unecessary as the full linac auto-phase already demonstrates this ability.

The initial calibration tune described earlier in Section 4.2.2 was completed with 38 cavities, and the autophase data collected in Section 4.3.2 was done with only 35 operational cavities, essentially both an autophase with a different beam as well as a failure recovery after 3 failed cavities. The cavities unavailable during the test were cavities 2, 5, and 9 as shown graphically below in Figure 4.19.

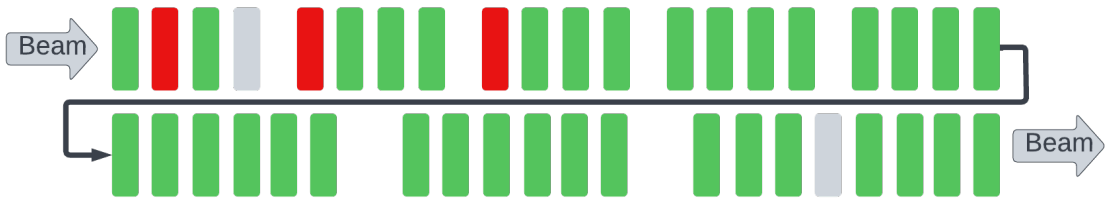

Figure 4.19: Diagram showing cavities available during autophase test

Shown in Figure 4.19 are the 40 cavities in the ISAC-II linac. The 35 cavities available during the autophase test are shown as green, cavities available during the calibration tune but omitted during the autphase test are shown in red, and cavities not available in either the calibration or autophase test are shown in grey.

Note that the cavities failed are very close to the start of the linac where the new approach is more sensitive to errors. Further to this, in regular operation the need for failure recovery from 3 cavities is not likely - much more common is for a single cavity to fail during delivery. Last, the second cavity having failed is suspected to be contributing significantly to the phase error of the third cavity discussed earlier in reference to Figure 4.17.

The final beam property of interest when recovering from once or more failed cavities is the output linac energy. The test with 3 failed cavities resulted in a final energy 9.822 MeV/u compared to a model prediction of 9.826 MeV/u, smaller than 0.1% error. It is clear however from Figure 4.17 that this 0.1% is more a result of using all available cavities as well as luck. The maximum output energy error compared to the model peaks at around 2%. For this reason it will be necessary to procedurally check and fine-tune the final beam energy when recovering from cavity failures during regular operation.

This chapter has presented the collected data alongside some initial discussion. The next chapter will cover some more detailed discussion about the results including some analysis of possible errors and their effects.

Chapter 5

Discussion

This chapter will provide some more detailed discussion and analysis of the data presented in the previous chapter. It will also provide some discussion around sources of error and similar projects at other labs.

5.1 Evaluating TRANSOPTR vs Trace 3-D

To benchmark the established TRANSOPTR model of the ISAC-II linac, it was first benchmarked against previous Trace 3-D simulations of the same beamlines. Trace 3-D simulations of the ISAC-II linac were done by M. Marchetto in three separate sections: the DSB injection beamline, acceleration through the SCL, and the post-linac transport section SEBT. Here envelopes with the same properties and initial conditions are modelled with TRANSOPTR for each of these three sections, and plotted together with the digitized Trace 3-D envelopes[1] for comparison.

Trace 3-D uses only magnetic field gradient for quads, while the implemented TRANSOPTR model actually varies the control system parameter which is the current driving this field. For this reason the current producing the same field gradient was used in TRANSOPTR. For the 40 accelerating cavities, the Trace 3-D model ran with each cavity providing an effective voltage, V_{eff} of 1.08 MV. This was likely done for convenience, because as the accelerating efficiency changes for each cavity as the beam velocity changes, this is essentially equivalent to setting the 'nominal' accelerating field of each cavity at different values. Nonetheless, a comparison to the Trace 3-D envelopes along the linac was still desired so the scaling of the longitudinal electric

[1]The Trace 3-D output envelope is graphical only in the version used. These envelopes were digitized and data provided for comparison by O. Shelbaya.

field E_z was varied for each of the 40 cavities to match the Trace 3-D assumption of $V_{eff} = 1.08MV$ for each cavity.

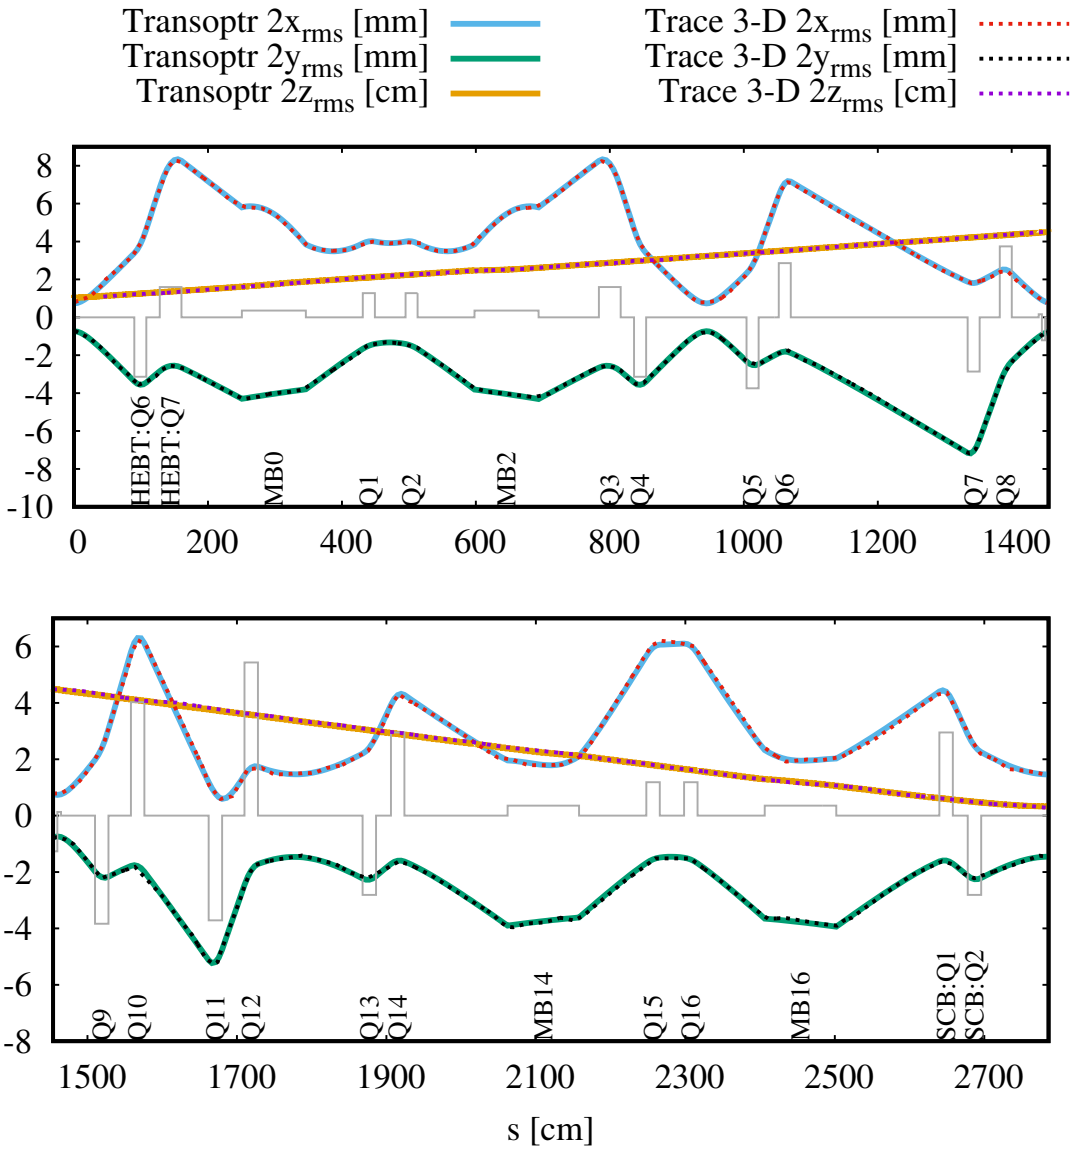

Figure 5.1: The 2rms envelope through the DSB section, comparing `TRANSOPTR` (solid lines) to Trace 3-D (dashed lines). Relative x focal strength [a.u.] in grey.

For the eight superconducting solenoids, Trace 3-D uses hard-edged solenoids where `TRANSOPTR` is integrating through the $B_z(s)$ field maps. To make a valid comparison, the integrated strength of the `TRANSOPTR` solenoids was matched to the hard-edged Trace 3-D cases.

All three comparison are done for a beam of mass 30.0 amu. For the DSB sec-

tion, both models are simulating charge state 5+ and energy/nucleon of 1.50 MeV/u. For the SCL section, both models are simulating charge state 10+ and input energy/nucleon of 1.50 MeV/u. For the SEBT section, both models are simulating a charge state 10+ and input energy/nucleon of 8.20 MeV/u.

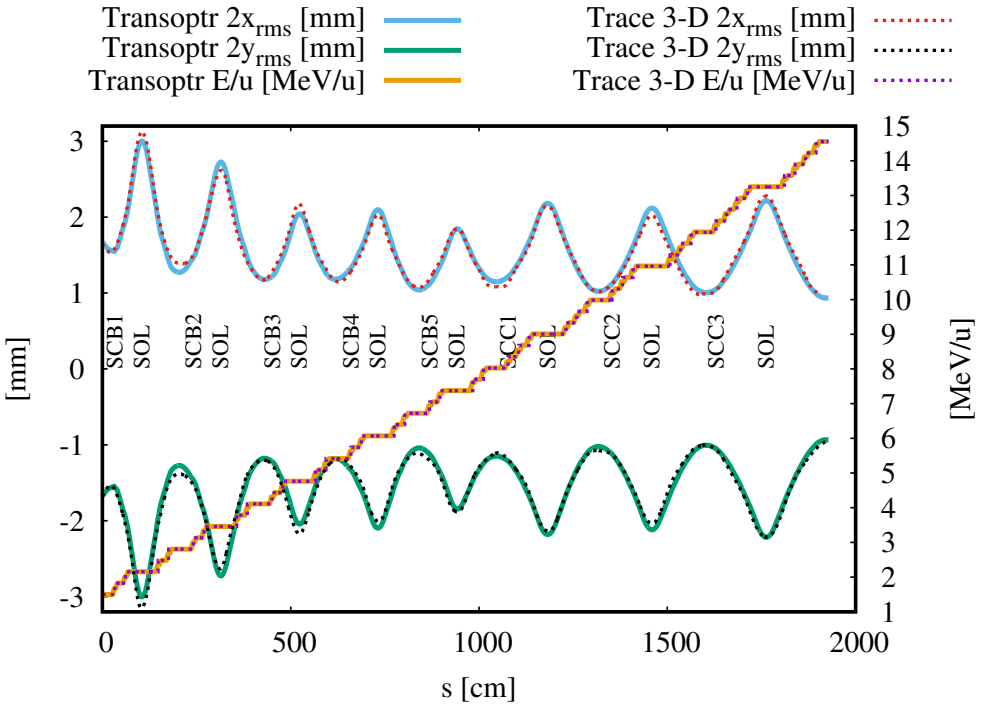

Figure 5.2: The 2rms envelope through the SCL section, comparing `TRANSOPTR` (solid lines) to Trace 3-D (dashed lines)

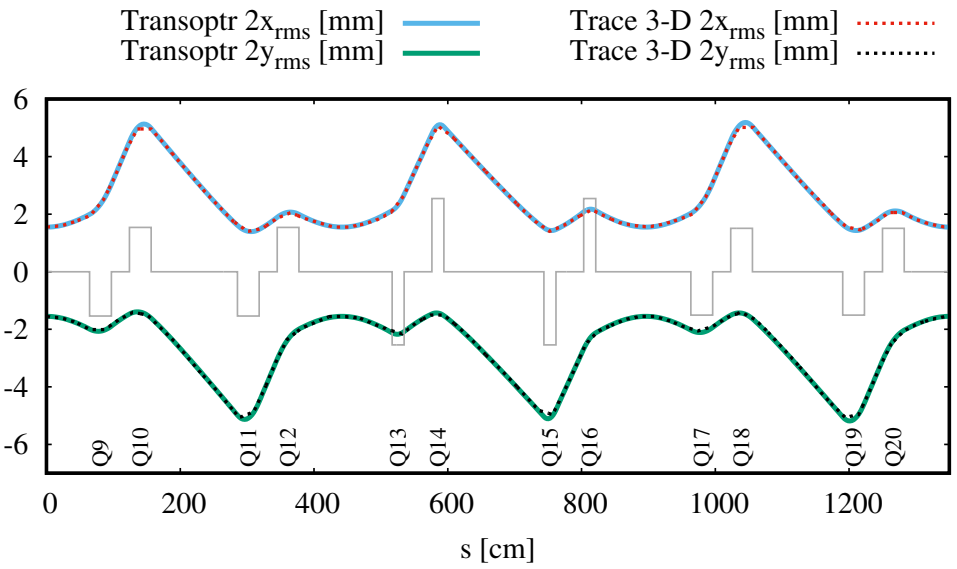

Figure 5.3: The 2rms envelope through the SEBT section, comparing `TRANSOPTR` (solid lines) to Trace 3-D (dashed lines). Relative x focal strength [a.u.] in grey.

As shown above in Figures 5.1, 5.2, and 5.3, the agreement between the two models was qualitatively quite good. The most notable divergence between the two models is at the centre of the first cryomodule at around 100 cm on the x-axis in Figure 5.2 as well as the subsequent beam waist near s = 200 cm. Despite these small discrepencies the beam optics further downstream remain in good agreement, as the integrated effect over the full lenfth of each solenoid remains the same between the two models.

5.2 Analysis of Calibration Data

The accuracy of the cavity calibrations to be used for future automated phasing methods is of critical importance and is discussed here. The possible errors affecting the calibrated parameters C and θ can be separated into four distinct categories. These include the flight time monitors themselves, cavity amplitude and phase setpoints in the control system, incoming beam properties, and the model.

Flight time monitors

The FTMs themselves have a time resolution of better than 100 ps [28]. An example of a typical profile acquired from a single FTM is shown below in Figure 5.4. This example is from the input energy of the ISAC-II linac, and has a larger time

spread than most measurements with ISAC-II cavities on, as the DSB buncher and ISAC-II cavities provide longitudinal bunching to the beam. From this example the mean, standard deviation, and error of the mean are summarized below in Table 5.1.

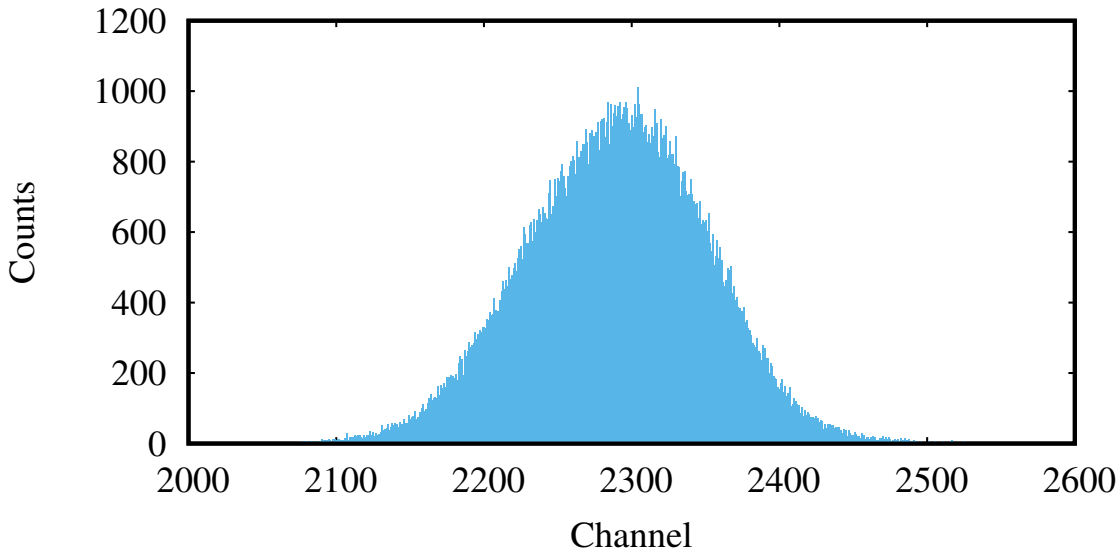

Figure 5.4: Typical time spectrum acquired of one of the Flight Time Monitors.

Sample size	146,002
Mean [channels]	2290.3
Standard dev. [channels]	79
Error of the mean [channels]	0.2
Mean [ns]	55.915
Standard dev. [ns]	1.93
Error of the mean [ns]	0.005

Table 5.1: Statistical quantities of interest from a typical FTM data set.

The bias of the MCP portion of the FTMs is typically automatically adjusted (via the click of a button) to maintain the count rate in the MCP to a value between 1,000 and 10,000 Hz. Thus a typical acquisition times on the order of 10 seconds at 5 kHz results in 5×10^4 data points. Given the largest standard deviation in time typically observed of 2 ns as in this case, this gives a standard error of the mean as calculated with

$$\sigma_{\bar{t}} = \frac{\sigma_t}{\sqrt{n}} \tag{5.1}$$

of approximately 0.01 ns. Measurements closer to the end of the linac typically have σ_t closer to 0.5 ns giving a standard error of the mean of approximately 0.002 ns. These are one or two orders of magnitude smaller compared to other sources of error to be discussed next, and thus the uncertainties related to the sample sizes are considered as insignificant in the total uncertainty of each FTM measurement.

To investigate possible unknown sources of uncertainty on the FTM time measurements, additional measurements were carried out for two different scenarios - one with the DSB buncher and all ISAC-II cavities off, and one with the DSB buncher and all available ISAC-II cavities on. For each scenario two data sets were collected - one with the FTM positions in a static position, and one with the FTM positions being varied. The results for the last FTM are summarized below in Table 5.2 which show a significant deviation.

Measurement	$\sigma_{\mu_{10}}\ [ns]$	$\sigma_{\mu_{12}}\ [ns]$	$\sigma_{\mu_{20}}\ [ns]$	$\sigma_{avg}\ [ns]$
Linac off, static FTMs	0.3	0.3	0.4	0.33
Linac off, moving FTMs	0.7	0.7	0.8	0.77
Linac on, static FTMs	0.02	0.03	0.07	0.04
Linac on, moving FTMs	0.07	0.17	0.09	0.11

Table 5.2: Deviation of FTM measurements under various conditions.

The source of the deviations between measurements is not known with certainty. It is clear that there is both a source of uncertainty with the FTMs left in static positions, as well as a more significant uncertainty when the monitors are moved. One theory is that the beam is coupled between the transverse dimensions x and y, which is possible with a mis-matched beam injected into the linac, as well as coupling between y and energy, caused by the phase-dependent vertical kicks that result from the asymmetry in quarter wave cavities. The FTMs are operated by moving the position of the $50\mu m$ Tungsten wire in the beam to maximize the count rate on the MCP - implicit in this method is that there are no correlations between the time spectrum and other spatial components. Basically, because the wire only intersects a small portion of the beam, estimated to be around 5%[28], it is assumed that this

small portion of the beam is representitive of the full structure of the beam, which is not necessarily true with unexpected correlations or aberrations in the beam.

While further investigations into the causes of these uncertainties present an interesting topic, they are outside the scope of this thesis. The uncertainty in the mean time measured by each FTM for the remainder of this chapter is estimated using these values, focusing on the case studies of all cavities off or all cavities on. It is suspected that the error associated with FTM measurements of intermediate configurations (a subset of all cavities on) fall between the errors estimated above.

It is also prudent to note that the correlations between individual FTM measurements in the last three cases of Table 5.2 are quite small. For the first case (Linac off, static FTMs) in Table 5.2, the last acceleration the beam sees is in ISAC-I, approximately 53 metres upstream of the FTMs. In this case, small variations in output energy from the DTL in ISAC-I are observable in the FTM measurements, which all share this common source of variation. The correlation coefficients in this case reflect this and are $\rho_{\mu_{10-12}} = 0.94$, $\rho_{\mu_{12-20}} = 0.99$, and $\rho_{\mu_{10-20}} = 0.93$.

As discussed in Chapter 2, the flight time between a pair of FTMs is related to the energy per nucleon by

$$E_{jk} = mc^2 \left[\frac{1}{\sqrt{1 - \left(\frac{D_{jk}}{c\Delta t_{jk}}\right)^2}} \right] \tag{5.2}$$

We can calculate the variance of the energy per nucleon with

$$\sigma_E^2 \approx \left|\frac{\partial E}{\partial t}\right|^2 \sigma_t^2 + \left|\frac{\partial E}{\partial D}\right|^2 \sigma_D^2 + 2\frac{\partial E}{\partial t}\frac{\partial E}{\partial D}\sigma_{tD} \tag{5.3}$$

Which if the distance D is not a random variable, or if the variance of the random variable D can be shown to be sufficiently small in comparison to σ_t, reduces to

$$\sigma_E^2 \approx \left|\frac{\partial E}{\partial t}\right|^2 \sigma_t^2 \tag{5.4}$$

The distances between the FTMs is a possible source of systematic error as well as random error, but they are indeed by comparison quite small. The possible systematic error in the distances would arise from an issue with the measured distances during installation - as stated in [19] these were measured with a laser tracker and measured to sub-mm precision.

A random error in the distance could result from the beam 'zig-zagging' along the beamline. The beam is transported through a total of two quadrupoles between the first and second FTM and eight quadrupoles the second and third FTMs. Assuming the beam centroid begins at half of the beamline aperture radius of 50.8 mm [2] and over some distance makes its way to half the aperture radius on the other side of the axis, the change in distance between two monitors can be approximated as

$$\Delta D \approx n\sqrt{(\frac{D}{n})^2 + R^2} - D \tag{5.5}$$

Where D is the total distance between a pair of monitors, D/n is the distance over which the beam goes from one side of the beam axis to the other, R is the aperture radius, and n is the number of oscillations around the beam axis. For a near worst case estimate between the second and third monitor with $n = 8$, $D = 9$, $R = 0.0254$ the result is $\Delta D \approx 2.3\,mm$. Assuming the mean distance would be the average of D and $D + \Delta D$, and conservatively that the extremes would be at 2σ of the distribution gives $\sigma_D \approx 0.57\,mm$.

Taking partial derivatives of Equation 5.2 we get

$$\frac{\partial E}{\partial t} = -mc^2 \frac{\beta^2}{(1-\beta^2)^{3/2}} \frac{1}{t} \tag{5.6}$$

and

$$\frac{\partial E}{\partial D} = mc^2 \frac{\beta}{(1-\beta^2)^{3/2}} \frac{1}{ct} \tag{5.7}$$

Using this to evaluate the first and second terms in Equation 5.3 shows that indeed the contribution to $\sigma^2_{E/u}$ by the uncertainty in the distance is 2 or 3 orders of magnitude smaller than the contribution from the uncertainty in the flight time, t. Continuing from Equation 5.4 to calculate the approximate uncertainty in the measured energy for the two scenarios (linac off, linac on) yields

[2] The beam centroid can't actually reach the aperture radius as half the beam would be lost in this case.

Measurement	σ_T [ns]	$\sigma_{\Delta T}$ [ns]	$\sigma_{E/u,1-2}$ [MeV/u]	$\sigma_{E/u,2-3}$ [MeV/u]	$\sigma_{E/u,1-3}$ [MeV/u]	$\sigma_{E/u,avg}$ [MeV/u]
Linac off	0.77	1.1	0.03	0.006	0.005	0.004
Linac on	0.11	0.16	0.04	0.01	0.008	0.006

Table 5.3: Deviation of FTM measurements under various conditions.

Where the uncertainty in the flight time between two monitors is $\sigma_{\Delta_T} = \sqrt{2}\sigma_T$ as there is no correlation between the meausurements on pairs of FTMs, since those depend on the position of each FTM as mentioned earlier.

There is one more possible source of systematic error in the FTM measurements - the electronics and timing of the FTM system. The signals from the three monitors must travel cables of the same length, and small discrepencies can introduce systematic errors. Fortunately this was investigated the last time the system was upgraded in 2011 [19] with a calibration laser - which travels at a very well known velocity. It is still possible that the calibration could have changed in decade since this last upgrade and thus this remains a possible systematic error in the analysis.

Control system phase and amplitude

In the fitting of the model to the measured E/u for a given cavity phase ϕ, the control system phase ϕ is treated as the independent variable. Thanks to the calibrations done on the control system phase shifts summarized earlier in Sections 3.2.4 and 4.2.1, the expected uncertainties in the phase values are quite small. The measurements were done with the Hewlett Packard 8508A Vector Voltmeter (VVM) which has a resolution of 0.1 degrees and is specified to have a phase accuracy of +/- 1 degree in the relevant frequency range. In the measurements made the only quantity used was the *difference* between two phases which would eliminate any constant systematic error in the phase readings. The random error was assessed with repeat readings for given control system settings and the rms error was estimated to be less than the VVM resolution of 0.1 degrees.

The possible systematic error introduced by the behaviour of the control system amplitude was mentioned briefly back in Section 3.2.5, which is related to non-linearities in the power meter used by the LLRF system that is maintaining a certain amplitude within a PID loop. For the auto-phase test presented in Section 4.3.2, only 2 cavities required scaling of the amplitude, as the large majority remained at the

same amplitude as used in the calibration tune, eliminating any possible scaling error for those. The two cavities that had their amplitudes adjusted were the 13th and 20th cavities, both of which are still operating well within the linear range of their respective power meters.

To ensure automatic phasing of the linac remains robust and reliable, it is **strongly recommended** that the RF group at TRIUMF adjust or add attenuation to the RF cables upstream of the power meters to ensure the signal remains within the linear range of the metres. As of writing, the RF group has already made this adjustment to twenty of the linac cavities, with adjustments to the remaining cavities expected during an upcoming shutdown period.

Errors in incoming beam properties

Implicit in all of the analysis so far has been the assumption that the energy and arrival time of the beam at the start of the ISAC-II linac are both fixed and do not change over time. Reality is not usually so kind and a number of different observations suggest that both beam properties do in fact vary. These include measurements showing that pickup signal from the cavities in the upstream DTL vary by up to 1 degree in RF phase for 1 degree in temperature in the ISAC-I experimental hall [31] and measurements that show the RF reference signal carried between different LLRF control stations varies by up to 3 degrees peak-to-peak over a 24 hour period [32].

The ISAC-I experimental hall contains the ISAC-I accelerator chain which serves as the input into the ISAC-II linac. The ISAC-I hall is a large building with a volume of approximately $3 \times 10^4\ m^3$ and is not rigorously temperature-controlled. It houses the reference RF cables travelling from each RF station to the next before reaching ISAC-II. It also houses the LLRF computers and associated electronics including LLRF boards and more. It also houses the power amplifiers for all ISAC-I room-temperature cavites.

Ambient air temperature in the ISAC-I hall can vary by as much as 10 degrees Celcius (peak to peak) in a 24 hour period in the summer, and by as much as 15 degrees Celcius (peak to Peak) between the hottest and coldest points in the year.

To try and quantify the possible changes in beam properties arriving from the ISAC-I accelerator chain, a beam was tuned through to the ISAC-II FTMs with all ISAC-II RF off. Shown below in Figure 5.5, the arrival time of the beam on the FTMs was measured approximately once per minute over a span of 30 minutes to observe any changes.

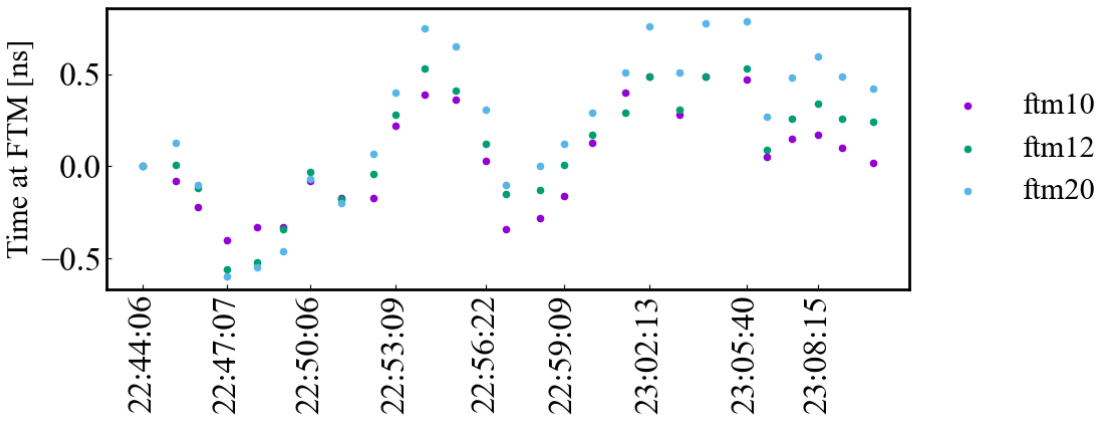

Figure 5.5: Change in relative arrival time at FTMs over 30 minutes.

It was observed that the meausued arrival time at all three FTMs seems to vary together with a fairly strong correlation - ruling out independent errors in individual FTM measurments due to position. The varying arrival time could be attributed to either a real affect of the beam varying in energy from the upstream DTL, from drifting RF signals (either reference signal for ISAC-I or the trigger for the FTMS), or a combination of the two. The FTMs use the RF reference signal as a trigger, but this signal travels along a significant distance in the temperature varying ISAC-I experimental hall. If the phase of the reference signal to the FTMs is varying, it would be reasonable then to also assume that the phase of the refence signal to the ISAC-II buncher and cavites is also changing with respect to the phase of the ISAC-I system. This would result in a similar effect as a real change in the beam energy - the phase of the linac would be shifted and relative to the linac the incoming phase of the beam would be changed. It is **strongly recommended** that as a future impovement, an independent check is carried out on the stability of the trigger signal used for the FTMs, as well as the phase of the ISAC-II LLRF system with respect to the ISAC-I system.

The arrival time varied by as much as 1 ns peak to peak, and had a standard deviation of 0.3 ns. A second test over a longer time period of 24 hours resulted in changes of up to 2 ns peak to peak with a standard deviation of 0.9 ns although the FTMs in this test were not as strongly correlated.

Under the assumption that these changes in arrival time at the FTMs were a result of changes in the energy of the beam exiting the DTL, one can estimate the

energy and arrival time uncertainties at the start of the ISAC-II linac and these are shown below in Table 5.4.

	$\sigma_{T,FTM}$ [ns]	$\sigma_{E/u}$ [%]	$\sigma_{E/u}$ [keV/u]	$\sigma_{T,CAV1}$ [ns]	$\sigma_{T,CAV1}$ [deg @ 106 MHz]
30-minute	0.3	0.03	0.4	0.2	6.5
24-hour	0.9	0.08	1.3	0.5	20

Table 5.4: Uncertainties in incoming beam properties

The uncertainties in energy for both cases are well less than the precision of both the TOF measurement as well as the energy diagnostic (an energy dispersive dipole followed by a profile monitor) for the upstream DTL. Thus the change in energy is difficult to verify independently of the arrival time at the FTMs. However, as a result of the approximately 30 metre drift from the DTL to the first superconducting cavity in ISAC-II, paired with the relatively short RF period of 9.43 ns, these effects would produce a significant change in the phase of the first cavity. These kind of changes have indeed been observed in the past after manual re-phasing of the first cavity after hours or a day of beam being on standby. Note that these observations would also be consistent with the phase of the ISAC-II RF reference signal varying.

With many ISAC-II cavities on one would expect the buncher and cavities to compensate for small energy deviations of the input energy and the beam would oscillate around the reference particle's trajectory. The result is small changes in the arrival times on the FTMs which are included in the uncertanties already summarized earlier in Table 5.3.

These uncertainties can be used now for Monte-Carlo simulations of the linac autophase to follow later in Section 5.4.

Linac model

A number of possible systematic errors could exist within the model of the linac including:

- The longitudinal electric field maps of the cavities.
- `TRANSOPTR's` calculation of the time and energy of the beam centroid via Runge-Kutte integration.

- Positions of cavities.

The cavity field maps were generated with CST Microwave Studios as noted earlier in Section 3.2.3. CST Microwave Studios is widely used around the world and any errors would likely have been due to the input cavity dimensions. However as the field maps are normalized and then their strength calibrated use the beam data, the integrated longitudinal effect of the cavities are accurate to the level of uncertainty on the cavity amplitude parameter C.

`TRANSOPTR` is widely used at TRIUMF and it has been benchmarked against a number of other envelope and particle tracking codes. The comparison shown in Section 5.1 shows good agreement between `TRANSOPTR` and Trace 3-D, including the transverse envelope in Figure 5.2, which shows that the linear effects of RF defocusing agree well between the two models, which in `TRANSOPTR` is extracted from the on-axis electric field.

The longitudinal positions of the cavities is an interesting possible source of error. An error in the distance between one cavity and the next could result in an error in the cavity phase offset parameter θ of all downstream cavities. However, as the overall length of the linac is known quite well, the most likely errors are shifts of individual cavities or cryomodules with respect to their design position. This means that the error in the distance between cryomodule 1 and 2, d_{12} would be negatively correlated to the distance error between cryomodules 2 and 3, d_{23}. So the net result would be a small error in the phase of a single cavity or cryomodule while leaving the error of further downstream cavities quite small. An interesting avenue of research in the future would be either detailed alignment mesurements using a laser tracker system, or to attempt to accurately measure this inter-cavity or inter-module distance using the beam as a diagnostic, as explored at other labs [33].

Errors in Calibrated Linac Parameters C, θ

The calibrated cavity parameters C and θ presented earlier in Table 4.2 resulted from a least squares fit to 10 observed phase and energy data points for each cavity with `TRANSOPTR`, using the internal Nelder-Mead optimizer. `TRANSOPTR` as a reminder is integrating through the on-axis electric field as shown in Equation 3.8. The same data was also fit to the simple TTF approximation (Equation 2.18) using SciPy's `curve_fit` and `least_squares` methods which were using a Trust Region Reflective algorithm. The result of these fits and residuals plots for two example cavities are

shown below in Figures 5.6 and 5.7.

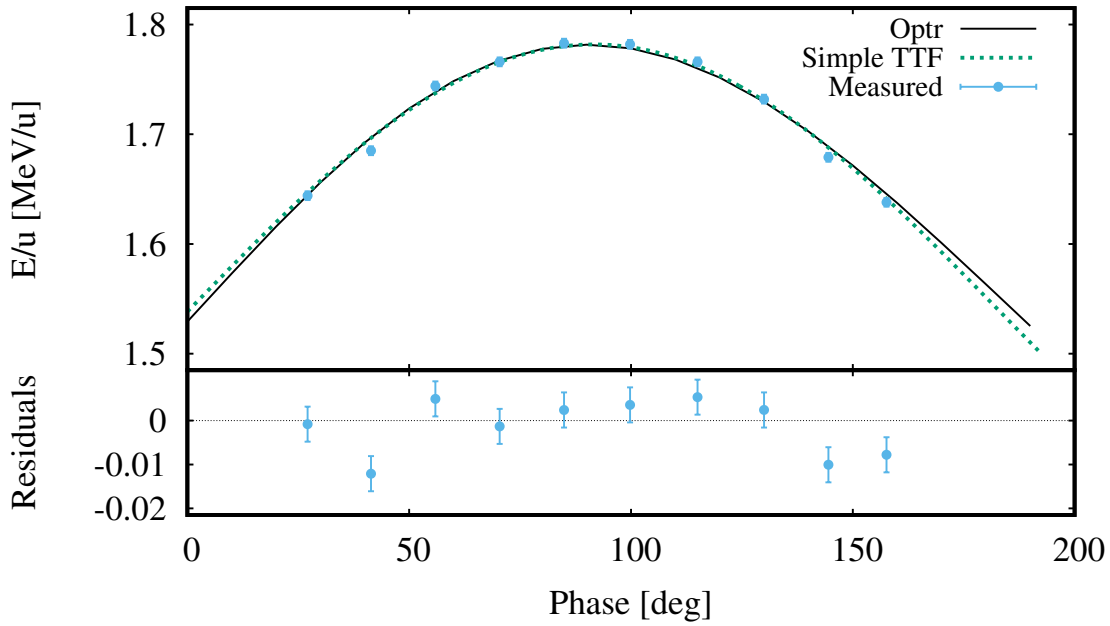

Figure 5.6: Measured energy compared to two different fits for linac cavity number 1 using data for $18O^{4+}$ from May 2021. Residuals are between measured energy and the TRANSOPTR model.

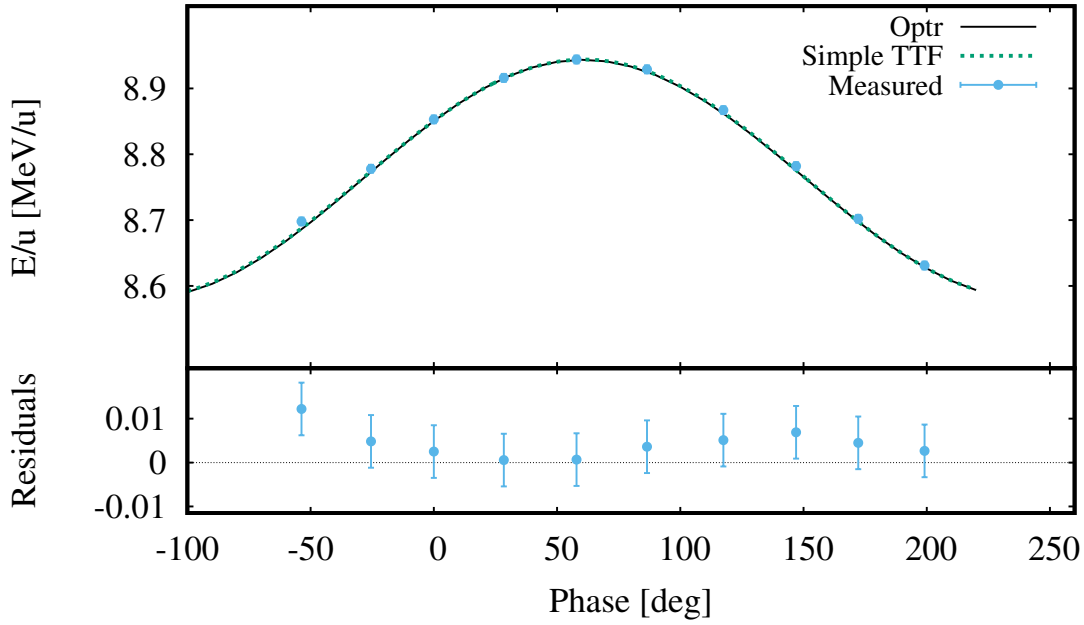

Figure 5.7: Measured energy compared to two different fits for linac cavity number 40 using data for $18O^{4+}$ from May 2021. Residuals are between measured energy and the `TRANSOPTR` model.

SciPy's `curve_fit` method also returns an estimate of the covariance matrix of the fit parameters using numerical methods to estimate the curvature of the objective function at the solution. The resulting 1-sigma uncertainty of the parameters for the simple TTF approximation are shown below in Table 5.5, and these can be considered as an upper limit on the uncertainty resulting in the fit to the more accurate `TRANSOPTR` model.

Cavity	σ_C [%]	σ_θ [deg]	$\rho_{C\theta}$
1	3.0	0.6	0.06
40	1.5	0.7	0.14

Table 5.5: Uncertainties in fit cavity parameters C and θ

These uncertainties are fairly representive of the of the remainder of cavities, which average $\sigma_C = 1.4\%$ and $\sigma_{theta} = 0.7°$. These uncertainties can now be used in Monte-Carlo simulations of the expected variance in energy along the linac for an automatically phased tune, as will be done in Section 5.4.

An attractive avenue for reduction in these uncertanties is to increase the number of data points collected during calibration and to have the cavity phase span over one full RF period. A detailed test was done to try and investigate if and to what extent the error on C and θ could be reduced such as to improve the automated phasing of the linac. In August of 2022 a beam of Carbon-12 at a 2+ charge state was used to re-investigate the uncertainties in calibration of the first few linac cavities. A total of 41 data points were collected for each cavity, and resulting fits for the first linac cavity are shown below in Figures 5.8 and 5.9.

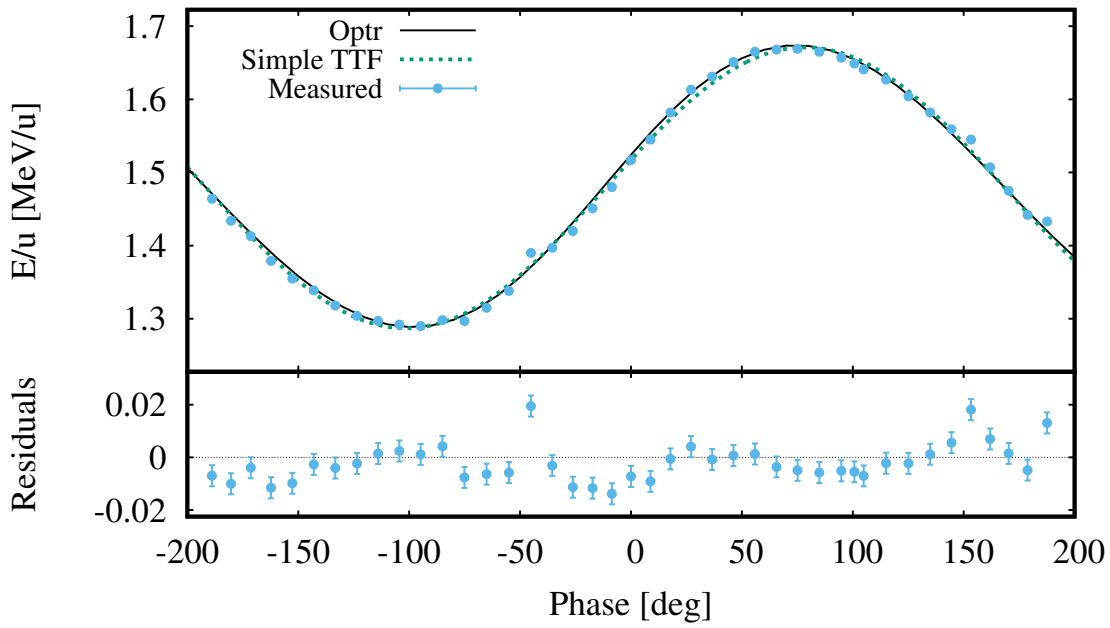

Figure 5.8: Detailed E vs ϕ fit using data for $12C^{2+}$ from August 2022. Residuals are between measured energy and the `TRANSOPTR` model.

The resulting uncertainties in the parameters from the E/u vs ϕ fit above were $\sigma_C = 0.5\%$ and $\sigma_\theta = 0.3°$, notably improved from the fits run on data sets with only 10 points as expected. Figure 5.9 uses a viable alternate method for cavity calibration - using ΔTOF vs ϕ on a single TTF monitor rather than E/u vs ϕ. In principle, running a least squares fit to find the unknown cavity parameters should yield the same result when using ΔTOF as the measured variable instead of E/u (as change in flight time is a direct result of the energy per nucloen). The ΔTOF method is beneficial in that it utilizes only a single FTM, which can be chosen to be the monitor closest to the linac exit (FTM10) minimizing the effects of potential losses along the

post-linac transport section. This method however is more prone to errors in the distance between the accelerating cavity and FTM, as well as input energy.

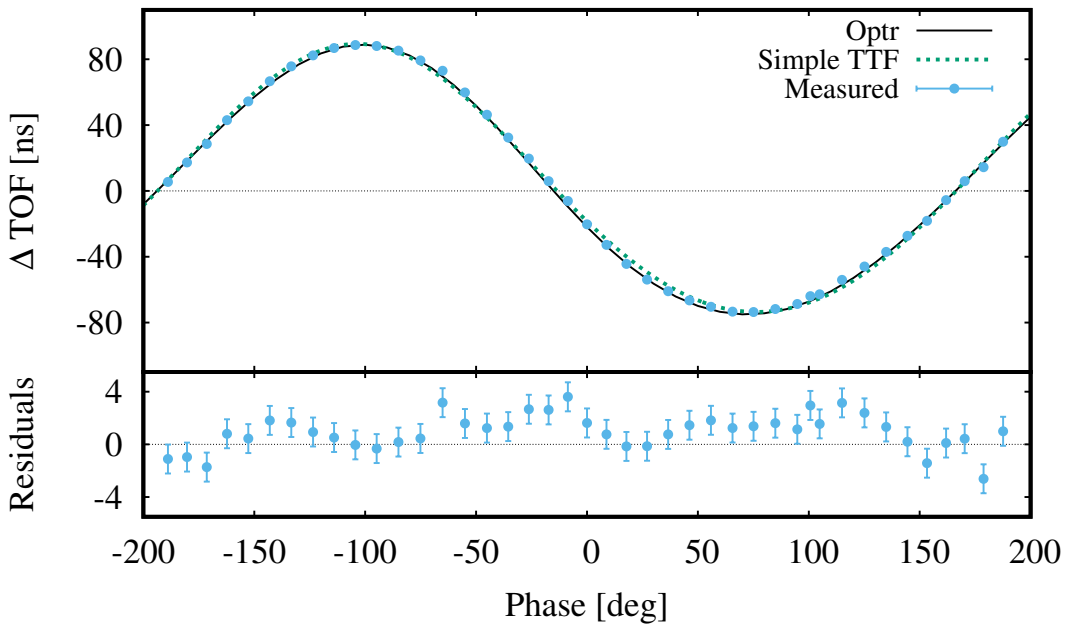

Figure 5.9: Detailed E vs ΔTOF fit using data for $12C^{2+}$ from August 2022. Residuals are between measured energy and the `TRANSOPTR` model.

While promising, the data collection for these cavities was quite time-intensive and required repeated optimization in the model to set the transverse optics to ensure efficient beam transport to the FTMs for the full span of the RF phase. To carry out future linac calibrations with more data points such as this example, the data collection process will need to be nearly fully automated.

These estimates of the errors of the fit cavity parameters can now be used to simulate expected results - how far do we expect the energy profile along the linac to vary given the estimated errors, and is the observed behaviour in the autophase test from Section 4.3.2 within the estimated uncertainties?

5.3 Discussion of Rephasing

Why rephasing is needed

Rephasing of independently controlled cavities is a highly desired capability inves-

tigated at various accelerator facilities. It is essentially synonomous with automatic phasing discussed in more detail in this thesis, but to be precise is a specific subset of automatic phasing where the incoming beam properties (mass, charge, energy, arrival time) are kept fixed, but the available accelerating field in one or more cavities has been reduced either partially or completely to 0. This typically happens when a cavity fails for a variety of reasons, often outside of regular business hours when there are no RF personnel on site. In these circumstances if the on-shift operator is unable to contact the on-call RF person, or if the RF expert is unable to fully restore the cavity, it is benefical to utilize the flexibility of the machine to re-configure around the new boundary conditions.

Given the flexibilty it provides for operation of an accelerator with many independently controlled cavities, rephasing is an interesting area of research toward high-reliability machines with redundancy, including Accelerator Driven Subcritical Reactors.

Time required for rephasing

The time required for the algorithm to re-phase the ISAC-II linac after failure of a single cavity depends on the number of cavities in use and the location of the failed cavity within the linac, but in general ranges in the same time scale as a full automated phasing of the linac. This typically ranges between one and two minutes. Given the expected accuracy with the current uncertainties this will be followed by a check and fine-tuning of the output energy by an operator prior to resuming beam delivery to an experiment, which could take up to 30 minutes. Even still, this is a vast improvement over other options such as manual re-phasing which can take as long as 10+ hours as demonstrated in the past.

Expected accuracy

As mentioned earlier in Section 4.3.3, with the model that is now in-hand following this work, the accuracy of the output energy is expected to be within 2% of the user-requested value for any configuration, and often will be much closer. Except for a few rare occurances, this will typically put any final corrections to be made firmly within the last accelerating cavity in use - a task that is (compared to a full manual tuning of the linac) quite easy for operators, and also reasonable automate fully.

Future considerations for rephasing

There are a number of very interesting ways the rephasing approach could be improved and expedited in the future, including:

- Segmenting the linac into smaller sections during optimization. Currently the full linac is modelled at once, meaning the code is integrating through every cavity and solenoid each time it runs. Thus while the optimization of the last cavity is running many times to find an optimal value it is also re-integrating the upstream cavities unnecessarily each time.

- Running cavities *intentionally below* their stable operating limit, such that cavities surrounding a failed cavity can be increased to ensure that the time-of-flight to as many downstream cavities as posible is kept unchanged, meaning for example only the amplitude and phase of 4 cavities need to change instead of 30 phases downstream.

- Using the drift-kick model as a first approximation prior to use of the more resource-intensive `TRANSOPTR` optimization. The drift-kick calculation is very fast (on the order of 5 ms on my desktop) and could be used to quickly calculate approximate phases. These could then be input into `TRANSOPTR` which would give a more precise output energy that could then be fine-tuned.

- Running the rephasing routine as a self-diagnostic of the linac and model - dropping single cavities and different combinations of cavities and measuring the output compared to model prediction for each. This will likely give some hints at which cavities are most accurately modelled and which may need more scrutiny.

5.4 Analysis of Autophase Data

With the cavity parameter errors estimated in the previous section, the resulting effect on automated phasing can be simulated. A Monte Carlo method was used to evaluate the expected uncertainty of the energy and arrival time of the beam at each cavity, based on the statistical uncertainties of the cavity amplitudes and phases. Uncertainties in incoming beam properties are also included in the simulation in some cases.

All plots and simulations in this section were done for the same beam as presented previously in Section 4.3.2, Argon-40 at a charge state of 11+ with 35 linac cavities

in use. For each of the cases presented below, 1000 trials were run. For each trial, the stated uncertainties are randomly applied to the fit cavity parameters from Section 4.2.2 under the assumption that they are distributed in a normal distribution. Next, the arrival time and output energy of each cavity are recorded. Once complete, the mean and 2rms of the energy and phase at each cavity over the 1000 trials are calculated and plotted for comparison with the measured data.

In the first case, the incoming beam energy and arrival time were assumed to be known precisely, and the amplitude and phase of each cavity were varied by values shown below in Table 5.6, which were approximate averages of the cavity parameter uncertainties summarized previously in Table 4.2.

	σ
Input energy	0.0 keV/u
Arrival time	0.0 ns
Cavity amplitudes	2.0%
Cavity phases	0.5 deg

Table 5.6: Uncertainties used in MC simulation shown in Figure 5.10

The resulting uncertainty in the energy and phase at each cavity is shown in below in Figure 5.10.

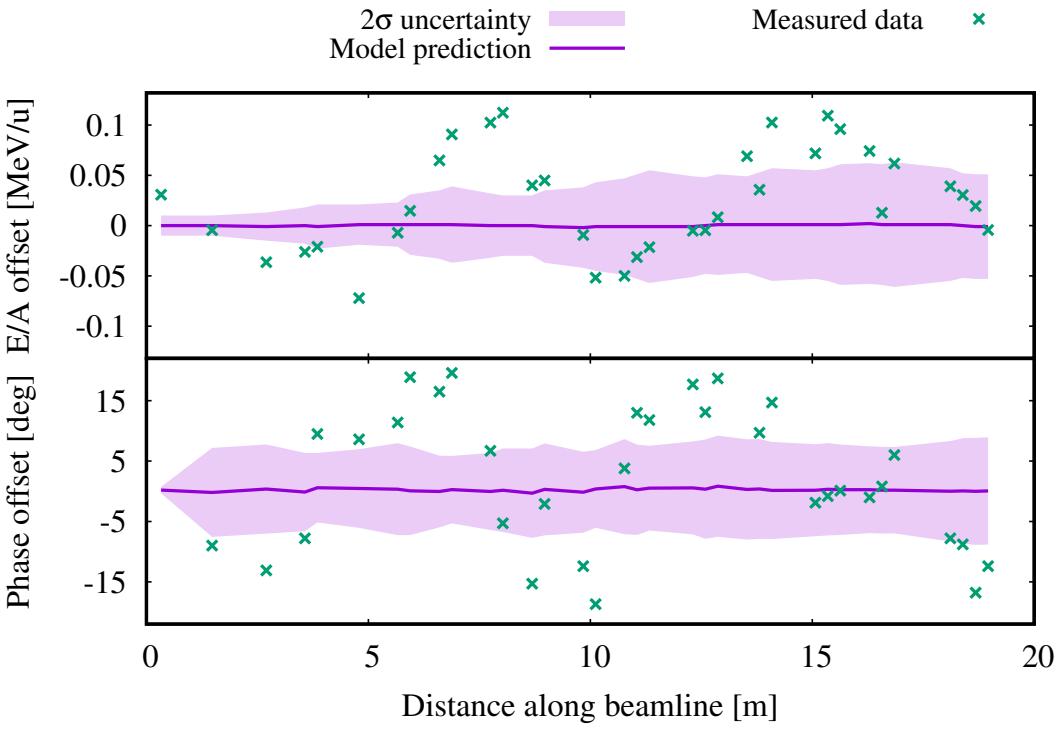

Figure 5.10: Measured offset in cavity phase and energy (relative to the model-predictions) compared to MC simulation for parameters given in Table 5.6

It is qualitatively quite clear immediately that the uncertainties used in this simulation are very unlikely to explain the observed behaviour (in green). The measured phase and energy error at each cavity are regularly well outside of the simulated uncertainty. If the model includes all important source of uncertainty, one would expect the measured offsets in green to mostly fall within the expected model uncertainty in purple.

To attempt to quantify the result, the 1000 trials were sorted based on the number of measured energy data points that are outside of the simulated mean +/- 2rms for each respective cavity. Focusing on the energy data alone, only 1 of 1000 trials had as many points or more outside of the simulated mean +/- 2rms region as the measured data, suggesting it is either a very anomalous result or that some sources of error are missing or need to be revisited.

Indeed, as discussed in Section 5.2, there are uncertainties in the arriving energy and phase of the beam at ISAC-II. Uncertainties estimated back in Table 5.4 were used alongside the independent statistical uncertainty in the fit cavity phase and amplitude parameters shown below in a second MC simulation with the result shown

in Figures 5.11 and 5.12.

	σ
Input energy	0.4 keV/u
Arrival time	0.2 ns
Cavity amplitudes	Various from Table 4.2
Cavity phases	Various from Table 4.2

Table 5.7: Uncertainties used in MC simulation shown in Figure 5.11

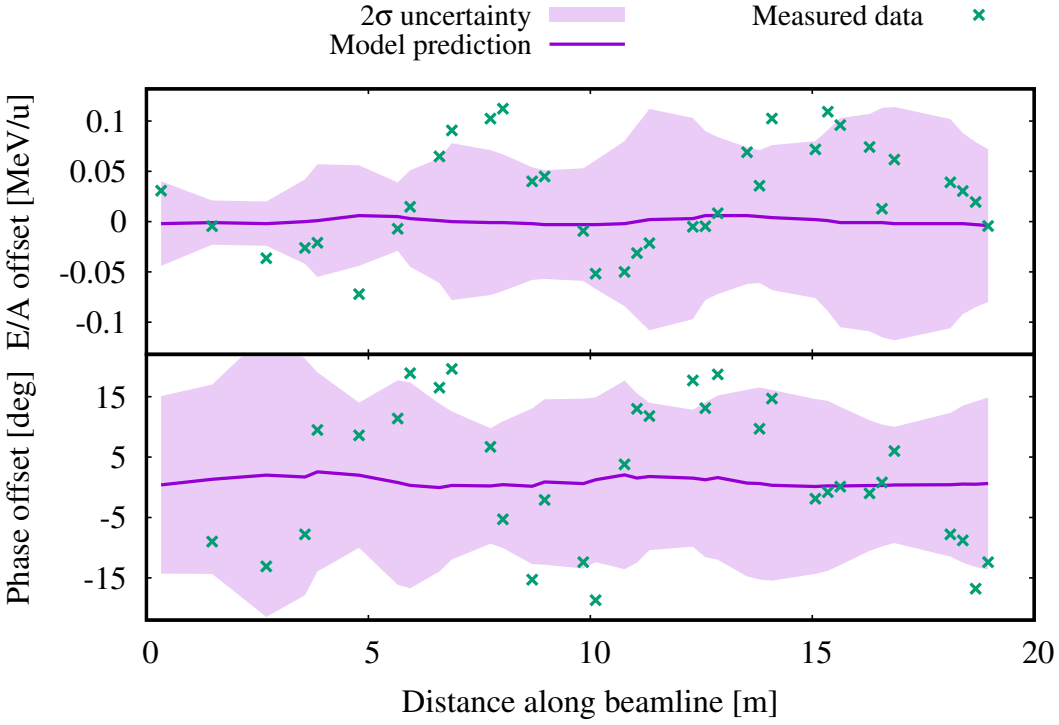

Figure 5.11: Measured offset in cavity phase and energy (relative to the model-predictions) compared to MC simulation for parameters given in Table 5.7

Qualitatively, the result is more in line with expectations. As with the first MC case, the 1000 trials were sorted based on the number of points that land outside of the simulated mean +/- 2rms. Again focusing on the energy data, in this case 17% of trials had as many points or more fall outside the simulated mean +/- 2rms as the measured data. This is not the most likely result but is not particularly abnormal either. Ideally one would repeat the experiment many times, measuring the phase and

energy error through the linac repeatedly and comparing the distribution of results with the distribution generated by the MC simulation. This is an interesting research path that can now be pursued thanks to the work done in this thesis. Each repeat of the experiment would need a new independent calibration of the linac (8+ hours) followed by an independent assesment of the accuracy of an automatic phase based on this calibration (8+ hours). Repeating 1000 trials would thus take close to 2 years if run continuously which is not currently possible.

Stepping back to view the uncertainty in energy gain and measured energies relative to the full energy gain of the machine, one can see that the energy uncertainty relative to total energy from using the model-based phasing approach remains quite small throughout the linac. Figure 5.12 below shows the model predicted energy with uncertainty compared to the measurements as done in Figure 5.11 over the full range of energies achieved through the linac, rather than just residuals. The 2rms relative energy uncertainty through the linac ranges between 0.4 and 1.2%, and is (in the case of all cavities in use) 0.4% at the linac exit.

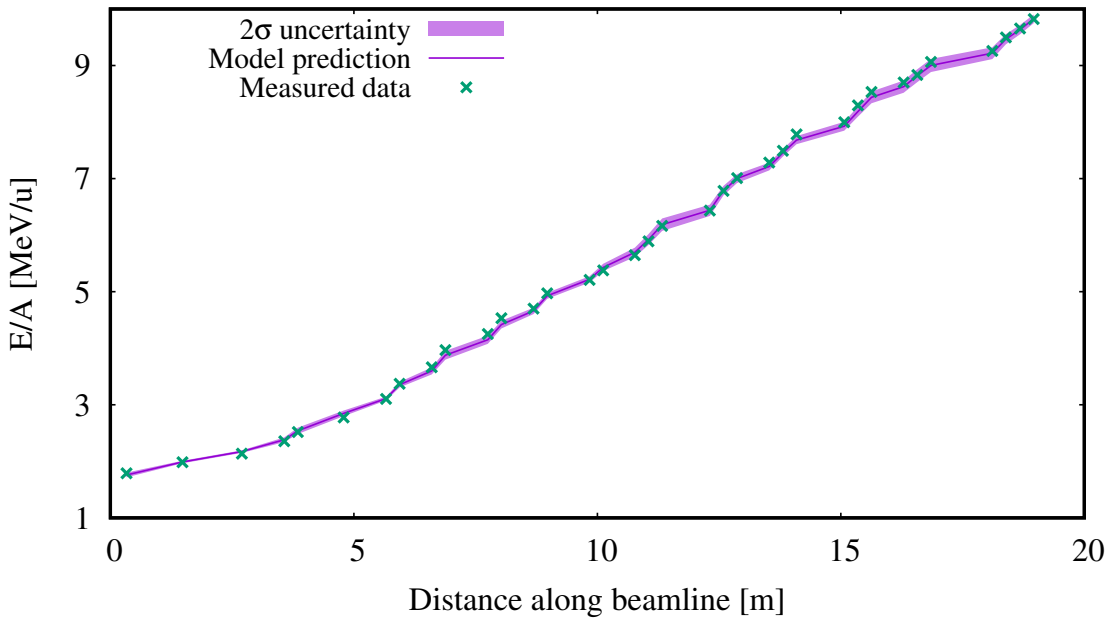

Figure 5.12: Measured offset in cavity phase and energy compared to MC simulation for parameters given in Table 5.7

To investigate avenues for futher improvement, the MC approach was repeated for a few other scenarios to identify which parameters were contributing most significantly. First, the same uncertainties as used above in Table 5.7 and Figure 5.11

were repeated with a single change of dropping the phase error to 0. Second, the same uncertainties again were repeated this time instead dropping the input energy uncertainty to 0. Both cases results in negligible change to the resulting energy and phase uncertainties from those plotted in Figure 5.11, suggesting that the plotted uncertainties are dominated by the error in the arrival time of the incoming beam and the cavity amplitudes. This can be easily shown analytically by taking the simple approximation of the energy gain:

$$\Delta E/u = Acos(\phi_{rel}) \tag{5.8}$$

Where I have assumed a fixed charge and input energy and substituted $A = \frac{q}{m}\mathcal{E}_0 TL$ for simplicity. Error propagation yields

$$\sigma^2_{\Delta E/u} \approx \left|\frac{\partial \Delta E/u}{\partial A}\right|^2 \sigma_A^2 + \left|\frac{\partial \Delta E/u}{\partial \phi}\right|^2 \sigma_\phi^2 \tag{5.9}$$

Assuming no correlation between A and ϕ. This gives

$$\sigma_{\Delta E/u} \approx A\sqrt{\cos^2\phi\left(\frac{\sigma_A}{A}\right)^2 + \sin^2\phi(\sigma_\phi^2)} \tag{5.10}$$

Plotting this for a typical cavity with $A = 0.2\, MeV/u$, $\sigma_\phi = 0.5\, deg$, $\frac{\sigma_A}{A} = 0.02$ while varying the cavity phase, one gets the result shown below (Figure 5.13).

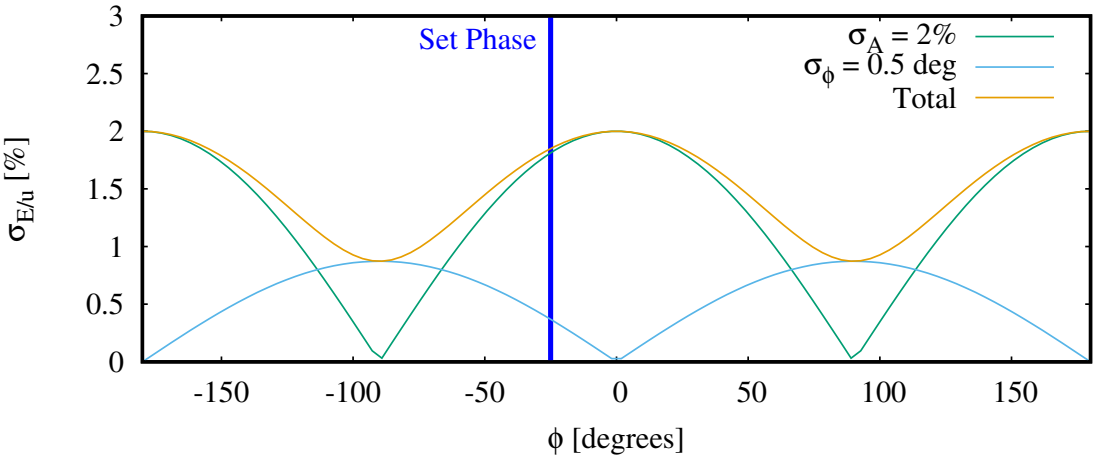

Figure 5.13: Approximate resulting uncertainty in the energy/nucleon gained in a single cavity for varying cavity phases

At the typical operational phase of -25 degrees the error in the output energy is

dominated by the uncertainty in the amplitude. Looking instead at the uncertainty in the total energy out of a cavity resulting from only the uncertainties in arrival time and input energy (assuming no uncertainty in cavity amplitude and phase) we get

$$\sigma_{E/u} = \sqrt{\sigma_{Ein}^2 + A^2(\sin^2 \omega t)\omega^2\sigma_t^2} \tag{5.11}$$

Which evaluated for $E_{in} = 1.50\ MeV/u$, $\sigma_t = 0.2\ ns$, $\sigma_{Ein} = 0.4\ keV/u$, $A = 0.2\ MeV/u$ gives a $\sigma_{E/u}$ strongly dominated by σ_t.

Using a detailed phase-energy scan of the first cavity prior to autophase, or with a time domain diagnostic sufficiently close and of sufficently high resolution, it is believed that in the future we could reach an accuracy of 1 degree (at 106 MHz, or 0.026 ns) in the arrival time of the ion beam. This paired with more precise parameters estimated as shown in Figures 5.8 and 5.9 would give very significant improvements to the expected energy and phase uncertainties of an automatically phased linac, as shown below in Table 5.8 and Figure 5.14.

	σ
Input energy	1.5 keV/u
Arrival time	0.026 ns
Cavity amplitudes	0.5%
Cavity phases	0.3 degrees

Table 5.8: Uncertainties used in MC simulation shown in Figure 5.14

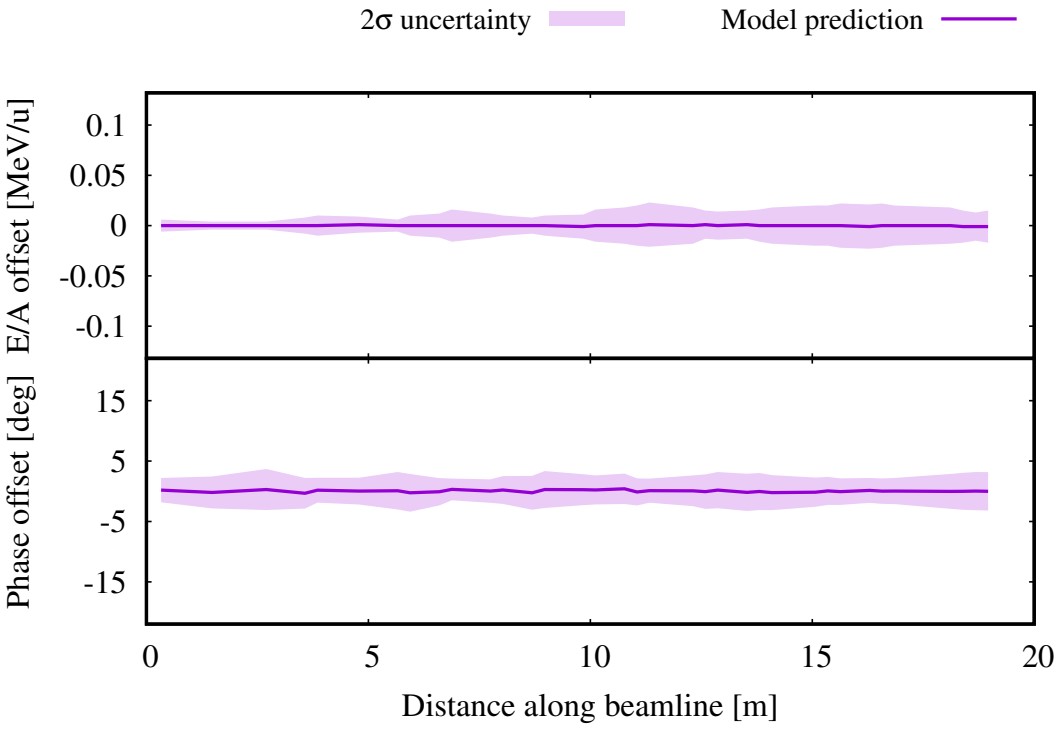

Figure 5.14: Monte Carlo simulation of linac energy and phase uncertainty for parameters given in Table 5.8

At this level of accuracy, the uncertainties of the energy and phase of the beam's centroid are expected be smaller than the statistical 2rms z and z' envelopes describing the spread of the individual particle's momenta and phases. To reduce uncertainties to this level, a few suggestions are made:

1. The cause of observed variations in arrival time at the ISAC-II FTMs must be identified. This should start with checks of how the RF reference signal is varying over both short (on the order of 10 minutes) and long (on the order of 24 hours) time periods.

2. A time-domain diagnostic at the start of the ISAC-II linac, as close as posible to the first cavity, should be used to measure the beam's arrival time and used as an input to the model.

3. Additional time-domain diagnostics between the cryomodules could greatly aid in both linac calibration and automated tuning. Even a single device in the centre of the linac would provide an opportunity to re-set the model at this location and make the new approach more robust.

Some of these suggestions are similar to what is currently done at other labs. Appendix A describes some of these and elaborates on other interesting differences.

Chapter 6

Conclusion

An understanding of the physics of linear accelerators, including beam transport elements and accelerating cavities, allows accurate simulations of modern accelerators like the ISAC-II superconducting linac at TRIUMF. This thesis has demonstrated a significant improvement in the understanding and control of the ISAC-II superconducting heavy ion linac at TRIUMF (Observations and Suggestions summarized in Appendix B), leading to improved reliability of the machine, and reduced overhead. The time to setup all available linac cavities for a given experiment has been reduced from 8 to 16 hours down to approximately 10 minutes.

The ISAC-II linac has been modelled in `TRANSOPTR`, and the model has been benchmarked to previous simulations done in Trace 3-D as well as measurements from the physical machine. Cavity parameters related to the field strength and phase have been calculated using beam measurements and included in the model, and an algorithm to use the model to calculate necesary settings for a desired output has been implemented and tested.

The claims made in the introduction to this thesis have been successfully demonstrated:

1. The measured transverse beam size at available diagnostics had qualitatively good agreement with envelope simulations, with the exception of 4 locations which require further investigation due to possible x-emittance growth. This is a prime example of a way in which the newly established trust in model predictions can lead to investigations and discovery of possible hardware issues.

2. In the autophase test of the linac, the output energy of each cavity was within 2% of the measured value. The phase of a number of cavites did exceed the

claimed goal of less than five degrees between model and measurement. However a possible cause of this has been identified and further investigations suggested in Section 5.4.

3. Final linac transmission for the full automatic phase test with 35 cavities on was 91%. The average transmission measured as each of the 35 cavities were turned on was 95%.

4. Final beam energy from the automatic phase test was within 0.05% of the model-predicted value.

What we can do now

This work has resulted in a wide range of benefits to the ISAC-II facility. Some examples of what we are able do now include:

- **Automated tuning** of the linac during initial setup of a beam through the accelerator. This has many knock-on effects including lower demand on operators, quicker turn-around time for experiments, and greater instituional knowledge.

- **Automated re-phasing**, or fast recovery from cavity failures. The same model used to automate setup of the linac can of course also be used to re-calculate when changes to the linac configuration occur. When a cavity fails and can't be recovered, the model can now be run to quickly re-optimize downstream cavity phases and add a cavity (if necessary) to recover the expected output energy. This can be done in as fast as a few minutes when similar failures in the past have cause 10+ hours of downtime.

- Higher available **total hours of beam** to experiments in the ISAC-II experimental hall, which could mean either longer or more frequent experiments.

- The ability to meet the demands of **more complex experiments**, for example an experiment requesting frequent energy changes. An interesting application of this would be for a regular and semi-automated campaign for charge state distribution measurements - measuring this every time the linac is configured for a new species/energy prior to delivering it to the experiment.

- **Reduced procedural complexity** for those operating the linac - this enables even fairly new operators to setup the accelerator for new experiments or to make changes for ongoing experiments.

- Enabled **monitoring and improving of beam quality**. Having regular measurements of beam emittance entering and exiting the linac and a model in-hand to predict and compare observations with are critical for understanding and improving the quality of the delivered beam. An example of an improvement that can now be pursued is the longitudinal beam envelope - the algorithm could now be altered to choose varying relative RF phases instead of a fixed phase. The algorithm could optimize cavity phases individually to minimize the increase in normalized longitudinal emittance.

The expected reduction in overhead of operating the ISAC-II facility is 1-2 weeks per year, allowing this time to be used to increase beamtime available for various nuclear physics experiments served by the ISAC-II linac. Further to this, the quality of the delivered beam can be more reliably predicted, monitored, and improved. This project has also provided a solid foundation for future studies of automatic tuning techniques and beam dynamics investigations at TRIUMF.

Bibliography

[1] J. A. Bagger *et al.*, "ARIEL at TRIUMF: Science and Technology," in *Proc. IPAC'18*, (Vancouver, Canada, Apr.-May 2018), JACoW Publishing, Geneva, Switzerland, pp. 6–11. DOI: 10.18429/JACoW-IPAC2018-MOXGB2.

[2] C. B. Barquest *et al.*, "Web-Based Control Room Applications at TRIUMF," in *Proc. IPAC'18*, (Vancouver, Canada, Apr.-May 2018), JACoW Publishing, Geneva, Switzerland, pp. 4832–4835. DOI: 10.18429/JACoW-IPAC2018-THPML078.

[3] O. Shelbaya *et al.*, "Autofocusing drift tube linac envelopes," *Phys. Rev. Accel. Beams*, vol. 24, p. 124 602, 12 Dec. 2021. DOI: 10.1103/PhysRevAccelBeams.24.124602.

[4] D. Wang *et al.*, "Accelerator tuning with deep reinforcement learning," in *Proceedings of the Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS)*, (Vancouver, Canada), 2021. [Online]. Available: https://ml4physicalsciences.github.io/2021/files/NeurIPS_ML4PS_2021_125.pdf.

[5] R. Laxdal and M. Pasini, "ISAC-II Optics Specifications," TRIUMF, Tech. Rep. TRI-DN-01-28, 2001.

[6] T. P. Wangler, *RF Linear Accelerators*. Wiley-VCH, 2008.

[7] M. Marchetto, S. Kiy, R. Laxdal, Y. Ma, O. Shelbaya, Z. Yao, and V. Zvyagintzev, "The ISAC-II Linac Performance," in *Proc. HIAT'15*, (Yokohama, Japan, Sep. 2015), JACoW Publishing, Geneva, Switzerland, pp. 175–179. [Online]. Available: https://jacow.org/HIAT2015/papers/WEM2I01.pdf.

[8] R. Kanungo, "IRIS: The ISAC charged particle reaction spectroscopy facility for re-accelerated high-energy ISOL beams," in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, R. K. J. Dilling and L. Merminga, Eds., Springer, 2014, ch. 28, pp. 235–240.

[9] B. Davids, "The electromagnetic mass analyzer EMMA," in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, R. K. J. Dilling and L. Merminga, Eds., Springer, 2014, ch. 26, pp. 225–228.

[10] G. Hackman and C. Svensson, "The TRIUMF-ISAC gamma-ray escape suppressed spectrometer, TIGRESS," in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, R. K. J. Dilling and L. Merminga, Eds., Springer, 2014, ch. 29, pp. 241–251.

[11] F. J. Sacherer, "Rms envelope equations with space charge," CERN, Tech. Rep. DL/70-12, 1970.

[12] E. Courant and H. Snyder, "Theory of the alternating-gradient synchrotron (reprinted from annals of physics, vol 3, pg 1-48, 1958)," *Annals of Physics*, vol. 281, pp. 360–408, Apr. 2000. DOI: 10.1006/aphy.2000.6012.

[13] D. V. Gorelov and P. N. Ostroumov, "Application of LANA Code for Design of Ion Linac," in *Proc. EPAC'96*, (Sitges, Spain, Jun. 1996), JACoW Publishing, Geneva, Switzerland. [Online]. Available: https://accelconf.web.cern.ch/e96/PAPERS/THPG/THP062G.PDF.

[14] D. Gorelov, "Beam Dynamic Simulation with 3-D Realistic Fields using LANA Code," TRIUMF, Tech. Rep. TRI-DN-97-6, 1997.

[15] D. P. Rusthoi, W. P. Lysenko, and K. R. Crandall, "Further Improvements in TRACE 3-D," in *Proc. PAC'97*, (Vancouver, Canada, May 1997), JACoW Publishing, Geneva, Switzerland. [Online]. Available: https://accelconf.web.cern.ch/pac97/papers/pdf/8P089.PDF.

[16] K. Brown, D. Carey, C. Iselin, and F. Rothacker, "TRANSPORT, A Computer Program for Designing Charged Particle Beam Transport Systems," CERN, Tech. Rep. 73-16, 1973.

[17] R. Baartman, "`TRANSOPTR`: Changes since 1984," TRIUMF, Tech. Rep. TRI-BN-16-06, 2016.

[18] R. A. Baartman, "Fast Envelope Tracking for Space Charge Dominated Injectors," in *Proc. LINAC'16*, (East Lansing, MI, USA, Sep. 2016), JACoW Publishing, Geneva, Switzerland, pp. 1017–1021. DOI: 10.18429/JACoW-LINAC2016-FR1A01.

[19] V. A. Verzilov, J. Lassen, R. E. Laxdal, and M. Marchetto, "Upgrade of the ISAC Time-of-flight System," in *Proc. IPAC'11*, (San Sebastian, Spain, Sep. 2011), JACoW Publishing, Geneva, Switzerland, pp. 1147–1149. [Online]. Available: https://jacow.org/IPAC2011/papers/TUPC065.pdf.

[20] M. Marchetto, R. Laxdal, and V. Zvyagintzev, "Beam Quality and Operational Experience with the Superconducting LINAC at the ISAC II RIB facility," in *Proc. PAC'07*, (Albuquerque, NM, USA, Jun. 2007), JACoW Publishing, Geneva, Switzerland, pp. 1392–1394. [Online]. Available: https://jacow.org/p07/papers/TUPAN003.pdf.

[21] D. Lanaia and M. M., "ISAC-II Phasing High Level Application in MATLAB," TRIUMF, Tech. Rep. TRI-DN-13-18, 2014.

[22] R. Baartman, "B-I Curve fits for TRIUMF Quads," TRIUMF, Tech. Rep. TRI-DN-07-4, 2007.

[23] ——, "First Order Fringe Field Effects in TRIUMF Quadrupoles," TRIUMF, Tech. Rep. TRI-DN-07-1, 2007.

[24] J. Nasser *et al.*, "Algorithm to Mitigate Magnetic Hysteresis in Magnets with Unipolar Power Supplies," in *Proc. IPAC'22*, (Bangkok, Thailand), JACoW Publishing, Geneva, Switzerland, Jul. 2022, pp. 156–159. DOI: 10.18429/JACoW-IPAC2022-MOPOST039.

[25] M. A. Fraser, J. C. Broere, S. Haastrup, D. Lanaia, D. Valuch, and D. Voulot, "Understanding the Error Tolerances Required to Automatically Phase the HIE-ISOLDE Linac," in *Proc. LINAC'14*, (Geneva, Switzerland, Aug.-Sep. 2014), JACoW Publishing, Geneva, Switzerland, pp. 496–498. [Online]. Available: https://jacow.org/LINAC2014/papers/TUPP031.pdf.

[26] R.E. Laxdal and M. Marchetto, "The ISAC post-accelerator," in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, R. K. J. Dilling and L. Merminga, Eds., Springer, 2014, ch. 8, pp. 79–97.

[27] M. Marchetto and R.E. Laxdal, "High energy beam lines," in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, R. K. J. Dilling and L. Merminga, Eds., Springer, 2014, ch. 9, pp. 99–109.

[28] V. A. Verzilov, R. E. Laxdal, M. Marchetto, and W. R. Rawnsley, "Time Domain Diagnostics for the ISAC-II Superconducting Heavy Ion Linac," in *Proc. DIPAC'07*, (Venice, Italy, May 2007), JACoW Publishing, Geneva, Switzerland, pp. 247–249. [Online]. Available: https://jacow.org/d07/papers/WEPB07.pdf.

[29] A. Terebilo, "Channel Access Client Toolbox for MATLAB," SLAC, Tech. Rep. SLAC-PUB-9339, 2002.

[30] S. Kiy, R. Baartman, O. Kester, and O. Shelbaya, "First Tests of Model-Based Linac Phasing in ISAC-II," in *Proc. HIAT'22*, (Darmstadt, Germany), JACoW Publishing, Geneva, Switzerland, Sep. 2022, pp. 113–117. DOI: 10.18429/JACoW-HIAT2022-TUP19.

[31] S. Kiy, "LLRF Phase Shifter Calibrations, October 2020," TRIUMF, Tech. Rep. TRI-BN-20-19, 2020.

[32] Z. Y. Yao *et al.*, "Status and Challenges at TRIUMF ISAC Facility," in *Proc. LINAC'22*, (Liverpool, UK), JACoW Publishing, Geneva, Switzerland, Oct. 2022, pp. 866–871. DOI: 10.18429/JACoW-LINAC2022-FR1AA03.

[33] A. P. Shishlo, "Benchmark of Superconducting Cavity Models at SNS Linac," in *Proc. IPAC'21*, (Campinas, Brazil, May 2021), JACoW Publishing, Geneva, Switzerland, pp. 671–674. DOI: 10.18429/JACoW-IPAC2021-MOPAB203.

[34] A. S. Plastun and P. N. Ostroumov, "Instant Phase Setting in a Large Superconducting Linac," (Albuquerque, New Mexico, USA, Aug. 2022), presented at NAPAC'22, Albuquerque, New Mexico, USA, Aug. 2022, paper THZD1, unpublished, [Online]. Available: https://accelconf.web.cern.ch/napac2022/papers/thzd1.pdf.

[35] P. N. Ostroumov *et al.*, "FRIB Commissioning," in *Proc. HIAT'22*, (Darmstadt, Germany), JACoW Publishing, Geneva, Switzerland, Sep. 2022, pp. 118–123. DOI: 10.18429/JACoW-HIAT2022-WE1I3.

[36] V. S. Morozov, C. C. Peters, and A. P. Shishlo, "Oak ridge spallation neutron source superconducting rf linac availability performance and demonstration of operation restoration with superconducting rf cavity off," *Phys. Rev. Accel. Beams*, vol. 25, p. 020 101, 2 Feb. 2022. DOI: 10.1103/PhysRevAccelBeams.25.020101.

[37] A. P. Shishlo, "Model and Beam Based Setup Procedures for a High Power Hadron Superconducting Linac," in *Proc. LINAC'14*, (Geneva, Switzerland, Aug.-Sep. 2014), JACoW Publishing, Geneva, Switzerland, pp. 41–45. [Online]. Available: https://jacow.org/LINAC2014/papers/MOIOC03.pdf.

Appendix A

Comparisons to other labs

A.1 Facility for Rare Isotope Beams (FRIB)

The Facility for Rare Isotope Beams (FRIB) in East Lansing, Michigan is a new facility just starting to produce radioactive beams for experiments. The facility can accelerate a wide range of stable ions all the way up to Uranium through a linac with over 300 independently phase superconducting cavities. These beams strike a target which fragments the primary beam into a variety of secondary beams, from which the desired secondary beam is separated and sent to experiments.

Extensive work has been done at FRIB towards the automatic setting of the cavity phases. As described in [34], commissioning began in 2018 with scans of cavity phase vs arrival time at a BPM (beam position monitor), similar to the historical cavity phasing techniques used at ISAC-II as described in Chapter 2. From there, they moved on to automate the phase scan procedure using a High Level Application (HLA) named ALPha (Automatic Linac Phasing). ALPha steps each cavity through a full RF period in 30 degree steps, measuring the phase difference between two downstream BPMs at each step. A cubic spline is fit to this measured data and evaluated to identify the phase giving the smallest phase difference between the two BPMs.

With over 300 cavities to phase, even with ALPha the time required to configure the full FRIB linac for a new beam took approximately 20 hours [34][35]. To address this, they moved to a model-based phasing approach very similar to that which is utilized in this thesis. Their approach, called Instant Phase Setting (IPS), uses the data collected from the ALPha approach to fit the unknown parameters C and θ

relating to scaling of the cavity field and phase offset respectively. With a calibrated model they then take an incoming beam of known energy and arrival time, and numerically integrate the longitudinal equations of motion (see Table 3.1) to find the new phase of each subsequent cavity and its output energy.

Results presented [34] show successful tests giving errors of less than 0.15 MeV/u between measurement and model for IPS-based phasing of 72 cavities in their second linac segment (LS2). It would be interesting to get some information on the uncertainties of the calibration on the cavity amplitude and phase offsets at FRIB, and the expected resulting uncertainty in their IPS approach. Particularly of interest would be the species and charge state(s) used for the calibration and resulting automatic phase tests, because it is suspected that the uncertainties in this model-based approach increases as the change in time of flight to a given cavity increases.

Considerations for ISAC-II

An automated program like ALPha for collecting cavity phase/energy data for calibration would be a significant asset as a future improvement of the suite of tools used for calibration and operation of the ISAC-II linac. The calibration data collected and presented in this thesis still involved a lot of manual optimization - the time-domain diagnostics at ISAC-II are all downstream of the 40 cavity linac, there are no time-domain diagnostics within the linac itself. Thus, especially near the start of the linac, changes in cavity amplitudes and phases have a substantial impact on both the transverse tune, as well as often introduce transverse kicks to the beam. Automating the necessary steering corrections and calculation and setting of transverse optics would significantly expedite the calibration process in the future.

With this said, the difference in available diagnostics between the FRIB linac and the ISAC-II linac are quite significant. FRIB has 144 BPMs [35] distributed along the linac, which measure both the transverse position and longitudinal phase of the beam. As shown below in Table A.1, in a 'per-cavity' comparison between the two facilities, ISAC-II has notably fewer longitudinal diagnostics. Installation of time-domain diagnostics between each cryomodule at ISAC-II would enable much easier calibration and studies of the linac. It is noted that ISAC-II is primarily a RIB facility, while FRIB is a stable beam accelerator that runs at much higher intensities. Nonetheless, when ISAC-II is operating with stable beams, as done for calibration and beam development, intensities can reach up to 100-500 enA if desired. The Off-line Ion Source (OLIS) and ISAC-I accelator chain have demonstrated the ability

to deliver over 1 euA for common gaseous species like 14N, 16O, and 20Ne. These intensities would be sufficient for use of BPMs, however the cost-benefit trade off would have to be considered as these devices would likely *not* be useful in regular delivery to experiments which run at more modest intensities in the range of 10 enA for stable beams at right down to 1,000's of particle/s for radioactive beams.

	ISAC-II	FRIB	SNS
Cavities	40	328	81
Longitudinal diagnostics	3	144	27
Cavities per Longitudinal diag.	**13.3**	2.3	3.0

Table A.1: Comparison of available diagnostics at ISAC-II, FRIB, and SNS.

With BPMs between each cryomodule, the automatic collection of calibration data would be made much easier - optimization of the transverse tune and steerers would not be required nearly as frequently. Observed drifts in phases[31], which could result in a longitudinal oscillation around the reference particle, could be much more easily identified using BPMs installed within the ISAC-II linac.

Also noteworthy is FRIB's use of digital LLRF boards, which are expected to have phase shifters which have much smaller errors than the analog boards currently in use at ISAC-II. Digital LLRF boards have already started to be installed at TRIUMF and promise to reduce or eliminate the need for the manual calibration of the phase shifters.

A.2 Spallation Neutron Source (SNS) at Oak Ridge National Laboratory (ORNL)

The Spallation Neutron Source (SNS) at Oak Ridge utilizes a pulsed superconducting 1 GeV H- linac [36] as its driver. The H- ions are accelerated initially to 186 MeV through a series of room temperature accelerators including a radio frequency quadrupole (RFQ), drift tube linac (DTL), and coupled cavity linac (CCL). Following this, beam is accelerated through a total of 81 6-cell elliptical superconducting cavities to a final energy of 1 GeV. The first 33 cavities have a design velocity of $0.61c$ while the remaining 48 have a design velocity of $0.81c$. All 81 SC cavities are operated

at 805 MHz. The SNS linac has one beam phase and position monitor (BPM) after each cryomodule, as well as several upstream and downsteam[36]. This puts the SNS in a very similar cavities-per-diagnostic regime as FRIB, as shown earlier in Table A.1.

SNS in many ways is less demanding than FRIB or ISAC-II on the required flexibility of the linac - the mass and charge of the accelerated ions does not change. They do of course have many different challenges that ISAC-II doesn't, including monitoring for beam loss and keeping it low to prevent damage and activation from the high power beam, space charge effects, and handling beam loading of cavities. Due to the different nature of the SNS linac, their main interest is in fast recoveries from cavity failures. As disussed in [36], high reliability is of particular importance in consideration of a proton beam used for possible future accelerator-driven subcritical reactors.

Like ISAC-II and FRIB, the SNS H- linac cavity tuning was intially performed using phase-energy scans [37], with the time of flight between two downstream BPMs giving a measure of the beam velocity (and thus energy since the beam species is known). Initially this took up to 8 hours for phasing of the 81 cavities[37], however with the implementation of their 'SCL Tuning Wizard' within the OpenXAL framework, this was reduced down to 40 minutes by automating the phase energy scans. This model has also been used to retune the SNS linac automatically, resulting in energies within 0.2% of the final total energy of 1 GeV[33]. It is not clear however what conditions these re-phasing exercises were done for - whether this was for a single failed cavity or just some changes to one or more cavity amplitudes. While also not completely clear in the available literature, the available material[33] suggests that the SNS model uses the simple drift-kick approach for re-calculating necessary cavity phases and output energy. The model is calibrated using phase-energy data collected at 25 data points from -180 to +180 at 15 degree intervals[36].

In 2021, investigations into the accuracy of the SNS linac model have shown uncertainties in the calibration of the cavity amplitudes of approximately 2%[33], similar to the estimated uncertainties of the ISAC-II cavity amplitudes in this thesis. They have also carried out some interesting work looking at the cavity position with respect to the longitudinal diagnostics[33] which could be applied to ISAC-II in the future. Most recently, they successfully carried out re-phasing tests of their linac, turning off the 41st cavity, using the model to rephase cavity 42 through 80, and turning on cavity 81 to recover the intial energy[36]. All this was done within 10 seconds,

which was limited by the time it took to move the tuner of the 41st cavity sufficiently far to avoid beam loading of this cavity. Results showed very similar patterns on beam loss monitors throughout the entire linac, suggesting the beam dynamics had changed very little, however the results of the initial and final beam energies were not presented.

Considerations for ISAC-II

As noted earlier, automating the phase-energy scans for ISAC-II, as SNS and FRIB have demonstrated, would be a significant asset. However, it is expected to be more difficult to implement due to the lower beam energy and larger relative energy gain per cavity, the cavity type (Quarter wave resonators, which can have a vertical steering effect), and current lack of orbit correction tool for ISAC-II.

Perhaps most interestingly, observations of anamolous changes to cavity phases have been observed at SNS, similar to those observed at ISAC-II. This included a test where all cavities were automatically phased with phase-energy scans, taking 20 minutes each, then repeating the measurement a total of 4 times[37]. While the phase of the first few cavities of the linac did not change by more than 1 degree, cavities further down the linac disagreed by progressively more, and by as much as 20 degrees at the last linac cavity. The author hypothesized that it could be an indirect result from varying temperatures in the linac tunnel, similar to suspicions raised in this thesis about temperature effects in ISAC.

Also of interest, Morozov, Peters, and Shishlo state[36] that the SNS linac calibration changes on an approximately weekly time scale due to maintenance of hardware. This seems rather concerning - on one hand if hardware needs to be replaced this frequently there must be some deficiency causing such frequent failures. However on the other hand it could be that they have in fact observed changes in calibration due to temperature dependent effects as hypothesized earlier at SNS and suspected now at ISAC-II.

Last, as discussed in the conclusion of [36], they seem to currently take the approach of pre-defining any possible linac configurations. Meaning that for their test, the calculation of necessary cavity phases and optics, etc, was already done offline before the test. They note that to cover scenarios with 1 failed cavity they need 81 saved states, and $81 \cdot 80$ saved states for 2 failed cavity scenarios, etc. In this aspect, ISAC-II is already in a more flexible position. With the HLA framework to read and write directly from the control system, and `TRANSOPTR` able to quickly re-

optimize cavity phases and transverse focal strengths, each of the 40! configurations of the linac can be calculated in near real-time, within 1 minute. It is also worth noting that the ISAC-II model can handle not only the 40! configurations of cavity availability, but also any changes made to the operating amplitudes selected by RF physicists.

Appendix B

Itemized List of Observations and Suggestions

Written more for the author and future staff members at TRIUMF rather than the general reader, it seems prudent to collect and condense some observations and suggestions made during the course of this thesis work. Some of these have already been mentioned while some have not.

Observations

- The calculation of the uncertainty in energy per nucleon previously used in the legacy matlab code for the ISAC-II linac was not correct, or was intentionally a rough approximation (without documentation of this decision). The width of the measured distribution was taken to be the uncertainty (using the standard deviation instead of the standard error of the mean).
- Another issue identified with the calculation of energy uncertainty was a missing factor of two in the relationship between relative energy error and relative time-of-flight error in Equation 2.31.
- There was an error in the legacy matlab code used to calculate the average energy per nucleon - it used the energy uncertainty of the meausurement between the second and third monitors as the uncertainty for the measurement between the first and third monitors.
- The legacy matlab code fit the ΔTOF vs ϕ data to a cosine curve which is not strictly correct - while approximately sinisoidal near full acceleration and in the

limit where the cavity provides a small change in velocity relative to the initial beam velocity this approximation is ok, but near the start of the linac and for scans spanning the full RF period it is better to use the correct expression.

- The energy per nucleon from the ISAC-I accelerator chain seems to have instabilities on the order of 0.03% over 30 minutes.
- It is suspected that the RF reference signal between the ISAC-I and ISAC-II accelerators varies with temperature in the ISAC-I experimental hall.
- Up until this (and other concurrent[24]) work, the effects of hysteresis and unipolar supplies on B-I curves have not been considered in the linac optics. This effect is most significant when quadrupoles are operating near the lower end of their operational range.
- The relationship between the control system setpoint and the actual effect varies between different cavities, and the effect is non-linear.
- The setpoint range over which the actual phase completes a full cycle ranges between around 374 and 384.
- The ISAC-II global phase appears to adjust the phase relative to the reference signal in the opposite direction of all other phase shifters.
- Power meters for measuring and locking on to ampltiudes of ISAC-II cavities are not linear over the full operating regime - typically for overperforming cavities an assumption that the physical field in the cavity will scale linearly with control system setpoint is not correct.
- SEBT:RPM8 is physically closer (approximately 15 cm) to the downstream SEBT:Q9 than indicated on design drawings.
- There is a possible source of emittance growth as the beam exits the ISAC-II linac.
- There is an unidentified source of error in the FTM measurement that seems to relate to the positions of the FTM monitors.
- It has been over 10 years since the timing of the FTM monitors was checked and calibrated, during which time they have been serviced.

Suggestions

- Upgrade the low-level RF (LLRF) boards to digitial phase shifters, which are expected to behave much more linearly than existing analog phase shifters.
- Adjust the attenuation on the return signal from cavity pickups to keep the signal within the linear range of the power meters.
- Add or improve time-domain diagnostics immediately upstream of the DSB buncher and SCB1:CAV1, as well as between each cryomodule.
- Carry out multiparticle simulations of the linac and comparison with observation.
- Further investigate errors in FTM measurements, including situations with non-negligible transverse-longitudinal coupling and with positioning of the FTMs themselves.
- Carry out simulations of steering effects from quarter wave cavities and make use of automated orbit correction methods.
- Attempt to fully automate the collection of calibration data used to setup the lianc model, similar to as done at FRIB.
- Use existing apertures or install new collimators in a position such that they can be used to cut the beam emittance and create a 'pencil beam'. This would be useful both for investigation into the transverse optics as well as longitudinal.
- Use the new model along with multiparticle simulations to investigate linac solutions with varying relative RF phases and amplitudes.
- Quadrupoles should be set in a reliably reproducible manner, such that effects of hysteresis on the uncertainty of the final quadrupole fields are minimized.
- Upgrade controls of the LPMs within the linac so that they data from these devices can be collected in an automated way.
- FTM monitor calibrations should be checked using the upstream laser system and re-checked on an annual basis.

- Measure and investigate any drifts in RF phase between ISAC-I and ISAC-II RF systems. Also measure between each different RF console. Maintaining a high level of beam quality and automating various processes requires a high level of stability in this system.
- Measure distances between the various cryomodules using the beam as a diagnostic, compare to design tolerances and do routine checks before/after cryomodules are serviced during extended maintenance periods.