

DESIGN AND COMMISSIONING OF A 16.1 MHZ MULTIHARMONIC BUNCHER
FOR THE REACCELERATOR AT NSCL

By

Daniel Maloney Alt

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Physics - Doctor of Philosophy

2016

ABSTRACT

DESIGN AND COMMISSIONING OF A 16.1 MHZ MULTIHARMONIC BUNCHER FOR THE REACCELERATOR AT NSCL

By

Daniel Maloney Alt

The ReAccelerator (ReA) linear accelerator facility at the National Superconducting Cyclotron Laboratory is a unique resource for the nuclear physics community. The particle fragmentation beam production technique, combined with the ability to stop and then reaccelerate the beam to energies of astrophysical interest, give experimenters an unprecedented range of rare isotopes at energies of nuclear and astrophysical interest. The ReAccelerator also functions as a testbed for technology to be incorporated in the upcoming Facility for Rare Isotope Beams linear accelerator, which will eventually in turn become the beam source for ReA.

This prototype nature of the ReAccelerator, however, dictated some design choices which have resulted in a final beam with a time structure that is less than ideal for certain classes of experiments. The cavities and RFQ used in ReA have an operating frequency of 80.5 MHz, which corresponds to a separation between particle bunches at the detectors of 12.4 ns. While this separation is acceptable for many experiments, sensitive time of flight measurements require a greater separation between pulses. As nuclear physics experiments rely on statistics, a solution to increasing bunch separation without simply discarding a large fraction of the beam particles was desired.

This document describes the design and construction of such a device, a 16.1 MHz multiharmonic buncher. The first chapter provides background information on the NSCL and ReA, and some basic concepts in accelerator physics to lay the groundwork for the project.

Next, more specifics are provided on the time structure of accelerated beams, and the experimental motivation for greater separation. The third chapter outlines the basic principles of multiharmonic bunching.

In order to evaluate the feasibility of any buncher design, the exact acceptance of the Radiofrequency Quadrupole (RFQ) of the ReAccelerator needed to be empirically measured. Chapter 4 describes the results of that measurement. Chapter 5 outlines the simulations and calculations that went into the design choices for this particular buncher, incorporating the results of the RFQ measurements. The next two chapters describe the construction, installation, and testing of the device, and give experimental results. Finally, Chapter 8 summarizes the project and the final steps which need to be undertaken to make the device a simple to use asset for future experimentalists at ReA.

Copyright by
DANIEL MALONEY ALT
2016

Dedicated to Leigh Van Handel for lorem ipsum and various other phenomena.

ACKNOWLEDGMENTS

The design and construction of this device, along with the writing of this document, represents an enormous amount of effort by a large number of people. This is a partial list only, and there are undoubtedly people who I have forgotten that deserve a place here.

First and foremost, my adviser Mike Syphers deserves my heartfelt and eternal thanks. Accepting a non-traditional student with my unusual background was a risk, and I appreciate his being willing to take a chance on me as a student. Throughout the entire process of design, approval, construction, and testing of the buncher, he was a source of excellent advice, an unflagging advocate, and a good friend.

Second, John Brandon and his colleagues from the NSCL / FRIB RF group, Dan Morris, Nathan Usher, Shen Zhao and Mike Holcomb. Together, they took my abstract optics, simulations, and plans, and transformed them into a device that actually works. Especially during testing, none of them flinched from late night stays to try and push the device the last inches over the finish line. I am deeply in their debt.

The administrators of the ReAccelerator project have always been solid supporters of this device and its sometimes bizarre demands. (You want to drill a hole in our clean room? Well, OK!) First Daniella Leitner and Walter Wittmer, and later Antonio Villari, were all extremely supportive, and offered valuable suggestions and experience which contributed to the success of the project. I am particularly grateful for Antonio's willingness to roll up his sleeves and assist with debugging the diagnostics for the final testing.

Sam Nash. It is impossible to overstate the importance of Sam Nash to all aspects of the ReAccelerator project. His knowledge of the entire facility is astonishing, and it is hard to imagine how it could continue to function half as well as it does without his constant

attention and dedication. His assistance and support in this new aspect of the facility has been utterly consistent with his support to all other aspects of the facility, which is to say, simply incredible.

The ReA3 operators, and in particular Shannon Krause and Randy Rencsok, have shown great patience in their support of this device. The critical measurements of the RFQ acceptance in particular would not have been possible without their skilled efforts.

Eugene Tanke's "Dynac" code was absolutely indispensable in the simulations for this project. I am especially grateful for his willingness to continually add functionality and documentation to the code to support my requests. My only regret in this area is that I did not have the time or resources to port DynacGUI to an open source format.

Other NSCL / FRIB personnel who were instrumental in this endeavor: Alain LaPierre, for not only providing me with EBIT beam for my experiments, but for his willingness to entertain the possibility of doing strange things to his time structures in the future, and his willingness to suffer some late nights during the final testing. Jeff Wendstrom, for the excellent mechanical design and adjustability of the buncher electrodes, the NSCL machine shop for the construction, and Dave Sanderson for the alignment of same.

Peter Ostroumov and his group at ANL provided valuable suggestions and the use of the TRACK code for aspects of the buncher design. Most importantly, the existence of their 12.125 MHz buncher was continued proof positive that such a device could be built.

The members of my committee: Norman Birge, Felix Marti, Wolfgang Mittig, and Yoshishige Yamazaki, for their support, suggestions, and willingness to submit to endless bouts of rescheduling on the final defense date.

Steve Lund, for stepping in as my line manager for the last year of the project, and his willingness to navigate countless layers of bureaucracy on my behalf, a fate I would wish on

no one.

My predecessors in accelerator physics at NSCL, Carla Benatti and Jeremiah Holzbauer. Without their suggestions on how to work more efficiently and survive graduate school, I would likely still be there. In a similar vein, my officemates Mike Jones, Mike Bennet, Vincent Bader, and Amanda Prinke all contributed in various ways to helping me maintain my sanity.

Also invaluable on the sanity front were my fellow officials, as well as the skaters and coaches of the Lansing Derby Vixens. Without the ability to spend time on roller skates and blow whistles at people, I honestly don't know how other people make it through graduate school.

Finally, and most importantly, my wife, Leigh Van Handel. Going back to school for a physics doctorate at age thirty-seven from a career as a middle school music teacher is objectively insane. Not only did she support my choice, but with her steadfast love and encouragement, she singlehandedly made it possible for me to believe in my own ability to do this crazy thing. I can never thank her enough.

TABLE OF CONTENTS

LIST OF TABLES	xii
LIST OF FIGURES	xiv
KEY TO ABBREVIATIONS	xx
Chapter 1 Introduction	1
1.1 Overview of NSCL and ReA	2
1.1.1 Ion Sources	3
1.1.2 K500 and K1200 Cyclotrons	4
1.1.3 Target Area and A1900 Fragment Separator	5
1.1.4 Beam Stopping	7
1.1.5 Electron Beam Ion Trap and Q/A separator	8
1.1.6 Reaccelerator ReA	11
1.1.7 Future Plans	12
1.1.7.1 ReA Expansion	12
1.1.7.2 FRIB	13
1.1.8 ReA Time Structure Issues	13
1.2 Overview of Essential Accelerator Physics	14
1.2.1 Coordinates	16
1.2.2 Transverse Dynamics	19
1.2.3 Transverse Emittance and CS parameters	23
1.2.4 Longitudinal Dynamics	27
1.2.5 Longitudinal Emittance	30
Chapter 2 Time Structure of Accelerated Beams	31
2.1 Original Time Structure of ReA	31
2.2 Motivation for Different Time Structures	34
2.3 Methods for Achieving Greater Bunch Separation	36
2.3.1 Modification of EBIT Time Structure	36
2.3.2 Subharmonic Bunching	37
2.3.3 Other Solutions	37
Chapter 3 Principles of Multiharmonic Bunching	39
3.1 Ideal Bunching	40
3.2 Simulation of a Sawtooth Wave Using Sine Waves	42
3.2.1 Fourier Synthesis	42
3.2.2 Window Functions	44
3.2.3 Consequences of a Finite Number of Modes	47
3.3 Other Deviations from the Ideal in a Real Buncher	47

3.3.1	Transverse Field Flatness	48
3.3.2	Transit Time Factor	49
3.4	Satellite Bunches	50
3.5	Case Studies	55
3.5.1	80.5 MHz Multiharmonic Buncher at ReA	55
3.5.2	12.4 MHz Multiharmonic Buncher at ATLAS (ANL)	56
Chapter 4	Measurement of the Longitudinal Acceptance of the ReA RFQ	58
4.1	Measurement Methodology	59
4.1.1	Total Transmission As A Function of Energy	60
4.1.2	Transmission as a Function of Phase and Energy	62
4.2	Analysis and Results	63
4.2.1	Energy Measurement Only	63
4.2.2	Combined Energy and Phase Measurement	66
4.2.3	Error Estimates	73
4.2.4	Implications for the Buncher	74
Chapter 5	Simulation and Selection of Design Parameters	76
5.1	ReA Beamline Simulation	77
5.2	Buncher Frequency	79
5.3	Voltage and Focal Length	82
5.3.1	Ideal Case	82
5.3.2	Numerical Simulation	84
5.4	Beamline Reconfiguration	88
5.5	Electrode Geometry	91
5.6	Effect on the Beam at the Detectors	95
5.7	Satellite Bunches	96
5.7.1	Simulation of Satellite Bunches with No Cleaning	96
5.7.2	Simulation of Satellite Bunches with Cleaning	97
5.8	Number of Buncher Modes	99
Chapter 6	Fabrication and Installation	100
6.1	Mechanical Design	100
6.2	RF Design	103
6.2.1	Resonators and Tuners	103
6.2.2	Low Level RF	105
6.3	Electrode Installation	106
Chapter 7	Commissioning: Methodology and Results	109
7.1	Timing Wire Detectors	109
7.2	Timing Wire Testing	112
7.3	Open Loop Bunching	117
7.3.1	Test Setup	117
7.3.2	Results	118
7.4	Closed Loop Bunching	123

7.4.1	Setup	123
7.4.2	Single Mode Results	125
7.4.3	Multiple Mode Results	129
7.5	Comparison with Simulation	133
Chapter 8	Conclusions	136
8.1	Summary	136
8.2	Future Plans	138
8.3	Satellite Bunches	140
8.3.1	Integer Ratio Cavities	141
8.3.2	Beam Chopper	142
8.3.3	EBIT Switching	144
APPENDICES	146
Appendix A	Comparison of Accelerator Codes	147
Appendix B	DynacGUI	179
BIBLIOGRAPHY	189

LIST OF TABLES

Table 1.1:	NSCL Primary Beam List as of December, 2014	6
Table 1.2:	Beam fractions for several definitions of emittance as a function of RMS beam size σ [1].	27
Table 4.1:	Summary of beam transmission through the ReA RFQ as a function of deviation from the reference energy.	63
Table 4.2:	Width of phase acceptance at various energies for the ReA RFQ, along with the measured phase offsets.	68
Table 4.3:	Calculated Δ_{tof} and Δ_{ϕ} values from the MHB to the RFQ relative to the reference particle.	71
Table 5.1:	Window functions with the highest bunching efficiency by frequency and initial beam energy spread.	87
Table 5.2:	Predicted minimum time and energy spread at target for beam energies from 0.3-3 MeV and 0.25-0.5 Q/A.	95
Table 6.1:	Lengths of the resonators and tuners for the 16.1 MHz buncher. . .	105
Table 7.1:	Open loop testing results for the first (16.1 MHz) mode of the buncher.	120
Table 7.2:	Open loop testing results for the second (32.2 MHz) mode of the buncher.	121
Table 7.3:	Open loop testing results for the third (48.3 MHz) mode of the buncher.	122
Table 7.4:	Closed loop testing results for the first (16.1 MHz) mode of the buncher.	126
Table 7.5:	Closed loop testing results for the second (32.2 MHz) mode of the buncher. No reflected power was observed at any of the chosen voltage settings.	126
Table 7.6:	Cavity settings for testing of the buncher with accelerated beam. All of the cavities in the third cryomodule were set off.	131

Table 7.7:	Summary of bunched beam properties. “Transmission” gives the percentage of the beam current measured after the linac relative to the beam current measured just before the buncher. “Bunching Efficiency” gives the percentage of the total beam detected at L110 in the main pulse.	133
Table A.1:	Comparison of Units Used in Several Accelerator Codes	178

LIST OF FIGURES

Figure 1.1:	Layout of the NSCL. Image Credit: Erin O'Donnell.	3
Figure 1.2:	Cutaway view of the Electron Beam Ion Trap.	9
Figure 1.3:	Layout of the ReAccelerator from the EBIT to the detector hall. . .	11
Figure 1.4:	The coordinate frame used to describe particles relative to a moving reference particle.	15
Figure 1.5:	The paraxial approximation means that x' can be approximated as θ for small values of θ	16
Figure 1.6:	Paths of three particles with identical energies, but different initial displacements through a dipole.	22
Figure 1.7:	Paths of three particles with differing energies but identical starting trajectories through a dipole.	22
Figure 1.8:	The phase space ellipse described by the Courant-Snyder parameters.	24
Figure 2.1:	Simulated longitudinal structure of the ReA3 beam at various points. From left to right: the exit of the EBIT, immediately after the 80MHz MHB, the start of the RFQ, the exit of the RFQ, the exit of the linac, and the final target.	33
Figure 2.2:	The uncertainty in the time of flight t for a single zero-width pulse is determined only by the pulse spread δt	34
Figure 2.3:	When multiple pulses are launched τ seconds apart, the measurement of the time of flight becomes uncertain by an integer multiple of τ . If δt is larger than τ , the measured time of flight can take any value and becomes completely uncertain.	35
Figure 3.1:	Comparison of the effect of an ideal buncher with a realistic bunching waveform. The upper plots show the beam bunched by a perfect sawtooth wave with zero reset time. The lower plots show a bunching waveform made up of four sine waves.	41
Figure 3.2:	A sawtooth wave synthesized using the first 2, 5, and 100 components of the Fourier series.	43

Figure 3.3:	Window function amplitude vs. normalized component for several window functions. Normalized component = (component number - 1) / (number of components - 1) Plotted here for $n = 10$ components.	45
Figure 3.4:	Side view of a pair of buncher electrodes with a conical profile. The arrow indicates the direction of beam travel	48
Figure 3.5:	A plot of the time and energy distribution of a simulated beam bunched at 16 MHz. The dashed lines indicate the 2 ns acceptance windows for a device operating at 80.5 MHz.	51
Figure 3.6:	The effect of shortening the initial EBIT pulse. At left, the shortened pulse after interaction with the bunch. At right, the beam at the buncher focus. No tails are present because the shorter initial pulse falls entirely within the linear portion of the bunching waveform. . .	52
Figure 3.7:	Bunch cleaning by an off frequency cavity. The red line is at a 4:5 frequency ratio with the blue one. Black dots represent beam particles at a the zero phase of both frequencies.	53
Figure 4.1:	A simple illustration of a longitudinal acceptance window.	59
Figure 4.2:	With the MHB turned off, the transmission through the RFQ was measured at a number of energies to give a picture of the relative transmission at each level.	61
Figure 4.3:	With the MHB turned on, the transmission through the RFQ was measured at a number energies and through a full 360 degrees of phase, giving a series of transmission curves.	62
Figure 4.4:	RFQ beam transmission vs. input energy for ReA RFQ.	64
Figure 4.5:	Naive RFQ acceptance window with no phase measurement or time-of-flight correction.	65
Figure 4.6:	Three phase scans at different energies.	65
Figure 4.7:	The simulated acceptance of the ReA RFQ. The white area at the center represents particles which are accelerated by the RFQ, and the red area shows particles which are not transported.	66
Figure 4.8:	All phase measurments superimposed. Each curve is normalized to a max of one, and phase units are unadjusted.	68

Figure 4.9:	The phase curves in figure 4.8 with the maximum transmission for each energy aligned at zero phase.	69
Figure 4.10:	Phase curves with zeroes aligned and scaled by overall transmission at each energy. The second figure is an interpolated surface generated from the data points in the first figure.	69
Figure 4.11:	A contour plot of the measured acceptance surface from Figure 4.10, with the 50% transmission contour emphasized, overlaid on the simulated acceptance.	70
Figure 4.12:	Measured and predicted phase advances relative to the reference energy for the phase measurements.	72
Figure 4.13:	For each energy, the offset from the predicted relative phase advance.	73
Figure 4.14:	The contour plot of the measured acceptance surface with the time-of-flight correction applied, overlaid on the simulated acceptance.	74
Figure 5.1:	A sample longitudinal model of ReA from the EBIT to the detector hall, showing x and y envelopes, particle count, reference particle energy, and the locations of beamline elements.	78
Figure 5.2:	Schematic illustration of bunching at several different subharmonics. Each plot shows energy vs. time. The first column shows the initial beam, the second column the beam immediately after bunching, and the third column shows the beam at the focus.	80
Figure 5.3:	Simulated longitudinal beam phase space for two different values of initial energy spread.	85
Figure 5.4:	Peak bunching voltage vs. focal length for a Monte Carlo model of a 16.1 MHz buncher.	86
Figure 5.5:	Reconfiguration of the ReA LEBT section. The double width diagnostic box (L052) after the stable ion source has been moved downstream of the first quadrupole doublet, and divided into a single diagnostic box and the 16.1 MHz buncher.	89
Figure 5.6:	A before and after comparison of the simulated beam envelope for the reconfiguration of the ReA LEBT line. The top figure shows the simulated envelope for the original configuration, and the lower figure the fitted envelope after the quadrupole move. See text for further explanations.	90

Figure 5.7:	Electrode design for the 16.1 MHz MHB	91
Figure 5.8:	Cross section of the 16.1 MHz MHB electrodes modeled in CST STUDIO, showing the accelerating field in the $x = 0$ plane.	92
Figure 5.9:	One simulation of the LEBT portion of the ReA beamline using a 3D field model for the 16 MHz buncher.	93
Figure 5.10:	Peak voltage for each mode for a range of electrode geometries. . . .	94
Figure 5.11:	Simulation of the central bunch and satellite bunches at the detector position for the middle ReA beamline. This simulation was of a $2H^+$ beam at 3MeV/u.	97
Figure 5.12:	Normalized histogram showing particles in satellite bunches relative to the main bunch when cleaned with two 96 MHz cavities located at the position of the present CM1. Time axis is in ns.	98
Figure 6.1:	The completed electrode assembly, showing the adjustment screws for the x and y axes and the alignment monuments. The center plate can also rotate to ensure the centerline of the electrodes is parallel to the beam axis.	101
Figure 6.2:	A schematic drawing of the connections for the RF system for the 16.1 MHz buncher.	105
Figure 6.3:	The control screen for the LLRF module, showing the controls for each of the three modes.	106
Figure 6.4:	The electrode assembly installed in the beamline, with the resonators visible	108
Figure 7.1:	Two views of the timing wire detector. The MCP is indicated by number 16, and the wire itself by 18. In the right hand view, the direction of beam travel is into the page, striking the wire before exiting on the far side of the “can.”	110
Figure 7.2:	The data acquisition path for the timing wire measurements.	111
Figure 7.3:	A example of a histogram from the TAC. This example is for beam bunched with the second (161.0 MHz) mode of the 80.5 MHz MHB at a control system setting of 11.	112

Figure 7.4:	The mean width of the bunched beam at a range of control system settings using the 80.5 MHz buncher. The 80.5 MHz mode is shown at left and the 161.0 MHz mode at right.	115
Figure 7.5:	Comparison of the FWHM curve for simulated bunched beams with a range of initial energy spreads for the first mode of the 80.5 MHz multiharmonic buncher.	116
Figure 7.6:	Measured beam FWHM from Figure 7.4 overlaid on simulated bunch widths. The X-axis of the measured data has been scaled to match the simulated beam voltage.	116
Figure 7.7:	The drive power setup for open loop testing of the 16.1 MHz buncher.	117
Figure 7.8:	Measured FWHM for the first (16.1 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match.	120
Figure 7.9:	Measured FWHM for the second (32.2 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match. . . .	121
Figure 7.10:	Measured FWHM for the third (48.3 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match. . . .	122
Figure 7.11:	Fitted peaks for manually tuned open loop bunching using modes 1 and 2.	123
Figure 7.12:	Measured FWHM for the first (16.1 MHz) mode of the buncher in closed loop mode, overlaid on a set of simulated curves, with the x -axis scaled to match.	127
Figure 7.13:	Measured FWHM for the second (32.2 MHz) mode of the buncher in closed loop mode, overlaid on a set of simulated curves, with the x -axis scaled to match.	127
Figure 7.14:	An illustration of an overfocused beam, showing the expected double peaked structure of the histogram.	128
Figure 7.15:	Initial beam time structure measurement taken at L110 with no phase adjustment.	132
Figure 7.16:	Best case bunching achieved at L110 with amplitudes 335 V and 35 V, and phases 22 and 202 degrees.	132

Figure 7.17:	Side-by-side comparison of the time structure of the beam at L110. The left hand image is the actual histogram measured on the timing wire, and the right hand image is a simulation of the same conditions.	135
Figure 8.1:	A simulated time distribution using all three bunching modes at L110. The calculated efficiency is 92.96%.	139
Figure A.1:	A sample emittance plot produced using DYNAC	151
Figure A.2:	An example of a longitudinal plot produced with COSY	156
Figure A.3:	An example of TRACK's output screen. (Phase vs. z plot not shown)	160
Figure A.4:	An example of a figure generated with MAD	164
Figure A.5:	The main screen of TRACE3D showing a profile graph and before and after emittance plots.	168
Figure A.6:	A sample plot generated by the GRAPHIC TRANSPORT FRAMEWORK interface to TRANSPORT.	172
Figure A.7:	Image Credit: xkcd.com. Used under Creative Commons Attribution - Noncommercial 2.5 License [2].	177
Figure B.1:	DYNACGUI Main Control Panel	183
Figure B.2:	DYNACGUI Main Z-Axis Plot	183
Figure B.3:	DYNACGUI Emittance Plot	185
Figure B.4:	DYNACGUI's fitting tool screen.	187

KEY TO ABBREVIATIONS

- ANL - Argonne National Laboratory (Argonne IL, USA)
- ARTEMIS - Advanced Room Temperature Ion Source
- ATLAS - Argonne Tandem Linear Accelerator System (ANL)
- AT-TPC - Active Time Target Projection Chamber
- CCF - Coupled Cyclotron Facility
- CERN - European Organization for Nuclear Research
- CS Parameters - Courant Snyder (Twiss) Parameters
- CSS - Control System Studio
- EBIT - Electron Beam Ion Trap
- ECR - Electron Cyclotron Resonance
- FRIB - Facility for Rare Isotope Beams (East Lansing MI, USA)
- FWHM - Full Width Half Maximum
- GUI - Graphical User Interface
- JENSA - Jet Experiment for Nuclear Structure and Stellar Astrophysics
- LEBT - Low Energy Beam Transport
- LEP - Large Electron Positron Collider (Geneva, Switzerland)
- LHC - Large Hadron Collider (Geneva, Switzerland)
- Linac - Linear Accelerator
- LLRF - Low Level Radio Frequency
- MCP - Micro-channel Plate
- MEBT - Medium Energy Beam Transport
- MCA - Multi-channel Analyzer
- MHB - Multiharmonic Buncher
- MSU - Michigan State University (East Lansing MI, USA)
- NSCL - National Superconducting Cyclotron Laboratory (East Lansing MI, USA)

- PSI - Paul Scherrer Institute (Villigen, Switzerland)
- Q/A - Charge to mass ratio
- ReA - Reaccelerator
- RF - Radio Frequency
- RMS - Root Mean Square
- ROCS - ReA Operations and Control Software
- SG - Signal Generator
- SuSI - Superconducting Source for Ions
- TAC - Time-to-Amplitude Converter
- TTF - Transit Time Factor

Chapter 1

Introduction

Accelerator or beam physics is concerned with producing beams of particles with necessary characteristics for basic research, medicine, industrial, and other applications. While the characteristics of each accelerator are different, the goal of the accelerator physicist is always to provide a beam which most closely matches the needs of the user, with the highest reliability and the lowest cost.

It is sometimes the case that over the course of the design, construction, or operation of an accelerator that there may be a demand for an additional capability not originally envisioned by the designers of the machine. In that instance, it is incumbent on the beam physicist to determine if this request is feasible, and if so how best to meet it.

This dissertation documents one such instance. As outlined below, the ReAccelerator (ReA) at Michigan State University's National Superconducting Cyclotron Laboratory (NSCL) was originally designed with a pulse repetition rate at the target of 80.5 MHz. Even while the facility was still under construction, a number of nuclear scientists expressed reservations about this rate, and a desire was articulated for beams with greater time separation. This document will outline the design and construction of a device to accommodate that request.

The remainder of this chapter will provide background information about the NSCL and ReA and an introduction to basic beam physics concepts. Chapter 2 will discuss general principles of beam time structure, and then explore the specifics of the time structures available

at ReA. In Chapter 3, the basic principles of multiharmonic bunching will be documented. Chapter 4 describes a preliminary measurement which was essential to the process of selecting design parameters for the constructed buncher, which process is documented in Chapter 5. Chapters 6 and 7 describe the construction and testing of the device, and Chapter 8 presents the conclusions which can be drawn from the project.

1.1 Overview of NSCL and ReA

The National Superconducting Cyclotron Laboratory at Michigan State University (MSU) is a research facility exploring basic questions in nuclear science using beams of radioactive isotopes. Stable beams from an ion source are accelerated to high energies in a pair of coupled cyclotrons and then impinged on a thin target. The resulting collision products include unstable isotopes of experimental interest. A beam of the desired isotopes is filtered from the collision products in a fragment separator and then sent to one of a number of target areas.

Since this projectile fragmentation method for producing radioactive beams results in radioactive beams with comparable energy to the original stable beam, for experimental applications which require lower energy beams the particles are first thermalized in a gas stopper or a reverse cyclotron stopper. If an energy greater than thermal but lower than the production energy is required, the beam is first stopped and then accelerated to the desired velocity in a secondary linear accelerator, the ReAccelerator.

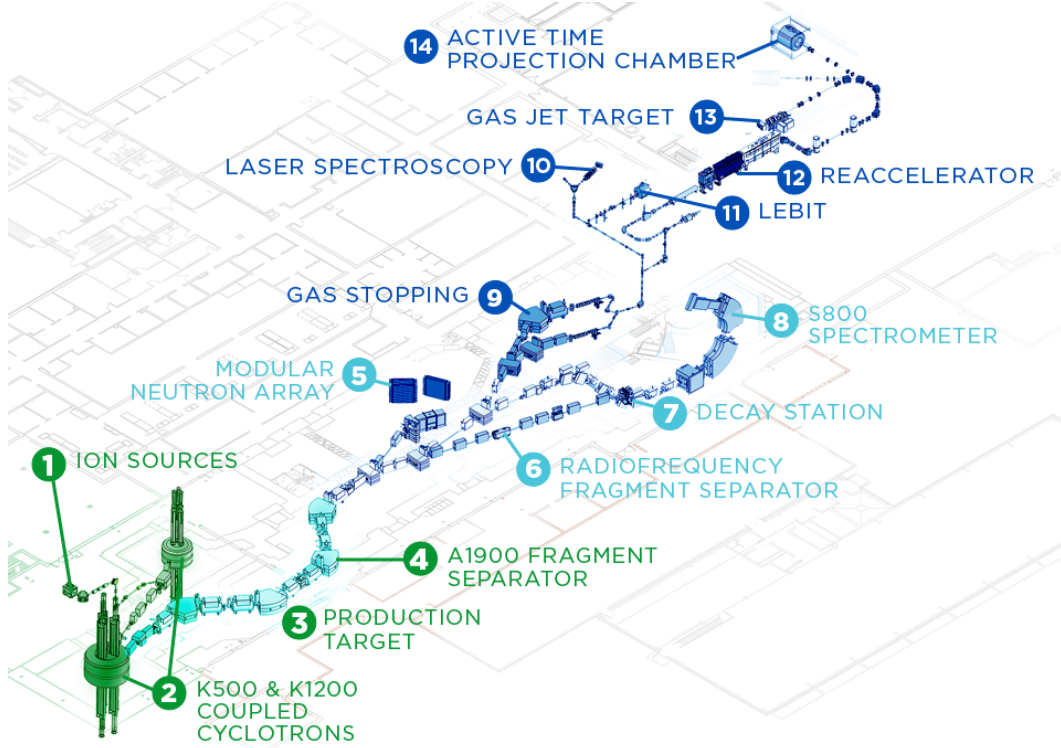


Figure 1.1: Layout of the NSCL. Image Credit: Erin O'Donnell.

1.1.1 Ion Sources

The first stage in producing the rare isotope beam at NSCL is the establishment of a beam of stable ions. NSCL has two primary ion sources, the Advanced Room Temperature Ion Source (ARTEMIS) [3], and the Superconducting Source for Ions (SuSI)[4][5]. Both of these are Electron Cyclotron Resonance (ECR) type ion sources.

In an ECR ion source, a low pressure gas of the desired ion is contained in a magnetic field. The gas is then excited by microwaves at a frequency corresponding to the cyclotron frequency of electrons in the field:

$$\omega_c = \frac{eB}{m} \quad (1.1)$$

where e is the charge on the electron, B is the magnitude of the field, and m is the mass of the electron.

This resonant excitation increases the energy of the free electrons in the gas, which collisionally ionize atoms of the source gas, leading to more free electrons in a positive feedback cycle. Radial confinement is provided by a hexapole magnetic field, and axial confinement by a solenoidal field. While SUSI is the more efficient source at NSCL, and is capable of producing much higher charge states, in the current coupled cyclotron operation mode, either source is capable of delivering most beams to the cyclotron. As such, the sources are generally used in alternation, with the source not currently providing beam to the cyclotrons used for future beam development.

1.1.2 K500 and K1200 Cyclotrons

The stable ion beam is accelerated by a pair of coupled superconducting cyclotrons, the K500 and the K1200. The K500 Cyclotron was completed in 1982, and was the first superconducting cyclotron in the world [6]. The K1200 Cyclotron was completed in 1988 and was at the time the most powerful cyclotron in the world. In 2001 the two cyclotrons began operation in a coupled mode, where beams from the ion source are first accelerated in the K500 and then further charge stripped before injection into the K1200 for final acceleration [7]. The pair of cyclotrons are now collectively referred to as the Coupled Cyclotron Facility (CCF)

The “K” number of a cyclotron is a measure of its total accelerating strength. It represents the kinetic energy in MeV a proton would achieve if accelerated at the maximum field of the cyclotron:

$$K = \frac{e^2}{2m_a} * (B\rho)^2 \approx \left(48 \frac{\text{MeV}}{(\text{T m})^2}\right) * (B\rho)^2. \quad (1.2)$$

Here e is the electron charge, m_a is an atomic mass unit, B is the maximum magnetic field

of the cyclotron and ρ is the maximum radius of the cyclotron.¹ Typical extraction energies from the K500 cyclotron in the coupled mode are on the order of 10-20 MeV/u (per nucleon) [8]. Final production beams from the K1200 cyclotron have energies on the order of 150 MeV/u. Table 1.1 shows the available primary beams from the coupled cyclotrons as of December, 2014 [9].

1.1.3 Target Area and A1900 Fragment Separator

Once the desired primary beam has been brought to its final energy, it is impinged on a production target. This target is usually, but not always, a thin foil of beryllium. While most ions in the primary beam will not be altered by the target, some will strike a target nucleus directly, leading to fragmentation. The beam after the production target is therefore a combination of the stable input beam, and a mix of stable and unstable fragmentation products [10].

Separating the desired reaction product from the other fragments and the unreacted beam is the task of the A1900 fragment separator. This separator consists of four superconducting dipole magnets and 24 superconducting quadrupole magnets with large transverse acceptances (8 milliradian) [11]. Between the first two dipoles, the beam is dispersed in horizontal space according to its rigidity, and the desired ($\beta\rho$) can be selected with the insertion of a slit. The desired mass is filtered from the remaining beam with a “wedge” energy degrader placed between the second and third magnet. Finally, a second slit is inserted between the third and fourth magnets to filter for only the desired isotope. The purity of the final beam depends on the distribution of nearby reaction products with similar charge

¹While in most accelerator contexts, the parenthesized quantity ($B\rho$) represents magnetic rigidity (momentum to charge ratio), and is a property of the beam, in this equation B and ρ are separate, and are properties of the accelerator.

Particle	Energy (MeV/nucleon)	Intensity (pnA)
^{16}O	150	175
^{18}O	120	150
^{20}Ne	170	80
^{22}Ne	120	80
^{22}Ne	150	100
^{24}Mg	170	60
^{36}Ar	150	75
^{40}Ar	140	75
^{40}Ca	140	50
^{48}Ca	90	15
^{48}Ca	140	80
^{58}Ni	160	20
^{64}Ni	140	7
^{76}Ge	130	25
^{82}Se	140	35
^{78}Kr	150	25
^{86}Kr	100	15
^{86}Kr	140	25
^{96}Zr	120	1.5
^{112}Sn	120	4
^{118}Sn	120	1.5
^{124}Sn	120	1.5
^{124}Xe	140	10
^{136}Xe	120	2
^{208}Pb	85	1.5
^{209}Bi	80	1
^{238}U	45	0.1
^{238}U	80	0.2

Table 1.1: NSCL Primary Beam List as of December, 2014

to mass ratios [12].

1.1.4 Beam Stopping

By the nature of the particle fragmentation process for rare isotope beam production, the purified beam of rare isotopes has nearly the energy of the primary beam when it was impinged on the target. At the NSCL, this can be up to 170 MeV/u (See Table 1.1). While this is acceptable for a number of experimental methods which rely on secondary collisions, there are also many experiments which require beams of radioactive isotopes at lower energies. The first step in producing any lower energy beam is thermalization in some type of stopping device, followed if necessary by reacceleration to the required energy. (Slowing the beam directly to an intermediate energy would be very difficult to accomplish without unacceptable emittance growth.) The NSCL has two primary beam stopping devices, a linear gas cell completed in 2004 [13], and a reverse cyclotron stopper with a projected installation date in 2016 [14].

The NSCL linear gas stopper consists of a 50 cm long gas chamber filled with a pure helium buffer gas. The incoming radioactive beam is passed through a set of degraders and wedges and enters the chamber via a thin beryllium window. Within the gas chamber, which is held at a pressure of 1 bar, collisional interaction with the gas slows the beam while simultaneously reducing its charge state back to 1+. The stopped ions are guided to the exit nozzle of the cylinder using a combination of electrostatic fields produced by a series of ring electrodes and gas pressure from the supersonic exit of gas through the nozzle.

The drawback of the gas stopper method is that as the incoming beam current increases, the helium buffer gas becomes ionized, and space charge effects begin to limit the efficiency of the device. As an alternative, the NSCL has developed a reverse cyclotron gas stopper

which is scheduled to be installed in 2016. As the name implies, the device essentially operates equivalently to the accelerating cyclotrons - a magnetic field confines the incoming ions to circular paths. Instead of being accelerated by an applied electric field, the beams are decelerated by helium buffer gas contained within the stopper. By contrast to the linear gas cell, since the ions in this device take a much longer path length to extraction, the gas pressure can be much lower, thus limiting space charge effects in the buffer gas. As the ions slow, they spiral to the center of the device, where they are extracted by a combination of static and RF electric fields.

Once extracted, the beams can be directed to the low energy experimental areas of the NSCL, or transferred to the reaccelerator for further acceleration. It is important to note that due to interaction with the buffer gas, the charge state of the extracted beam is decreased to $1+$, just as in the linear gas cell.

1.1.5 Electron Beam Ion Trap and Q/A separator

Once the thermalized beam is extracted from the stopper, it must be charge bred back up to the desired level of ionization before it can be reaccelerated. This is accomplished in the Electron Beam Ion Trap (EBIT) [15]. (Figure 1.2)

Radial confinement within the trap is provided by a combination of two superconducting magnet structures - a pair of Helmholtz coils and a longer solenoid magnet which can be adjusted independently. Axial confinement is provided by a series of 22 ring shaped electrodes.

The $1+$ charge state beam from the gas stopper is injected into the trap where it interacts with a coaxial beam of electrons produced from a gun at the far end of the trap. Interaction with the electron beam increases the ionization state of the incoming ions, increasing their




Figure 1.2: Cutaway view of the Electron Beam Ion Trap.

level of confinement within the trap, and preventing their escape. Due to this configuration, beam can be injected continuously into the trap from the stopping area. The particles within the trap continue to increase in ionization until the desired charge state is reached.

Once the ions are at the required charge to mass ratio, the electrode potential at the exit of the trap is lowered and the beam is directed towards the reaccelerator. The trap is biased at the voltage required to achieve a final beam energy of 12 keV/u. This energy can be achieved for ions with charge to mass (Q/A) ratios from 1:2 to 1:5. The time and energy structure of the beam injected into ReA is thus critically determined by the parameters of the EBIT: the thermal distribution of the beam in the trap, and the time structure of the lowering of the trapping potential. Typical energy spreads are on the order of 0.3% [16]. If the trap is simply opened and allowed to empty, the typical time structure is on the order of tens of microseconds. There has been exploration of using alternate modes of quickly

opening and closing the trap electrodes to create shorter pulse structures, with successful extraction down to pulses of $2\mu\text{s}$. However, the capacitance of the system, combined with the size of the trapping potential (on the order of hundreds of volts), places a practical limit on the speed with which the electrodes can be switched open and closed.

After the beam of radioactive ions has been extracted from the EBIT, it is passed through a Q/A analyzer to ensure contaminants from the stopping and charge breeding process are minimized. This section consists of a 90 degree electrostatic bend followed by a 90 degree magnetic bend. Between the two bends, the beam is brought to a transverse focus with a horizontal spread dominated by energy dispersion, and a slit is inserted in the beam to select the desired beam energy. A second slit is inserted after the magnetic bending section where the transverse beam spread is dominated by momentum dispersion. Taken together, these two bends and slits constitute a mass separator [17].

Two important figures of merit for such a mass separator are the mass resolution R and the achromaticity. The mass resolution is defined as

$$R = \frac{\Delta M}{M} \tag{1.3}$$

where M is the mass of the desired particle and ΔM is the minimum difference in M at which a contaminant beam can be successfully separated. The achromaticity is the amount by which desired particles can deviate from the reference energy and still be returned to a common focus at the end of the separator. This characteristic was of particular interest for the ReA separator given the relatively wide ($\pm 0.2\%$) predicted energy spread of the beam from the EBIT. The final design of the Q/A section was calculated to have an R of 100 and achromaticity within $\pm 1.5\%$ of the nominal beam energy [18].




Figure 1.3: Layout of the ReAccelerator from the EBIT to the detector hall.

1.1.6 Reaccelerator ReA

The layout of ReA, including the EBIT and Q/A separator is shown in Figure 1.3. The purified beam of radioactive ions emerging from the Q/A separator is transported through the Low Energy Beam Transport (LEBT) line prior to injection into the accelerator itself. In addition to electrostatic quadrupoles used for focusing, the LEBT as designed contains an 80.5 MHz Multiharmonic Buncher (MHB) used to bunch the beam prior to injection and a final focusing room temperature solenoid. This area is also where the offline stable ion source used for commissioning and testing connects to the beamline, and is where the 16.1 MHz buncher which is the topic of this document was inserted. The MHB will be discussed in greater detail in the following chapters.

The first accelerating stage of ReA is a four vane, room temperature Radio Frequency Quadrupole (RFQ) structure. The RFQ is designed to accelerate the beam from 12 keV/u up to 600 keV/u for all Q/A values within the design specification of ReA [19]. Following

the RFQ are a series of three cryomodules, containing a total of 15 superconducting niobium quarter wave cavities. The first seven of these cavities are optimized for $\beta=0.041$, and the last eight for $\beta=0.085$, where β is the particle velocity relative to the speed of light. The RF frequency for the MHB, the RFQ, and the cavities is 80.5 MHz in each case, meaning that groups of particles (“bunches”) emerge from the accelerator at that frequency. The maximum final design energy of the beam following the third cryomodule is 3 MeV/u for uranium ($A=238$).

After the third cryomodule, the beam is transported via a final Medium Energy Beam Transport (MEBT) line to the experimental hall. This line is approximately 30 meters in length. As originally installed, this line contains no further elements for longitudinal correction of the beam, but one potential upgrade to this line would be one or two rebunching cavities to compress the longitudinal structure of the beam prior to reaching the experiments.

As of 2016 there were three main experimental lines at the end of the MEBT. The first connects to the Active Target Time Projection Chamber (AT-TPC) experiment, the second to the Jet Experiment in Nuclear Structure and Astrophysics (JENSA) and the third is a utility beamline which can be connected to a variety of detectors as needed.

1.1.7 Future Plans

1.1.7.1 ReA Expansion

The initial construction of ReA was completed in 2014 with the installation and commissioning of the third cryomodule. This incarnation of the accelerator is commonly referred to as “ReA3”, after the maximum design energy for the heaviest isotopes. Plans are under development for additional cryomodules which will bring the final energy to 6 or 12 MeV/u,

and these future stages are referred to as “ReA6” and “ReA12”.

Other possible improvements to ReA include the addition of an ECR ion source for injection of stable beams, and the addition of rebunching cavities between the third cryomodule and the detector stations to provide greater control over the time and energy spread of the beam at the detectors.

1.1.7.2 FRIB

In 2008, the U.S. Department of Energy approved the construction of the Facility for Rare Isotope Beams (FRIB) as a successor to the cyclotrons at NSCL. FRIB will consist of a superconducting linac to produce high intensity beams of rare isotopes at intensities up to 400 kW [20]. The FRIB beamline will connect to the existing NSCL experiment halls, including ReA, thus allowing the existing detectors to be used with the new, more intense beams.

1.1.8 ReA Time Structure Issues

The ReAccelerator was specifically designed to make use of 80.5 MHz prototype cavities built as part of the planning and development process for FRIB. Since the facility was conceived to make use of existing hardware, rather than being designed from the ground up, not all potential user needs were necessarily anticipated.

The selection of a bunching frequency equal to the accelerator frequency was a logical choice. For reasons that will be outlined in the following chapters, this buncher requires relatively little power and physical space, and no provisions needed to be made for dealing with “satellite” bunches.

However, the cost of this bunching method was a relatively short repetition period of

12.4 ns between bunches. In particular, time-of-flight measurements involving reaccelerated beams are made extremely difficult with this short period [21]. Future plans for an Isochronous Large Aperture Spectrometer (ISLA) will also be adversely affected by this time separation [22]. Chapter 2 will expand on the time structure of the ReAccelerator as constructed, and the reasons for desiring the option of a longer time separation. Before this explanation is given, some basics of accelerator physics are first outlined here.

1.2 Overview of Essential Accelerator Physics

At the most fundamental, accelerator physics is simply the application of the Lorentz force law:

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}) \quad (1.4)$$

There are, of course, a number of other complex effects which can come into play, including space charge, wakefields, misalignments, and so forth. Many sophisticated and powerful tools have been developed for analysis of beam dynamics (see Appendix A). However, it is useful to bear in mind that this simple equation is still at the heart of any analysis of the behavior of an accelerator.

With that in mind, the remainder of this chapter will explore some of the fundamental principles of accelerator physics that are used in this document. As we are concerned here with a linear accelerator, little to no mention will be made of issues that are primarily of interest in circular machines, such as stability of closed orbits or resonances. Additionally, since ReA will almost entirely transport beams with low charge currents, even when operated in conjunction with FRIB, there will be no discussion of issues associated with high current beams such as space charge effects.




Figure 1.4: The coordinate frame used to describe particles relative to a moving reference particle.

In order to describe the motion of a particle in an accelerator, it is first necessary to define the coordinate system in which this motion will be analyzed. Particle coordinates in an accelerator are usually described relative to a reference trajectory. The reference trajectory is the path followed by a reference particle which enters the accelerator with no deviation from the desired energy, transverse position, or transverse angle. In addition, for accelerators with time-varying fields, the reference particle also defines the ideal time of entry into the accelerator relative to the phase of the fields.

With the reference trajectory thus defined, an orthogonal set of planar x , y , and z coordinates can be defined at each point on the beamline. The variable z points along the reference trajectory in the direction of travel, and x and y are coplanar in the plane normal to z with their origin on the reference trajectory (Figure 1.4). While x and y can be selected to meet the conditions of a given beamline, the usual practice is to choose x and y right handed with x in the horizontal plane and y vertical. It is important to note that the orientation of these coordinate vectors can change relative to 3D Cartesian space as the origin is moved along the beamline.




Figure 1.5: The paraxial approximation means that x' can be approximated as θ for small values of θ .

1.2.1 Coordinates

There are a number of possible definitions of the z coordinate for an individual particle. While the x and y coordinates are always measured in distance from the reference trajectory, z may be measured either relative to an absolute point on the beamline, such as the starting point, or it may be measured relative to the position of a reference particle along the beamline. Whether z is measured relative to a fixed point or relative to a moving reference, it may still be expressed either as an absolute distance, a time of flight (for a given particle velocity), or a phase (for a given reference particle velocity and system frequency).

In addition to x , y , and z , a full description of each particle must also include the instantaneous rate of change of each of these quantities. For the transverse coordinates x and y , we designate the rates of change as x' and y' . (Note that some sources refer to these quantities as a and b .) Here we make use of the “paraxial approximation”, which applies when the particle’s trajectory is nearly parallel to the reference trajectory (Figure 1.5). In this case, the small angle approximation may be used to replace $\tan(\theta)$ with θ , and express x' and y' as angles. Further, in the case that any magnetic field present is perpendicular to the beam direction, x' can also be expressed as the ratio of the transverse momentum to the

longitudinal momentum.

$$x' = \frac{dx}{dz} = \tan\theta \approx \theta \quad (1.5)$$

$$\approx \frac{p_x}{p_z} \quad (1.6)$$

The rate of change of z has even more possible variants than z itself, starting with the symbol to be used, which can be either z' or δ . Conceptually, it is important to remember that this quantity is simply an expression of the rate of change of the particle's position along the direction of travel - in other words, the velocity of the particle. Less simple is how this value is to be expressed. As before, this quantity can be described either in absolute terms, or relative to a reference particle. In absolute terms, the velocity may be given directly, or the kinetic energy, total energy, or momentum of the particle specified (for a given particle mass.) If the velocity is rather to be specified relative to a reference particle, that difference may be specified as an absolute deviation in velocity (Δv), energy (ΔE), kinetic energy (ΔW), or momentum (Δp) from that of the reference particle. Alternately, the relative deviation in velocity ($\frac{\Delta v}{v}$), energy ($\frac{\Delta E}{E}$), kinetic energy ($\frac{\Delta W}{W}$), or momentum ($\frac{\Delta p}{p}$) from the reference particle may be specified.

As a final point of confusion, some accelerator codes, notably COSY [23], normalize the longitudinal coordinate by a factor of $\frac{\gamma}{1+\gamma}$, which must be borne in mind when attempting an apples-to-apples comparison between these and other codes. (This normalization corresponds to a conversion from relative kinetic energy to relative momentum [24].)

This document will, unless specified otherwise, express z and z' relative to a reference particle. The longitudinal coordinate z will usually be specified in units of time (typically ns), and z' in either absolute (ΔW) or relative ($\frac{\Delta W}{W}$) kinetic energy deviation from the

reference.

Once a suitable set of coordinates has been decided upon, each particle can be specified as a six element vector in phase space:

$$\vec{s} = \left\{ \begin{array}{c} x \\ x' \\ y \\ y' \\ z \\ z' \end{array} \right\} \quad (1.7)$$

Assuming the equations of motion can be linearized, the effects of beamline elements acting on the beam can be expressed as a series of matrices multiplying this phase space vector. For a simple example, a particle moving through a drift of length L experiences the following changes in the phase space coordinates (Continuing to assume the paraxial approximation.):

$$\begin{aligned} x &= x_o + L * x' \\ x' &= x'_o \\ y &= y_o + L * y' \\ y' &= y'_o \\ z &= z_o + L * z' \\ z' &= z'_o \end{aligned} \quad (1.8)$$

This can be expressed compactly in matrix notation as:

$$\vec{s} = \begin{bmatrix} 1 & L & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & L & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & L \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \vec{s}_o \quad (1.9)$$

It is important to reiterate that the choice of z and z' matrix elements (and by extension, equation of motion coefficients) depends on the units chosen for these quantities. Care must always be taken that units are understood and consistent when dealing with longitudinal dynamics.

1.2.2 Transverse Dynamics

If beamline elements which affect the energy of the beam are not considered, the matrix approach outlined above allows for a relatively simple calculation of the effects of a given set of elements on a particle. The linear nature of the equations of motion means that the matrices for each element on a beamline can be combined to give a single transfer matrix which describes the effect of that line on the phase space coordinates of a given particle.

For many common beamline elements, the x and y equations of motion are not coupled - the initial x and x' coordinates do not affect the final y and y' components, and vice versa. In matrix terms, this means that the off-diagonal 2×2 blocks within the transfer matrix are

zero:

$$\mathbf{T} = \begin{bmatrix} * & * & 0 & 0 & * & * \\ * & * & 0 & 0 & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \end{bmatrix} \quad (1.10)$$

For elements which are not affected by particle energy, this means that the equations of motion can be reduced to two simple 2x2 matrices, one for each transverse direction x and y . The simplest such element is a quadrupole, which focuses the beam in one transverse direction while defocusing in the other. The matrix for a quadrupole approximated as a thin lens is

$$\mathbf{T} = \begin{bmatrix} 1 & 0 \\ \pm \frac{1}{f} & 1 \end{bmatrix} \quad (1.11)$$

where f is the focal length of the quadrupole, and the sign of the lower left element is defined to be negative in the focusing direction and positive in the defocusing direction.

The only other purely focusing element used in the ReA beamline is the solenoid. Since the solenoid rotates the beam about the beam axis, x and y are no longer independent after passing through a solenoid, and the off diagonal elements of the transfer matrix are nonzero.

The x/y transfer matrix for a solenoid is [25]

$$\mathbf{T} = \begin{bmatrix} C^2 & \frac{1}{k}SC & SC & \frac{1}{k}S^2 \\ -kSC & C^2 & -kS^2 & SC \\ -SC & -\frac{1}{k}S^2 & C^2 & \frac{1}{k}SC \\ kS^2 & -SC & -kSC & C^2 \end{bmatrix} \quad (1.12)$$

where

$$k = \frac{B_s}{2(B\rho)}, \quad (1.13)$$

$$C = \cos(kL), \text{ and} \quad (1.14)$$

$$S = \sin(kL). \quad (1.15)$$

Here, L is the length of the solenoid, B_s is the magnetic field of the solenoid, and $(B\rho)$ is the rigidity of the beam.

The other important type of element in the ReA beamline is the bending element, which bends the trajectory of the beam using electric or magnetic fields orthogonal to the beam's direction of motion. Since there is no rotation of the transverse coordinates, the off diagonal blocks are once again zero. However, due to the curved path of the reference particle through the dipole, the transverse motion of the particle in the plane of the bend becomes correlated with the longitudinal coordinates of the particle.

For a dipole which bends in the xz plane, a particle which enters a dipole at the reference energy but with a displacement from the reference trajectory in x will see either a shorter or greater path length through the magnet, and will therefore see a change in its longitudinal position relative to the reference particle (Figure 1.6). Similarly, a particle entering on the




Figure 1.6: Paths of three particles with identical energies, but different initial displacements through a dipole.




Figure 1.7: Paths of three particles with differing energies but identical starting trajectories through a dipole.

reference trajectory but away from the reference energy will find itself traveling away from the central trajectory for which the magnet is calibrated (Figure 1.7). This is used deliberately in mass spectrographs such as the A1900, which are calibrated to prevent the passage of off energy particles.

This correlation between the bending plane coordinates and the z coordinate can be seen in the nonzero elements in the fifth row and sixth column of the transfer matrix for a bending magnetic dipole with bend radius ρ and magnetic field in the y direction B_y [25]

$$\mathbf{T} = \begin{bmatrix} C_x & \frac{1}{k_x} S_x & 0 & 0 & 0 & h \frac{1-C_x}{k_x^2} \\ -k_x S_x & C_x & 0 & 0 & 0 & h \frac{S_x}{k_x} \\ 0 & 0 & C_y & \frac{1}{k_y} S_y & 0 & 0 \\ 0 & 0 & -k_y S_y & C_y & 0 & 0 \\ -h \frac{S_x}{k_x} & -h \frac{1-C_x}{k_x^2} & 0 & 0 & 1 & -\frac{1}{\rho^2} \frac{(k_x \Delta s \beta^2 - S_x)}{k_x^3} + \frac{\Delta s}{\gamma^2} \left(1 - \frac{1}{\rho^2 k_x^2}\right) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (1.16)$$

where

$$n = - \left[\frac{\rho}{B_y} \frac{\partial B_y}{\partial x} \right]_{x=0, y=0}, \quad (1.17)$$

$$k_x = \sqrt{(1-n)h^2}, \text{ and} \quad (1.18)$$

$$k_y = \sqrt{nh^2}. \quad (1.19)$$

Here α is the bend angle of the dipole, k_x and k_y relate to the strength of the dipole (k_y is zero for a pure vertical field with no gradient), Δs is the path length of the reference particle through the dipole, h is $\frac{1}{\rho}$ and $C_x = \cos(k_x \Delta s)$, etc. Additional matrices can also be applied to account for effects due to the curvatures and angles of the entrance faces for dipole elements.

1.2.3 Transverse Emittance and CS parameters

Since accelerators are rarely tasked with accelerating a single particle at a time, it is valuable to extend the dynamics developed in the previous section to a description of the motions of ensembles of particles. Of particular value in this analysis is Liouville's theorem, which states that for a canonically conjugate pair of coordinates, the area occupied by a system in the phase space defined by that pair is conserved. Classically, the canonical pair of variables in each dimension is position and momentum. As indicated in Section 1.2.1, in situations where the paraxial approximation holds and all magnetic field components are perpendicular to the beam axis, x' may be expressed as a ratio of p_x to p_z and may therefore be considered as conjugate to x . In the longitudinal direction, a classically conjugate form is $z = v * (\Delta t)$, where Δt is the time of flight and $z' = \delta = \frac{\delta p}{p}$ [26].




Figure 1.8: The phase space ellipse described by the Courant-Snyder parameters.

With these coordinate pairs, x / x' , y / y' , and z / z' , chosen to be canonically conjugate, it is customary to plot each pair and examine the phase space represented. The total area in each phase space occupied by the beam is referred to as the “emittance” of the beam for that axis. Assuming the motion of the particles in each dimension is uncorrelated, these emittances are ideally conserved by Liouville’s theorem. This is only strictly true for unaccelerated motion.

An important caveat to the conservation of emittance is that as the momentum of the beam particles increases, the beam becomes more “rigid” and oscillations reduce in amplitude, thus reducing the emittance. If the emittance is normalized by the following formula [1],

$$\epsilon_N \equiv \epsilon \left(\gamma \frac{v}{c} \right) \quad (1.20)$$

the resultant normalized emittance should ideally remain constant in all cases. The primary reasons why it still might not, such as space charge and beam-beam interaction, are not of great concern in ReA, as mentioned above.

The equations of motion for transverse oscillation in an accelerator are of the form:

$$x'' + K(s)x = 0 \quad (1.21)$$

where s represents the position of the reference particle along its trajectory, and $K(s)$ is a piecewise function that represents the action of the magnetic and electric fields along the accelerator. (For a full derivation, see Edwards and Syphers [1].) This differential equation is known as Hill's equation, and the solution can be written as:

$$x(s) = A\sqrt{\beta(s)} \cos[\phi(s) + \delta]. \quad (1.22)$$

Here, $\beta(s)$ is a function of the arrangement of fields in the accelerator, $\phi(s)$ represents the total phase advanced of the oscillating particle up to point (s) in the accelerator, and A and δ are constants of integration. This can be rearranged to the form:

$$\epsilon\pi = \gamma x^2 + 2\alpha x x' + \beta x'^2 \quad (1.23)$$

which is the equation of an ellipse with an area of $\epsilon\pi$ (Figure 1.8). (Note that β in this equation is not the same as in equation 1.22.) Collectively, the values α , β , and γ are known as the transverse Courant-Snyder (CS) parameters for the beam on the x axis, and ϵ is called the emittance. Conservation of momentum constrains individual particles to remain on a given phase space ellipse as they move along the beamline, however, the ellipse itself changes shape as different fields are applied. In a drift, the particles in the upper half with positive x' values will move to the left and those with negative x' to the right, causing the major axis of the ellipse to shear in a counterclockwise direction. In a thin lens quadrupole, the x

values will remain unchanged while x' values are altered, causing the ellipse to scale along the vertical axis, etc. Thus, the CS parameters of the beam evolve along the accelerator with the changing shape of the ellipse.

Intuitively, these parameters can help give a quick assessment of the beam properties at a point - β is related to the x size of the beam, γ is related to the spread in trajectories on the x axis, and α relates to the angle the major axis of the ellipse makes with the origin. (Since x is minimized when the axis is vertical, α is zero when the beam reaches a waist.) If ϵ is changed without altering the other parameters, the ellipse scales larger or smaller in size but maintains the same basic shape.

Of the parameters, α , β , and γ , only two are independent, with the third being determined by the relation

$$\beta\gamma - \alpha^2 = 1. \quad (1.24)$$

An important note about emittance: While it is universally defined as a phase space area, there are many different conventions for actually calculating a number for the value of the emittance of the beam. The π in equation 1.23 may either be incorporated into the units (as $10 \pi^* \text{mm}^* \text{mrad}$) or explicitly included in the number (as in $31.4 \text{ mm}^* \text{mrad}$). Another common definition is to take the root mean square (RMS) beam size σ at a location where the value of β is known, and use these quantities to define a fraction of the total beam, as in Table 1.2.

As a result, it is critically important to determine which definition of emittance is in use when using a specific value for the emittance, rather than simply using the general principle of emittance conservation.

One other caveat - treating the x and y emittance as separate is only fully valid if the

ϵ	F (%)
$\frac{\sigma^2}{\beta}$	15
$\frac{\pi\sigma^2}{\beta}$	39
$\frac{4\pi\sigma^2}{\beta}$	87
$\frac{6\pi\sigma^2}{\beta}$	95

Table 1.2: Beam fractions for several definitions of emittance as a function of RMS beam size σ [1].

beam never passes through elements which correlate the two coordinates, such as solenoids. After passing through a solenoid, it is still possible to define a 4D conserved phase space area in $x/x'/y/y'$ space, but this is a far more complex prospect, and no consensus exists on how best to define such an emittance or how to apply it.

1.2.4 Longitudinal Dynamics

In some ways, longitudinal dynamics (dynamics in the direction of beam travel) are treated similarly to transverse dynamics. Despite the wide range of possible choices for expressing z and z' , these are chosen to be canonically conjugate variables, and thus display phase space conservation except in accelerating fields, or in situations (such as dipoles) where z and δ can become correlated to motion in other dimensions. However, there are some important distinctions.

As will be discussed in more detail in Chapters 2 and 3, there are two primary types of particle accelerator, those which use static electric fields for acceleration, and those which use radio frequency (RF) accelerating fields. The dynamics for a simple static accelerating gradient are simple - the Lorentz force law applied over a distance gives the change in energy. However, even this simple case means that the beamline can no longer be modeled

as a simple series of matrix multiplications, as many matrix elements are dependent on beam energy. Once the particle has been accelerated by a static accelerating gap, the matrices for subsequent elements must be recalculated for the new energy.

The problem becomes even more difficult for RF acceleration. The necessary phase information can be extracted from the z coordinates of the particle distribution, but since the field is changing with time, even an infinitely narrow accelerating gap will change the energy of particles that arrive at different times by different amounts. Most commonly, beamlines with accelerating elements will be modeled by a combination of transfer matrices for portions of the beamline where energy is not changing, and direct calculation of the change in energy for each particle where it is.

The simplest form of the energy change for a particle traversing an accelerating gap is given by:

$$E = E_o + qV_o \sin(\phi) \quad (1.25)$$

where q is the charge on the particle, V_o is the maximum voltage across the gap and ϕ is the phase of the RF at the time of the particle's arrival, with $\phi = 90$ degrees giving the maximum acceleration and $\phi = 0$ degrees giving no acceleration. While this would seem to be relatively straightforward to represent as a matrix,

$$\begin{bmatrix} z \\ E \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ qV_o \sin(\phi) & 1 \end{bmatrix} \begin{bmatrix} z_o \\ E_o \end{bmatrix} \quad (1.26)$$

a problem arises because ϕ is not a fixed value, but changes with time. As such, the entire transfer matrix for an RF gap becomes time dependent. (This simple example is also an oversimplification, because while z and $\frac{dp}{p}$ are canonically conjugate, z and E are not.)

The standard practice with RF accelerating structures is to fix the desired value of ϕ for the reference particle and adjust the accelerator so that the reference particle passes each accelerating station at the same phase, known as the “synchronous phase.” This does make the transfer matrix above time-independent. However, the problem returns when we try to describe the motion of an arbitrary particle with a phase or energy different from the reference particle. In that case, the equations of motion become

$$(E - E_s)_f = (E - E_s)_i + qV_o(\sin\phi - \sin\phi_s) \quad (1.27)$$

$$\phi_f = \phi + C(E - E_s)_f \quad (1.28)$$

where E and ϕ are the energy and phase of the particle in question, $_i$ and $_f$ represent the initial and final states before and after the gap, and C is a factor which is related to the geometry of the accelerator and the energy of the reference particle [1]. With the time dependence on the phase reintroduced, this situation is once again not representable by a static matrix.

A further complication is introduced by the fact that even the transverse matrix elements for magnetic elements are dependent on particle velocities, due to the $\vec{v} \times \vec{B}$ dependence of the Lorentz force law. As such, a single matrix representation for an accelerating structure isn’t easily generated, since the matrices for later beamline elements depend on the properties of the *particles* input to prior elements.

The accelerator simulation codes referenced in the document which DO treat complex accelerating structures such as bunchers, do so by transporting distributions of particles using the matrix formalism through non-accelerating elements, and then evaluating equations 1.27 and 1.28 individually for each particle at each accelerating element.

1.2.5 Longitudinal Emittance

As mentioned above, z and z' are always selected as canonically conjugate variables, so the beam area in this phase space is generally conserved for unaccelerated beams (except when correlated to another dimension, such as in a dipole). However, in actual practice, a beam's longitudinal emittance is often not well described by an ellipse, and as such CS parameters are less often used in the z direction. For example, the continuous beam emitted by the ReA3 EBIT appears as a simple rectangle on the phase space plot with a z width of 360 degrees and a z' height corresponding to the energy spread in the beam.

While the CS approach to describing the non-elliptical shape of the z/z' phase space is not always useful, Liouville's theorem is not dependent on any particular geometry on the phase plot, so the area of the plot (and thus emittance) is still conserved. This can be extremely useful in making simple estimates of the effect of bunching or accelerating elements on the beam time and energy spreads. For example, if a molecular hydrogen beam ($A=2$) has a longitudinal emittance of 2.4 eV ns per period, for the case of the continuous beam with period 12 ns the energy spread is

$$\frac{2.4 \text{ eV ns}}{12 \text{ ns}} = 0.2 \text{ eV}. \quad (1.29)$$

If that beam were then compressed to a length of 1 ns, the beam energy would expand to a spread of

$$\frac{2.4 \text{ eV ns}}{1 \text{ ns}} = 2.4 \text{ eV}. \quad (1.30)$$

Chapter 2

Time Structure of Accelerated Beams

All accelerating structures can be divided into two broad classes: DC and AC. DC accelerators are the simplest conceptually - an electrostatic field is used to provide acceleration in accordance with the Lorentz force law. These types of accelerator suffer from the drawback that they are limited to accelerating voltages on the order of a few megavolts [1], beyond which electrostatic breakdown occurs, placing a limit on the final energy of such a machine. Above that limit, radio frequency (RF) electric fields must be used.

The advantage of using RF fields for acceleration is that by using resonant cavities significantly higher accelerating gradients can be produced. For example, the cavities planned for FRIB will have peak electric fields at the average accelerating voltage from 26.5 - 30.8 MV/m [27]. The disadvantage is that the accelerated beam must be bunched in order to avoid deceleration during the phase of the electric field when its direction is negative.

2.1 Original Time Structure of ReA

There are three sources of beam acceleration in ReA: First is the EBIT, which ejects the beam at 12 keV/u via a DC bias. Second, the RFQ, which accelerates the beam to 600 keV/u. Finally, the linear accelerator (linac) itself, which accelerates the beam to a final energy up to 3 MeV/u for A=238. The initial acceleration from the EBIT is electrostatic, while the RFQ and linac use oscillating electric fields for acceleration. The RF frequency of

the RFQ and linac is 80.5 MHz.

In order for the alternating fields of the linac to provide acceleration to more than half of the beam particles, they must first be bunched before acceleration [28]. While there are RFQ designs capable of providing bunching as well as focusing [29], it has been found that the longitudinal emittance of the beam after the RFQ can be greatly improved with the addition of an external buncher before the beam reaches the RFQ [30]. For this reason, the LEBT section of ReA includes a multiharmonic buncher designed to longitudinally bunch the beam prior to injection into the RFQ.

In principle, an external buncher can operate either at the RF frequency of the accelerator or at any integer divisor (subharmonic) of that frequency. Bunching directly at the the RF frequency is the simplest case, and this was implemented for the design of ReA. Using the 80.5 MHz MHB, the beam is bunched with a period of 12.4 ns, and therefore one bunch occupies each RF “bucket” of the accelerator.

Within each bunch, the longitudinal structure is determined by a number of factors. In an ideal perfectly monoenergetic beam, all the particles travel at the same speed. However, in reality, there is always a range of energies in the beam. This leads to the beam spreading out along the beam axis, as the faster particles move ahead and the slower ones fall behind.

The specific case of ReA is illustrated in figure 2.1. When the beam leaves the EBIT, there is an initial energy spread ($\frac{\Delta W}{W}$) on the order of 0.3% [16]. The time structure of the initial beam is determined by the rate at which the ions leave the trap. However, since the trap empties on a characteristic time scale on the order of milli- or micro- seconds depending on the scheme used, the beam can be treated as continuous in time from the perspective of a buncher with a period of tens of nanoseconds. When the beam leaves the buncher, the energy spread of each bunch has been increased, and the time spread begins to decrease,

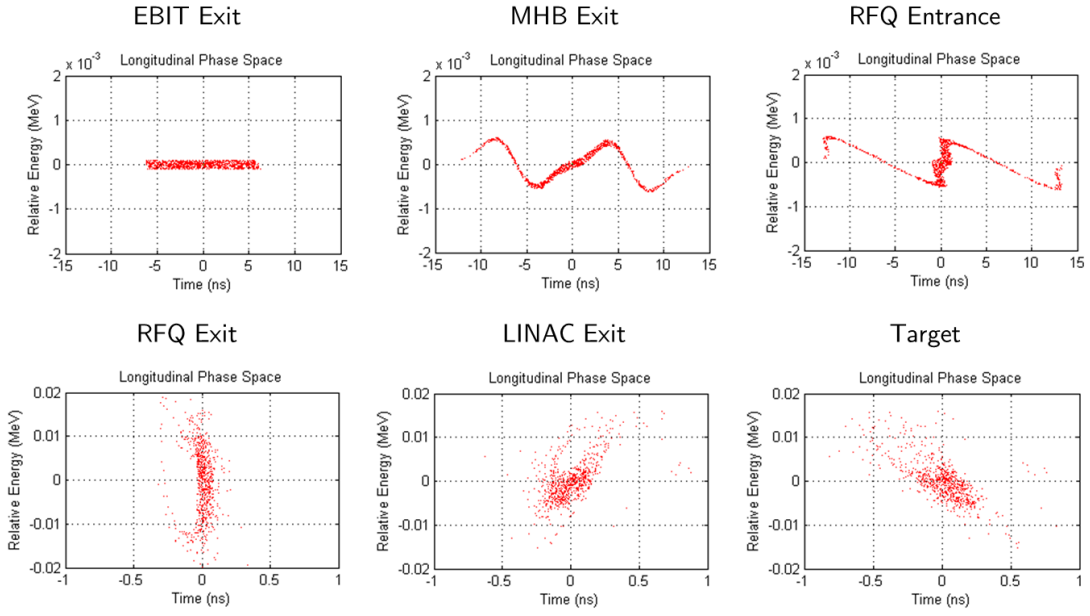


Figure 2.1: Simulated longitudinal structure of the ReA3 beam at various points. From left to right: the exit of the EBIT, immediately after the 80MHz MHB, the start of the RFQ, the exit of the RFQ, the exit of the linac, and the final target.




Figure 2.2: The uncertainty in the time of flight t for a single zero-width pulse is determined only by the pulse spread δt .

reaching a focus at the entrance of the RFQ.

As the total beam energy is increased in the RFQ and linac, the spread relative to that energy decreases. Finally this spreading can be reduced further by using the final cavities of the linac as rebunchers - the cavities are run at their non-accelerating phase in an attempt to equalize the particle energies within each bunch as much as possible. However, the long distance from the end of the linac to the experimental halls (30 meters) means that even with the lowest achievable energy spreads, some beam spreading in time is inevitable. The effect is more pronounced at lower energies, because the longer travel times mean that equivalent fractional energy differences have longer times to spread out than at higher energies. The worst case scenario occurs when the time spread of the bunch exceeds the bunching period (12.4 ns) and it becomes impossible to temporally resolve individual bunches at the detectors.

2.2 Motivation for Different Time Structures

Many experimental protocols using reaccelerated beams involve making a measurement of the time of flight of the particle from the source to the detector [21]. If a single bunch with no time spread is launched from the accelerator, then the primary source of uncertainty in the time of flight, t , of a particle within the bunch is the time spreading of that bunch (due to its initial velocity spread) between the launch and the detector (Figure 2.2). If δt is the full time width of the bunch at the detector, then the uncertainty in the measured time of




Figure 2.3: When multiple pulses are launched τ seconds apart, the measurement of the time of flight becomes uncertain by an integer multiple of τ . If δt is larger than τ , the measured time of flight can take any value and becomes completely uncertain.

flight is:

$$|t_{measured} - t| \leq \frac{\delta t}{2} \quad (2.1)$$

However, if multiple bunches are launched with a separation of τ , then the final time of flight is also uncertain by an integer multiple of τ , representing the inability to determine to which initial bunch the detected particle belongs. (Figure 2.3)

$$|t_{measured} - t| \leq n * \tau + \frac{\delta t}{2} \quad (2.2)$$

To minimize the difference between $t_{measured}$ and t , the period between successive bunches, τ , must be increased to the point that the value of n can be determined based on considerations extrinsic to the immediate measurement. Once n is known, the time of flight calculation reduces to the first case. At the same time, δt must be made as small as possible, and certainly smaller than τ . A value of δt larger than τ , as mentioned above, will make it impossible to resolve individual bunches at all.

Discussions by nuclear scientists at NSCL and elsewhere concluded that even with a low enough value of δt at the ReA detectors for temporal resolution, the original 12.4 ns value of τ for the bunch separation is not enough for an unambiguous determination of n - which launched bunch corresponds to a detected particle at the end of the beamline [31].

Finite detector reset times between bunches also contribute to the desire for a greater bunch separation.

2.3 Methods for Achieving Greater Bunch Separation

Several possible methods are considered here for lengthening the gap between bunches at the detectors of ReA. An ideal solution would have the same transmission efficiency as the 80.5 MHz buncher case, would have no particles in outlying bunches between the primary bunches, and would minimize the time and energy spread on target. Obviously, a preferred solution would also not consume undue resources, in terms of either money (either for equipment or operational costs), installation time, or physical space required.

2.3.1 Modification of EBIT Time Structure

Conceptually, the simplest method for achieving greater bunch separation would be to simply release the pulses from the EBIT at the desired time separation. If the trapping potential on the EBIT could be opened and closed again within the MHB period of 12.4 ns, with an acceptably low repetition rate between openings, that would immediately allow for any desired bunch separation with no additional loss of particles from the bunching method. However, this method turns out to be extremely difficult to achieve in practice.

In order to open and close the trap quickly enough, a square pulse would need to be applied to the trapping electrodes with a combined rise / sustain / fall time under 12.4 ns. For comparison, the fastest available switch from Behlke at the time of writing had a combined time of 62 ns with a maximum repetition rate of 3 MHz when operated in liquid cooled mode [32]. A prototype driver developed at CERN by Mauro Paoluzzi could

potentially have a very short pulse time on the order of 2 ns, with a voltage slew of up to 650 V [33]. However, the device operates with a 1 ms burst at a repetition rate of 2 Hz, giving it a 0.2% duty cycle, which was determined to be unacceptably low.

Due to these technical limitations, pure EBIT switching is not likely to provide a short term solution to the bunch spacing issue. However, some form of switching may be used in the future to produce improved results combined with another device.

2.3.2 Subharmonic Bunching

As constructed, ReA bunches the continuous output beam from the EBIT into the RFQ using a 80.5 MHz MHB. This fills each accelerating “bucket” of the RFQ and Linac. If a buncher were to operate instead at an integer divisor n of the RF frequency, then it could in principle bunch the beam into every n th bucket, increasing the spacing between bunches to n times the RF period of 12.4 ns [34].

There are two main tradeoffs to such an approach. The first is that this method of bunching necessarily increases the energy spread of the beam. The second is that unless the bunching waveform is a perfect sawtooth wave, some particles will end up in satellite bunches. Despite these potential drawbacks, this option was determined to be the most practical of the available options, and the design and construction of such a device for ReA is the subject of the remainder of this document.

2.3.3 Other Solutions

Two other possible approaches to increasing the bunch spacing were considered. The first was a beam chopper. A chopper in its simplest form simply deflects unwanted beam away

from the transverse acceptance of the accelerator. Regardless of the specific mechanism of the deflection, this approach was deemed to be undesirable because choppers by definition reduce beam transmission in proportion to the amount by which the spacing is increased. In other words, to increase beam spacing by a factor of 5, a chopper would need to discard 4/5 of the beam. This level of decrease in beam transmission was deemed to be unacceptable.

A more exotic solution would be to modulate the amplitude of the accelerating cavities along the entire length of the linac in such a way as to bunch the beam to a focus at the detectors. While intriguing, it is unclear that the cavity amplitudes could be modulated accurately on these sub-microsecond time scales, and the resulting energy spread would likely be quite large.

Chapter 3

Principles of Multiharmonic Bunching

The oldest and simplest types of particle accelerators are simple electrostatic devices. Particles are emitted from a source and are accelerated across a static voltage gap. However, these types of accelerators are limited by electrostatic breakdown to gap voltages on the order of 10 MV. For more powerful machines, it is necessary to use time varying (RF) electric fields for particle acceleration.

Since these time dependent fields alternate between accelerating and decelerating gradients over the course of a single period, it is necessary to first bunch the beam longitudinally before injection into such a device for acceleration. (Some RFQs are an exception to this rule, providing both bunching and acceleration in one device. The RFQ in ReA is not of this type.) In the simplest case, the continuous beam is bunched at the frequency of the accelerating device, and the bunch is timed to arrive near the peak of the accelerating gradient.

As will be discussed below, it is also possible to bunch the beam at any integer divisor of the accelerating frequency. It is a reasonable metaphor to think of each accelerating period of an RF accelerator as a “bucket”, carrying particles within that period forward through the accelerator. If the beam is bunched at half the frequency of the accelerating device, for example, then every other “bucket” of the accelerator will be empty.

Of critical interest in the design of any buncher is the acceptance of the subsequent beamline. The acceptance is defined as the region in phase space within which particles

will be successfully transported through the accelerator. In particular, the longitudinal acceptance in time / energy phase space is critical; particles must be compressed along the time dimension from their initial continuous distribution far enough to enter the time acceptance of the machine without being given so much energy they exceed the width of the energy acceptance.

3.1 Ideal Bunching

The basic idea of a buncher is simple: particles which would arrive earlier than the optimum phase of the accelerating device should be slowed down, and particles which would arrive later should be sped up, to a degree proportionate to their time separation from the center of the period. (As illustrated in Figure 2.1.) In the distance from the buncher to the accelerator, the beam will then come to a time focus, as the (now) slower particles at the front of the bunch and the faster particles at the rear converge on the unaffected particle at the center. This speeding up and slowing down is accomplished by means of an oscillating electric field directed parallel to the direction of beam travel.

An idealized buncher would have the following properties: The amplitude of the longitudinal electric field would follow a pure “sawtooth” waveform with a perfectly linear ramp and an instantaneous return to the starting voltage at the start of each period. The longitudinal field would be perfectly uniform in the transverse direction; there would be no deviation in field strength away from the beam axis. In addition, the boost from the field would be applied over an infinitely short distance, so that each particle would see no change in field in time taken to pass through the buncher. Finally, the buncher would be completely transparent to particles.

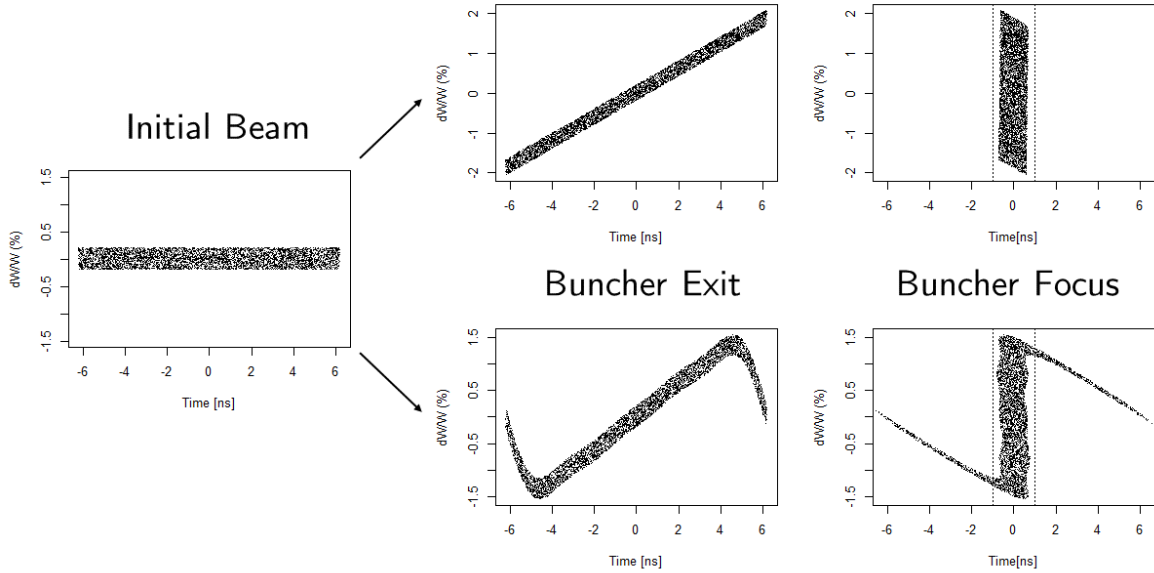


Figure 3.1: Comparison of the effect of an ideal buncher with a realistic bunching waveform. The upper plots show the beam bunched by a perfect sawtooth wave with zero reset time. The lower plots show a bunching waveform made up of four sine waves.

Therefore, the idealized buncher would be a pair of infinitely wide plates infinitely close together applying an infinite bunching voltage (to counteract the zero bunching distance) via a pair of infinitely thin infinitely permeable planes, with an infinitely fast reset time for the applied voltage. Thus far, no supplier has been located for such a device.

The pure sawtooth wave, in particular, is important. If the voltage could be applied with an instantaneous change of field direction at the start of each period, then all of the initial beam could in principle be brought to a focus in the final bunched pulse (Figure 3.1). The fact that this instantaneous change in field direction is impossible in practice is the direct cause of the “tails” of the particle distribution in longitudinal phase space visible in the second line of plots. The effects of these unwanted tails will be discussed below.

3.2 Simulation of a Sawtooth Wave Using Sine Waves

At first glance, it may appear that the simplest way to apply the bunching voltage to the beam would be to connect a function generator producing a sawtooth wave to an amplifier. However, the expected peak-to-peak voltage for our application is in the range of several kilovolts, so power requirements render this approach impractical. Instead, the waveform is typically approximated using a combination of sine waves which can be produced via a resonant cavity or resonant circuit.

3.2.1 Fourier Synthesis

Constructing a periodic function using a series of sine waves is the well-known technique of Fourier synthesis. To produce the desired function, a series of sinusoidal waves at integer multiples of the fundamental frequency of the original function are added together at calculated amplitudes to produce the desired function. The list of amplitudes required to produce a given function are referred to as the Fourier series for that function. While some functions can be completely constructed with a finite number of terms, the Fourier series for many periodic functions are infinite.

As an aside, Fourier series, which involve adding together only sine and cosine waves at integer multiples of a fundamental frequency, can technically only reproduce infinitely repeating periodic signals. To express a non-periodic signal in terms of sinusoidal components requires sine waves over a continuous spectrum of frequencies, with a continuous, rather than discrete, amplitude function. The extension of the technique from a sum over discrete frequencies to an integral over a frequency spectrum is referred to as a Fourier transform. For the purposes of the present discussion, the waveform to be represented will be assumed to

Fourier Synthesized Sawtooth Wave

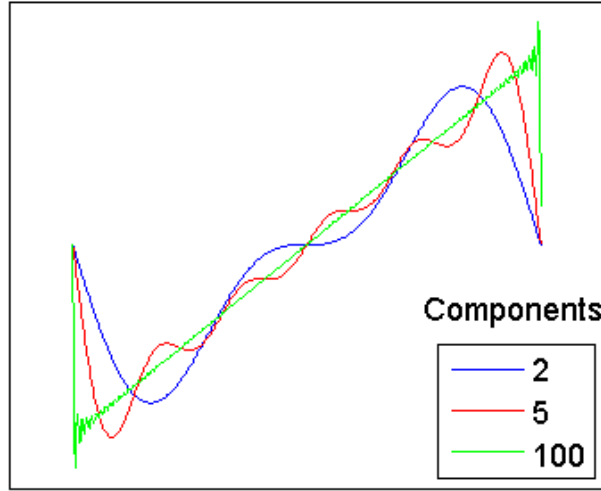


Figure 3.2: A sawtooth wave synthesized using the first 2, 5, and 100 components of the Fourier series.

be infinite, since in practice the total length of the bunching signal is much greater than the period of an individual repetition.

The ideal sawtooth bunching waveform desired for particle bunching is an example of a function with an infinite Fourier series expansion. If the zero point of the bunching waveform is taken to be the time origin, the Fourier series can be expressed solely in terms of sine waves. The amplitudes of each mode relative to the first mode are 1, -1/2, 1/3, -1/4, etc. In order to produce a completely perfect version of a sawtooth wave using sine waves, a device would require an infinite number of resonators with amplitudes set according to this series:

$$f(t) = \sin(\omega t) - \frac{1}{2}\sin(2\omega t) + \frac{1}{3}\sin(3\omega t) - \frac{1}{4}\sin(4\omega t) \dots \quad (3.1)$$

$$= \sum_{k=1}^{\infty} (-1)^{k-1} \frac{1}{k} \sin(k\omega t) \quad (3.2)$$

The effect of using fewer components can be seen in Figure 3.2. With only a few sine waves, the central region is not very linear. Further, the time extent of the “tail” area, where the value of the function is returning from the extrema to zero, is longer. As more components are added, the linear region becomes more linear and occupies more of the function period. Even with 100 components, fast oscillations are observed near the discontinuities at the end of the waveform, an artifact known as the “Gibbs Phenomenon” [35].

3.2.2 Window Functions

Since a real device can only use a limited number of resonators, the Fourier series must be truncated. This is the equivalent of applying a low-pass filter to the bunching waveform. The question arises: are the ideal amplitudes for a truncated series the same as those for the complete series? If they are not, what is the appropriate way to alter the amplitudes for optimum bunching?

Let us represent the change to each component as a series of coefficients a_k :

$$\sum_{k=1}^{\infty} a_k * (-1)^{k-1} \frac{1}{k} \sin(k\omega t) \quad (3.3)$$

In the simplest case, if we completely eliminate all truncated modes while keeping the remaining modes at their full Fourier amplitudes, then a_k is:

$$a_k = \begin{cases} 1 & k \leq n \\ 0 & k > n \end{cases} \quad (3.4)$$

where n is the number of retained modes. Effectively, this applies a rectangular low-pass filter to the sawtooth waveform. While this is the easiest case to understand, there is no

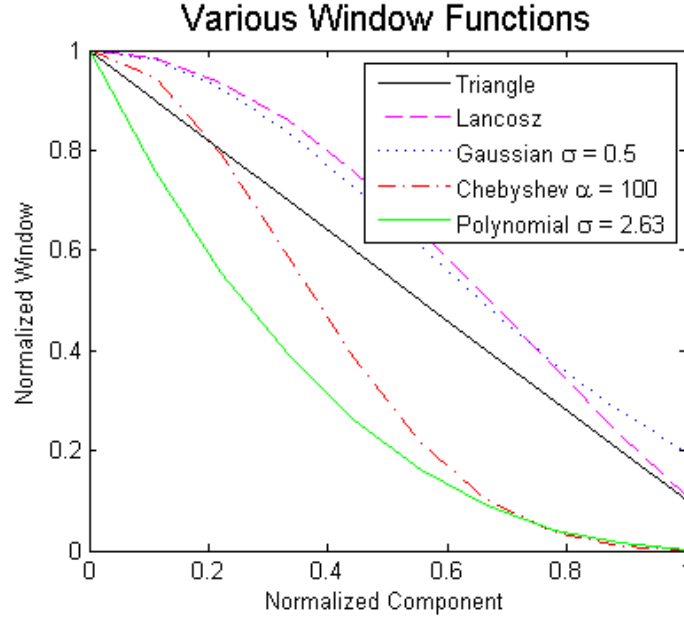


Figure 3.3: Window function amplitude vs. normalized component for several window functions. Normalized component = (component number - 1) / (number of components - 1) Plotted here for $n = 10$ components.

requirement that a_k be such a simple step function, and other functions may in fact give better results. Any arbitrary discrete function a_k may be used to describe the shape of the effective applied filter, and this function is known as a “window function.” (In the continuous case, a_k becomes $a(\omega)$.) The window function is always normalized to a maximum value of 1 and has a value of zero outside the frequency domain of interest.

There are a number of possible window functions described in the signal processing literature [36]. Selection of an appropriate window function is made based on the desired characteristics of the output signal. The stepwise function shown above is known as a rectangular window function, based on its shape. A number of window functions are illustrated in Figure 3.3.

A triangular window function, consisting of a linear decrease in the weight of each com-

ponent, would be expressed:

$$a_k = \begin{cases} \frac{n+1-k}{n} & k \leq n \\ 0 & k > n \end{cases} \quad (3.5)$$

Another commonly used window function is the Lancosz window:

$$a_k = \begin{cases} \frac{\sin(\frac{\pi(k-1)}{n})}{\frac{\pi(k-1)}{n}} & k \leq n \\ 0 & k > n \end{cases} \quad (3.6)$$

Several window functions have tunable parameters which can adjust the shape of the window. These include the Gaussian window:

$$a_k = \begin{cases} e^{-\frac{(k-1)^2}{2\sigma n}} & k \leq n \\ 0 & k > n \end{cases} \quad (3.7)$$

The polynomial window:

$$a_k = \begin{cases} (\frac{n-k+1}{n})^\sigma & k \leq n \\ 0 & k > n \end{cases} \quad (3.8)$$

And the Chebyshev window:

$$a_k = \begin{cases} (-1)^{k-1} \frac{\cos[n \cos^{-1}[\beta \cos(\frac{\pi(k-1)}{n})]]}{\cosh[n \cosh^{-1}(\beta)]} & k \leq n \\ 0 & k > n \end{cases} \quad (3.9)$$

$$\beta = \cos[\frac{1}{n} \cosh^{-1}(10^\alpha)] \quad (3.10)$$

By adjusting σ or α in these windows, the exact size of the window can be varied while

keeping the fundamental characteristics of each shape.

The selection of a window function can have a significant impact on the characteristics of the final synthesized waveform. The goal in the present case is to select a window that allows the greatest percentage of the beam to enter the longitudinal acceptance of the RFQ after the buncher. Different window functions result in longer or shorter lengths of the linear portion of the final bunching waveform, and in more or less “ringing” at the ends of the waveform. In addition, the best choice of window can vary depending on the energy spread of the initial beam; a beam with a very narrow energy spread can tolerate more “ringing” at the corners before losing bunched beam from the acceptance window.

3.2.3 Consequences of a Finite Number of Modes

Regardless of the specific window function chosen, a consequence of using finite number of modes to simulate the sawtooth bunching waveform is that there is a nonzero transition time from the maximum bunching voltage at the end of one period to the minimum at the beginning of the next. During that transition time while the electric field is slewing back to the starting value for the next period, the beam particles passing through the buncher are smeared into “tails” in phase space, the consequences of which are explored in Section 3.4.

3.3 Other Deviations from the Ideal in a Real Buncher

The previous section explored the synthesis of the bunching waveform, the reasons why it will deviate from the ideal, and the consequences of that deviation. In addition to deviating from a perfect sawtooth waveform, a real buncher deviates from the ideal buncher described above in several other important ways.




Figure 3.4: Side view of a pair of buncher electrodes with a conical profile. The arrow indicates the direction of beam travel

3.3.1 Transverse Field Flatness

The ideal buncher consists of two planes of infinite transverse extent, infinitely transparent to particles. These two planes, being infinite, will have a perfectly uniform accelerating field at all points between them. Since any real means of applying voltage in the direction of beam travel must include some sort of aperture to allow the beam to pass, this raises the question of how best to equalize the field in the longitudinal direction across the aperture.

Since an electric potential is to be applied in the beam direction, bunchers generally consist of a symmetric pair of electrodes with an opening in the center to allow for beam passage, for example, Figure 3.4. These electrodes are generally conical or cylindrical, and can employ several strategies for field equalization across the aperture, as well as across the gap between them. The aperture can be filled with grids, bars, or other conducting structures which equalize the applied voltage, at the cost of some impediment to the beam. Alternately, the aperture can be left open, and the deviations from homogeneity with distance from the beam axis simply tolerated. Which approach is preferable depends on the circumstances of the individual device.

3.3.2 Transit Time Factor

Another deviation from ideality is the finite time taken by the particles to traverse the accelerating voltage region. An ideal buncher would impart the bunching “kick” to each particle instantaneously, but a real electric field must be applied across a gap. And since the accelerating voltage is changing in time, the actual effect of the buncher must be modified by considering the effect of the change of field in the time it takes the particle to traverse the gap.

This effect, known as the “transit time factor” (TTF) is mathematically calculated as a coefficient multiplied to the overall energy gain of the particle across the gap. A TTF of 1 indicates that the particle sees no change in voltage across the gap, and is therefore only possible for DC fields. A TTF of zero indicates that the particle sees no net acceleration across the gap, and occurs when the time taken by the particle to cross the field is equal to an integer number of periods.

The formula for the TTF is:

$$TTF = \frac{\sin(\frac{\omega d}{2v})}{\frac{\omega d}{2v}} \quad (3.11)$$

where ω is the frequency of the applied field, v is the speed of the particle, and d is the length of the gap. This assumes a hard edged field with no field outside the gap, and Δv small enough that it can be neglected during the time the particle is traversing the gap. For the buncher which is the topic of this document, Δv is on the order of 4%.

Physically, the TTF is directly related to several factors: As the particle travels faster, the field has less time to cycle, and the TTF goes up. As the physical length of the field (determined by the gap between the electrodes) increases, the particle takes longer to cross

the gap, allowing the field to cycle further, lowering the TTF. For velocities which are high compared to ωd , as here, the TTF decreases monotonically as the frequency increases. As such, lower frequency modes have better TTFs than higher ones. Selection of electrode physical parameters such as aperture size and gap width is most directly affected by this parameter.

3.4 Satellite Bunches

When particles are bunched at the same frequency as the accelerator into which they are being injected, any particles which are not successfully moved into the time / energy acceptance of the accelerator will not be transported through the accelerator. However, if the buncher is operating at a subharmonic of the accelerator frequency, there is a potential for particles to enter an acceptance window other than the desired one. For example, an accelerator with a frequency of 60 MHz being injected with a 12 MHz buncher has five longitudinal acceptance windows for each period of the buncher.

Figure 3.5 illustrates the case of bunching at the 5th subharmonic. While the majority of the particles will be bunched into the desired acceptance window, and most of the rest will fall outside of any acceptance window and be lost, some particles will enter one of the undesired acceptance windows and be transported through the remainder of the accelerator. These particles form bunches which will be referred to as “satellite bunches.” Since the point of a subharmonic buncher is to increase separation between bunches transported through the linac, these satellite bunches are potentially problematic if their particle count is too high relative to the main bunch. What constitutes “too high” is determined on a case-by-case basis.




Figure 3.5: A plot of the time and energy distribution of a simulated beam bunched at 16 MHz. The dashed lines indicate the 2 ns acceptance windows for a device operating at 80.5 MHz.

There are a number of potential strategies for eliminating satellite bunches. While the detailed design and construction of a bunch cleaner for ReA is beyond the scope of this dissertation, presented here are a few possible approaches to such a device.

- Modification of Source Beam

The most direct way to eliminate satellite bunches would be to simply shorten the initial pulse delivered to the buncher (Figure 3.6). This approach would combine the subharmonic buncher with the trap switching described in Section 2.3.1. Instead of a constant DC beam, if the buncher were instead fed with a pulsed beam shorter in extent than the linear portion of the bunching waveform, there would be no particles passing the buncher while the sawtooth wave was transitioning from maximum to minimum. As such, no particles would be present to be smeared into tails.

- Off-Frequency Cavities

A buncher with a given frequency will create a beam structure at that frequency.




Figure 3.6: The effect of shortening the initial EBIT pulse. At left, the shortened pulse after interaction with the bunch. At right, the beam at the buncher focus. No tails are present because the shorter initial pulse falls entirely within the linear portion of the bunching waveform.

Such a beam can only be transported by subsequent devices with frequencies which are integer multiples of the buncher frequency. For example, the ReA RFQ frequency is five times the frequency of the buncher described in this document. This principle holds true for all subsequent devices as well - all must be integer multiples of the original frequency to transport the bunched beam. This can be exploited by utilizing devices in the beamline at two *different* integer multiples of the original frequency.

An example is illustrated in Figure 3.7. In this illustration, the filled circles represent particles at the fundamental frequency of the buncher and zero amplitude is taken to be the desired point on the accelerating wave. If the blue line represents a device which is at a 5:1 frequency ratio with the buncher, and the red line a second device at a 4:1 frequency ratio, the beam structure from the buncher would be successfully transported through either one in isolation. This is illustrated by the filled circles falling at the zero amplitude point of both waves.

The hollow circles in the illustration represent satellite bunches produced in the first device. Since they are at the zero amplitude point of the first device (blue line), they would be transported successfully by that device alone. However, the second device




Figure 3.7: Bunch cleaning by an off frequency cavity. The red line is at a 4:5 frequency ratio with the blue one. Black dots represent beam particles at a the zero phase of both frequencies.

has a 4:5 frequency ratio with the first one. This means that, assuming the acceptance windows are aligned in phase for the primary bunch to transit both devices, they will be out of phase for the satellite bunches. Particles which arrive at either device out of phase will not receive the same amount of acceleration as the primary bunch.

Since these bunches are unwanted, this accomplishes the goal of cleaning satellite bunches by moving unwanted bunches away from the target energy of the primary bunch. These bunches, now off-energy, will fail to traverse subsequent bending elements which are tuned for the reference energy. This is the bunch cleaning strategy employed by the ATLAS accelerator complex at ANL (See 3.5.2 below).

- Beam Choppers

A more direct approach is to simply push the undesired bunches off axis using an electric deflector. Such a device generally consists of a pair of parallel plates applying

a deflecting voltage that varies from zero at the main bunch to a maximum at the undesired bunches. The deflected bunches are then stopped by a slit inserted farther down the beamline. The device must be able to apply sufficient deflecting voltage to resolve the satellite and main bunches at the location of the slit, and must be able to cycle quickly enough to leave the main bunch untouched while deflecting the satellites.

The simplest chopping waveform conceptually would be a constant deflecting voltage that turns off with a square pulse for the bunch to be transmitted. Due to the fast rise times it is often more practical to use a sinusoidal chopping waveform and accept the additional bunching and concomitant energy growth of the main bunch. A more complicated form of the transverse chopper is the “stripline chopper” [37], which allows the square pulse to be applied over a longer distance, thus lowering the required voltage.

- Storage Rings

The most complicated approach to be mentioned here is a storage ring. A small storage ring inserted into the beamline would allow trains of ions from the ion source to be stored and compressed in time before being sent to the buncher. This has the same advantage of the modification of the source beam approach above, in that the buncher would not receive any particles outside of the linear portion of the bunching waveform. It also allows for a longer spacing of final bunches, since the ring can release the stored particles at a subharmonic of the buncher itself. It has the disadvantage of being logistically very complex.

3.5 Case Studies

Descriptions of two existing multiharmonic bunchers are presented here as illustrations.

3.5.1 80.5 MHz Multiharmonic Buncher at ReA

The ReAccelerator at NSCL is described in Section 1.1.6 above. The source beam is provided by either via an Electron Beam Ion Trap (EBIT) or a test stable ion source. In particular, the EBIT is used to store and charge breed radioactive ion beams from the coupled cyclotron facility at NSCL, and in the future from the linear accelerator of FRIB.

In either case, the input beam has an energy of 12 keV per nucleon (keV/u). The beam is first accelerated by a 4-rod type RFQ and then in a superconducting RF linear accelerator, both at a frequency of 80.5 MHz. The input beam pulse from either the test source or the EBIT is sufficiently long that it may be considered continuous compared to the 12.4 ns period of the accelerator RF.

Prior to injection into the RFQ, the beam is bunched by a multiharmonic buncher at the 80.5 MHz frequency of the accelerator. The MHB synthesizes its bunching waveform from the fundamental frequency and the first harmonic (161 MHz) of the accelerator frequency. Each mode is produced in a $\lambda/4$ mode coaxial resonating cavity. As initially designed, the 80.5 MHz cavity was also intended to be excited in the $3\lambda/4$ mode to produce a second overtone (241.5 MHz), but this mode proved difficult to tune properly, and is no longer used.

The electrodes used to apply the voltage to the beam have also undergone several changes from the original design. As originally conceived, the electrodes were a pair of copper cylinders with a fine wire mesh across the faces to equalize the field. The wire mesh proved

susceptible to damage from the beam, and in 2013 the electrodes were replaced with cylinders with solid faces and a 1 cm aperture in the center of each. These apertures contained very narrow bars aligned edge on to the beam, intended to equalize the electric field across the aperture. As the required voltage for beam bunching did not change after the replacement of the electrodes, it can be inferred that the bars are indeed providing an equivalent field. However, the smaller aperture and physical dimensions of the bars did pose a more substantial challenge for beam tuning.

Since the MHB is at the RF frequency, this buncher has no issues with satellite bunches - 12.4 ns of input beam are bunched into each RFQ period, or “bucket.” The small portion ($\approx 15\%$) of the beam which sees the bunching waveform while it is resetting from high to low forms a tail which is largely not transported by the RFQ.

3.5.2 12.4 MHz Multiharmonic Buncher at ATLAS (ANL)

The ATLAS accelerator at Argonne National Laboratory is also a heavy ion accelerator, fed by either an ECR stable ion source or by the CARIBU radioactive ion source. In either case, the input beam has an energy of 30.5 keV/u and is effectively DC. The accelerator consists of a 60 MHz RFQ, a series of split-ring type SRF cavities at 72.75 MHz and 97 MHz, and then cryomodules containing a number of $\lambda/4$ SRF resonating cavities at 72.75 MHz and 109.125 MHz [38].

Before the RFQ, the beam is bunched by a MHB with a frequency of 12.125 MHz [39]. Unlike the 80.5 MHz MHB in ReA, this buncher is at the fifth subharmonic of the RFQ. The lower frequency allows the bunching voltage to be generated using a resonant tank circuit, rather than a coaxial cavity. This buncher employs the first three harmonics of the bunching frequency in addition to the fundamental (12.125, 24.25, 36.375, and 48.5 MHz) to synthesize

the bunching waveform.

Since the bunching period is longer than the RF period for the RFQ, some beam does enter the acceptance window for adjacent buckets to the main bucket, and thus forms satellite bunches. However, since the first 15 accelerating cavities are at either a 6:5 or 8:5 frequency ratio with the RFQ, the out of phase bunches are strongly deflected from the energy of the central bunch, and do not transport through the bend after the first accelerating section.

The electrodes of the ATLAS MHB have both an aperture width and an electrode-to-electrode gap distance which is larger than the ReA3 80.5 MHz MHB. The ATLAS electrodes are also completely open, with no grids or bars to equalize the field. Practical experience with the device has shown that good bunching can still be achieved with this electrode design in this frequency range.

Chapter 4

Measurement of the Longitudinal Acceptance of the ReA RFQ

The most significant potential limitation on the parameters of a beam buncher is the acceptance of the devices downstream. As long as the longitudinal phase space area is conserved, it is possible to compress a portion of beam to an arbitrarily short duration, at the cost of an increased spread in beam energy and increased bunching voltage. However, the longitudinal acceptance of a typical accelerating device is not solely one-dimensional in time; there is also an energy component. Particles which deviate too far from the reference phase OR the reference energy will not be transported. Thus, in order to determine acceptable parameters for a bunching device, it is essential that the full two-dimensional longitudinal acceptance of downstream devices be well understood.

In the case of ReA, the next device with a longitudinal acceptance window downstream from the proposed buncher location is the room temperature Radio Frequency Quadrupole (RFQ) which serves as an injector to the superconducting linac. While the RFQ had been extensively simulated prior to this work [19], the acceptance had not been empirically measured. An important prerequisite to buncher design was therefore a direct measurement of the shape of the RFQ acceptance window in longitudinal phase space. Since a buncher has in principle very little effect on the transverse dynamics of the beam, direct measurement of




Figure 4.1: A simple illustration of a longitudinal acceptance window.

the transverse acceptance was not performed.

Thus, the objective of this measurement was to determine the longitudinal acceptance of the RFQ; that is, if an incoming particle has phase ϕ relative to the accelerating phase of the RFQ and energy E , for what values of E and ϕ will it be successfully accelerated by the RFQ, and for which values will it fail to be accelerated? This can be visualized in terms of a “window” on a plot of energy vs. phase. (Figure 4.1) The results of this measurement were presented as a poster at the International Particle Accelerator Conference 2014 in Dresden, Germany [40].

4.1 Measurement Methodology

Measurement was conducted using a stable beam of H_2^+ from the offline test source located upstream of the RFQ and the existing 80.5 MHz MHB. When unbunched, this beam has a very low energy spread (~ 18 eV RMS for H_2^+), and an adjustable central energy, making it a sensitive probe of the energy width of the acceptance window of the RFQ. With the buncher active, the bunches have a narrow time spread at the entrance of the RFQ, making

them good probes for the phase width of the acceptance.

4.1.1 Total Transmission As A Function of Energy

The first phase of the measurement used this beam as a probe for the overall transmission of the RFQ as a function of beam energy. To minimize the energy spread of the beam, the MHB was turned off for this portion of the measurement. This would never be done under normal conditions, as much of the unbunched beam falls outside of the phase acceptance of the RFQ and therefore is not accelerated. In this case, however, ample beam current was available from the test source and the narrow energy spread was essential for a precise measurement of the energy acceptance.

The beam was first set to the design energy of 12 keV/u and the RFQ, linac, and subsequent beamline tuned for maximum transmission. The beamline between the end of the linac and the measurement point included two magnetic dipoles. This point is important, because it allowed the dipoles to be used to filter out unaccelerated beam particles. The dipoles were tuned for a beam with a rigidity corresponding to an energy of 600 keV/u, the design output energy of the RFQ. The beam current was then measured on a Faraday cup designated “L130”. The accelerating cavities in the linac were turned off for the entirety of this measurement, although the solenoids were used for focusing the beam. At the time of the measurement, the third cryomodule of the ReAccelerator had not yet been installed.

These measurements were then repeated at a number of other beam energies (Figure 4.2). At each energy, the beamline was retuned for maximum transmission to L130. By measuring total transmission for the unbunched beam at each energy, a picture of the phase width of the longitudinal acceptance at various energies was built up. While the beamline was tuned for each energy, the phase measurement described in 4.1.2 below was also performed.

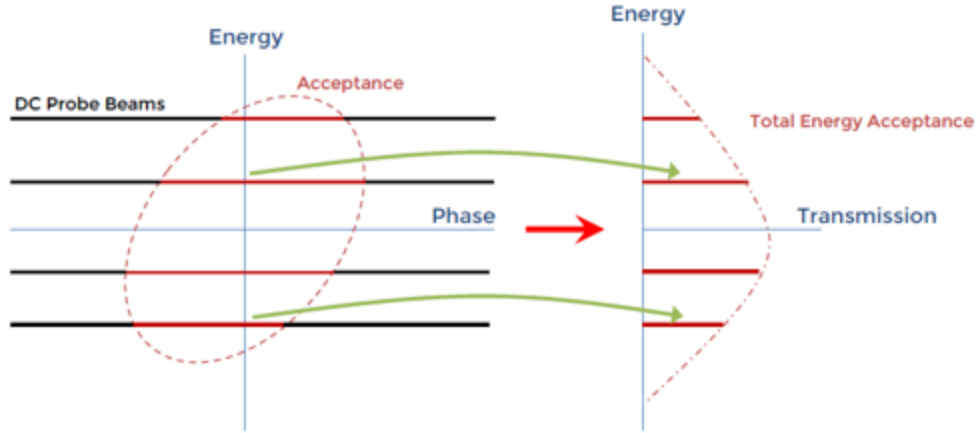


Figure 4.2: With the MHB turned off, the transmission through the RFQ was measured at a number of energies to give a picture of the relative transmission at each level.

On the first day of measurements, after the reference current at 24 keV/u was established, the input beam was adjusted until the measured transmission on the L130 Faraday cup reached a desired fraction of the original beam, such as 50%. Since adjusting the energy of the input source also affects the rigidity of the beam, and thus the tuning of the accelerator, this was actually an iterative process to reach the desired level of transmission - the source energy was changed, the beamline retuned, the source energy changed, etc. until the operator was confident that the lowest possible beam energy with 50% transmission had been reached.

Once good tunes were established for the accelerator at a range of energies using this process, more measurements were taken by simply scaling the beamline and input source voltage up and down. While not iteratively retuning for each measurement may have resulted in some beam loss, periodic checks were made to see if increased transmission could be achieved beyond simple scaling, and no significant increase was found.




Figure 4.3: With the MHB turned on, the transmission through the RFQ was measured at a number of energies and through a full 360 degrees of phase, giving a series of transmission curves.

4.1.2 Transmission as a Function of Phase and Energy

After the beam was tuned for each energy and the total transmission measured, the MHB was turned back on. This altered the beam from one with a continuous phase distribution and narrow energy spread to one with a much wider energy distribution and a narrow time spread. The relative phase between the MHB and the RFQ was then swept through a full 2π of phase. (In practice, it proved easier to adjust the phase of the RFQ than the MHB, but since there were no other RF elements involved in the measurement, this distinction is irrelevant.) The bunching amplitude of the MHB was not altered for measurements at different energies.

The transmission to the L130 Faraday cup was measured as a function of this relative phase (Figure 4.3). This gave the shape of the transmission as a function of phase for that energy level. It is important to note that while the energy probe was extremely narrow ($\Delta E/E \approx 0.001$), the time probe was not as focused. ($\Delta\phi/2\pi \approx 0.16$) The exact shape of the bunched beam in longitudinal phase space was also likely not a pure vertical rectangle.

$\Delta E/E_o$ (%)	Transmission (%)
-11.5	25
-8.8	50
-6.8	70
-4.5	90
-4.0	102
-2.0	99.8
0	100
1.99	99.9
3.99	98.9
5.87	90
9.19	66
12.47	50
15.59	25

Table 4.1: Summary of beam transmission through the ReA RFQ as a function of deviation from the reference energy.

Despite these limitations, a reasonably clear picture of the relative transmission as a function of phase for each energy level was obtained.

4.2 Analysis and Results

4.2.1 Energy Measurement Only

The first stage of building a picture of the window uses only the measurements of the unbunched beam. The transmission percentage at each energy is shown in Table 4.1 and Figure 4.4. This percentage is normalized to 100% transmission at the reference energy.

With just these measurements, it is possible to construct a first approximation of the acceptance window. Since no phase information is yet being incorporated, this first model must assume that the transmission as a function of phase is perfectly square. In other words, if 50% of the particles are transmitted, the model assumes that the phase window is 180 degrees wide, and symmetric about zero phase. It also assumes a hard edge - all particles




Figure 4.4: RFQ beam transmission vs. input energy for ReA RFQ.

within the phase window transmitted, and none outside the window. For example, at the reference energy of 12 keV/u, the ratio of beam current at L130 to the injected current was measured to be 0.4. Therefore, the equivalent width of the transmission window at the reference energy is taken to be $360^\circ * 0.4 = 144^\circ$ or 4.96 ns.

The error on the transmission is estimated as follows: The dominant source of error in the transmission at each energy is assumed to be tuning inefficiency. While the beamline was being tuned by an experienced operator, due to time restrictions, once the transmission appeared to be near a maximum for each energy, tuning was halted and a measurement was taken. A reasonably conservative estimate for the further refinement possible at each level is an increase of up to 5% in transmission. Since this error was also present for the original measurement against which the others are compared, an error of $\pm 10\%$ is assumed in the final transmission amount. (This also explains how a *relative* transmission percentage greater than 100% can be computed for a measurement away from the reference energy.)

This naive acceptance window is shown in Figure 4.5.

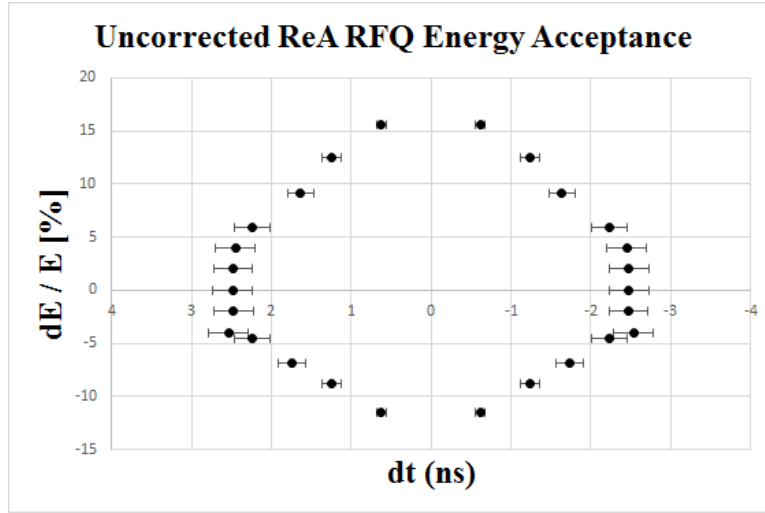


Figure 4.5: Naive RFQ acceptance window with no phase measurement or time-of-flight correction.




Figure 4.6: Three phase scans at different energies.

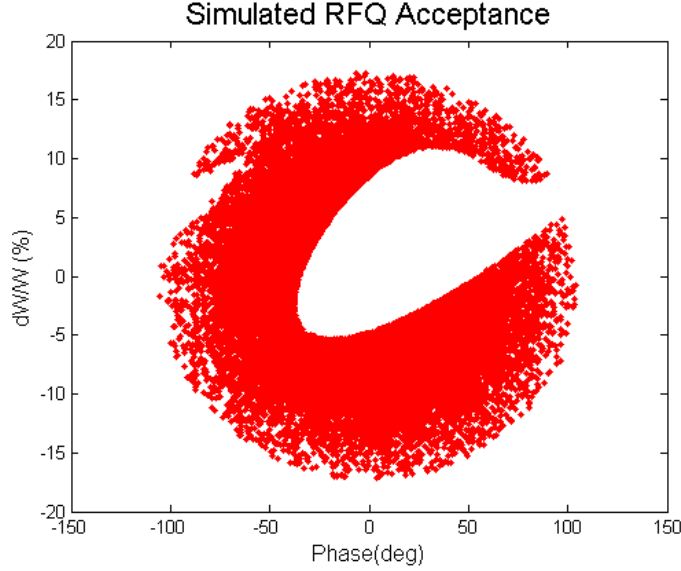


Figure 4.7: The simulated acceptance of the ReA RFQ. The white area at the center represents particles which are accelerated by the RFQ, and the red area shows particles which are not transported.

4.2.2 Combined Energy and Phase Measurement

Plotting the acceptance window based entirely on total transmission gives a reasonable first approximation of the energy window, but the phase dependent data gathered at each energy level allows for a somewhat more detailed picture. As the bunched beam at each energy was swept in relative phase, the beam transmission was measured as a function of that phase. Figure 4.6 shows several measurements of the transmission during a phase sweep. The transmission curve for each energy is normalized to a maximum of 1. One interesting feature that is apparent in the highest energy curve is second peak in the phase distribution. This additional peak is predicted by the DYNAC simulated acceptance window shown in Figure 4.7, but had not previously been directly observed for this machine.

There are a number of imperfections in this measurement that are not present in the energy measurement. A perfect version of this measurement would require beam bunches with both very narrow energy and time spreads to give the best possible resolution to the

measured data. When the beam is bunched, the time distribution of the beam is an approximately Gaussian shape with a width of $\approx 2\text{ns}$. As such, even if the RFQ phase acceptance for a given energy were a perfect step function, we would still expect to see a gradual rise and fall in transmission as the beam phase was swept across the acceptance. Further, the bunched beam has a computed energy spread on the order of $\approx \pm 2\%$. Since the phase acceptance width narrows farther from the reference energy, the measured transmission will see a greater contribution from the portion of the bunch with energy closer to the reference.

Both of these deviations from the ideal measurement could in principle be resolved given a perfect knowledge of the exact longitudinal phase space shape of the probe beam. Such a shape could then in principle be deconvolved with the measured transmission to give a fully accurate profile of the RFQ window. However, no simple way to measure the full 2D profile of the beam in longitudinal phase space was available.

Instead, the phase acceptance width will be defined here as the full width of the transmission curve at which the transmission is one half of the maximum transmission at the design energy. In other words, if the maximum current read on the Faraday cup at the design energy is 500 pA, then the phase acceptance width $\delta\phi$ is defined to be the difference between the phases ϕ_- and ϕ_+ where the current read is 250 pA. This value of current is applied at every energy level, so that $\delta\phi = \phi_+ - \phi_-$ becomes smaller as the beam energy is moved away from the design energy, and eventually reaches zero when the maximum measured current drops below half of the value measured at the design energy. The measured phase widths by energy are shown in Table 4.2. The following paragraphs are a visual illustration of the same information.

Using this information, a full 2D model of the acceptance can be constructed. First, the phase curves are superimposed (Figure 4.8). Then, the phases are adjusted so that the

$\Delta E/E_o$ (%)	Phase Width (°)	Phase Width (ns)	Phase Offset (ns)
-11.5	0	0	0.739
-8.8	40	1.4	-0.284
-6.8	75	2.6	0.325
-4.5	105	3.6	-0.165
0	140	4.8	-
5.87	95	3.3	-0.875
9.19	70	2.4	-0.456
12.47	50	1.7	-1.16
15.6	0	0	-0.291

Table 4.2: Width of phase acceptance at various energies for the ReA RFQ, along with the measured phase offsets.

maximum transmission of each curve is aligned at zero (Figure 4.9).

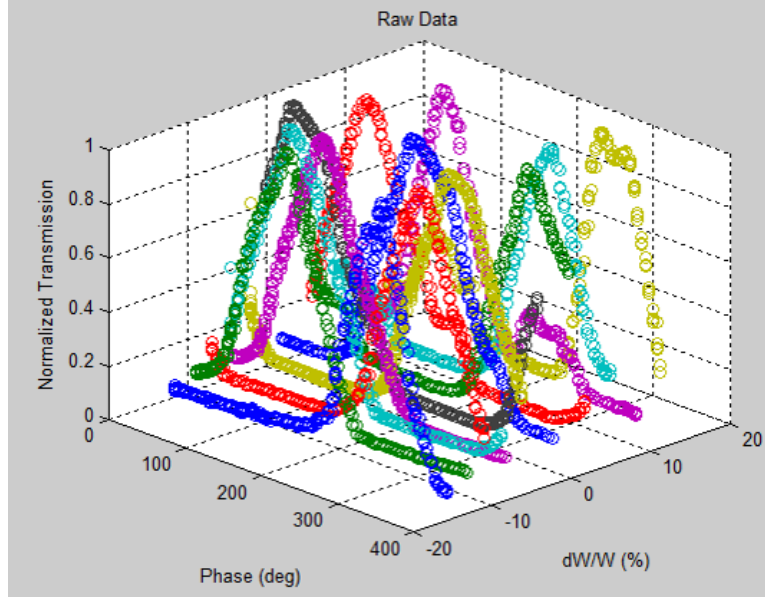


Figure 4.8: All phase measurements superimposed. Each curve is normalized to a max of one, and phase units are unadjusted.

Next, each individual curve is scaled down by the relative transmission at that energy level from Figure 4.4, as shown in Figure 4.10.

Finally, the 50% transmission contour of the interpolated surface in Figure 4.10 can be extracted. This allows Figure 4.5 to be updated to show measured deviations from zero phase, instead of assuming a symmetric, hard-edged transmission profile. The result is shown

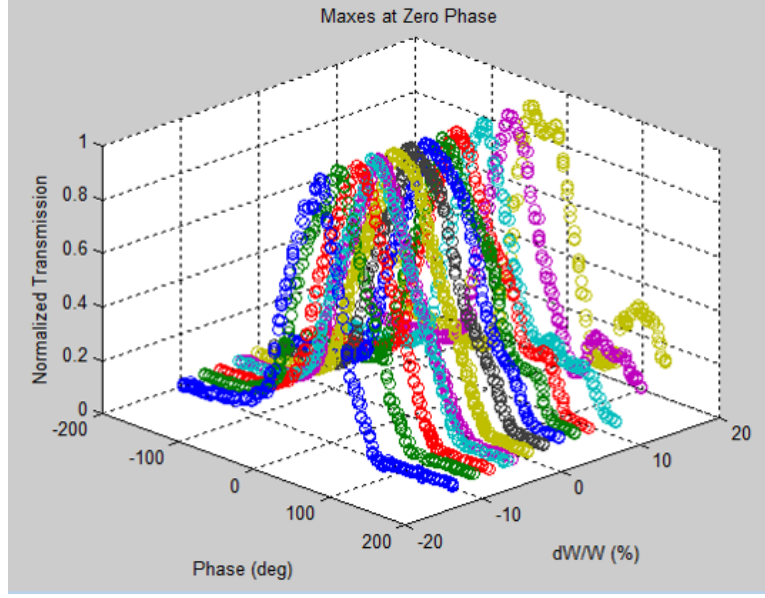


Figure 4.9: The phase curves in figure 4.8 with the maximum transmission for each energy aligned at zero phase.

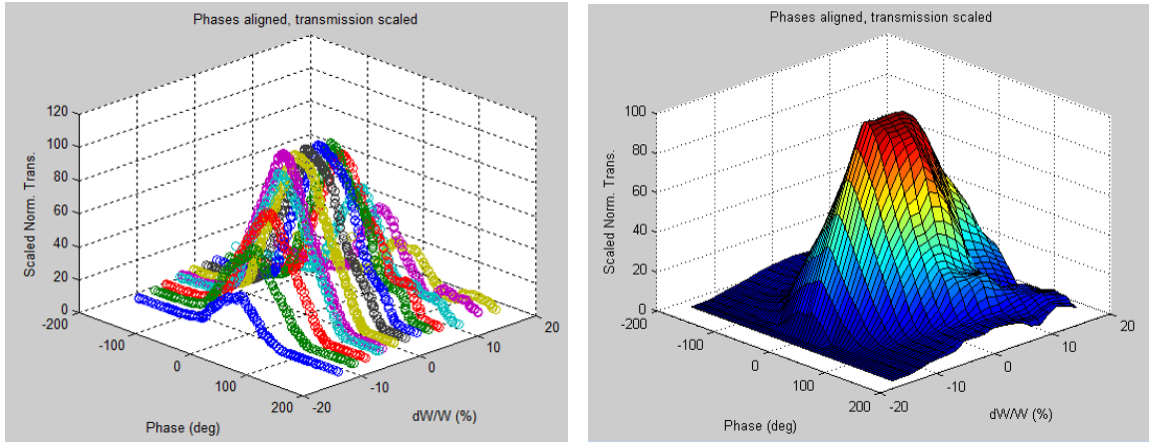


Figure 4.10: Phase curves with zeroes aligned and scaled by overall transmission at each energy. The second figure is an interpolated surface generated from the data points in the first figure.

in Figure 4.11, overlaid on the simulated RFQ acceptance window from DYNAC which was used in the remainder of this project. This contour is the equivalent of interpolating between the data points listed in Table 4.2.

One further correction can be applied. Beam particles at different energies are naturally traveling at different velocities, meaning that their arrival times at the RFQ are also different.

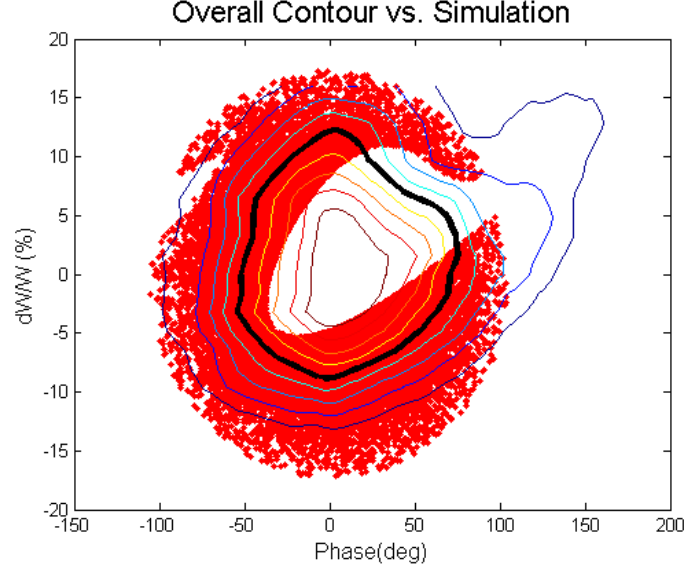


Figure 4.11: A contour plot of the measured acceptance surface from Figure 4.10, with the 50% transmission contour emphasized, overlaid on the simulated acceptance.

The actual time of flight from the buncher to the entrance of the RFQ for each energy can be calculated from basic dynamics. As such, for each energy away from the reference energy, a predicted *difference* in time of flight, Δ_{tof} , can be calculated based solely on the difference in energy. This is then converted to a predicted difference in RF phase (in degrees), Δ_ϕ

$$\Delta_\phi = \Delta_{tof} * f * 360. \quad (4.1)$$

where f is the RF frequency of the accelerator. Δ_{tof} is found by

$$\Delta_{tof} = \frac{L}{v_o} - \frac{L}{v}. \quad (4.2)$$

Here L is the distance from the buncher to the RFQ, v_o is the velocity of the reference

dE/E (%)	γ	$v(\text{m/s})$	MHB-RFQ tof	Δ_{tof} (ns)	Δ_ϕ (°)
-11.5	1.000011402	1431618	508.306	-30.094	-872.127
-8.8	1.000011746	1453032	500.815	-22.602	-655.034
-6.8	1.000012005	1468975	495.38	-17.167	-497.522
-4.5	1.000012299	1486865	489.419	-11.207	-324.786
-4.0	1.000012366	1490877	488.102	-9.890	-286.617
-2.0	1.000012623	1506327	483.096	-4.883	-141.537
0	1.000012883	1521711	478.212	0	0
2.0	1.000013139	1536760	473.529	4.683	135.7151
4.0	1.000013396	1551753	468.954	9.258	268.3036
5.9	1.000013638	1565715	464.772	13.440	389.4957
9.2	1.000014066	1590073	457.652	20.559	595.8216
12.5	1.000014489	1613776	450.93	27.281	790.624
15.6	1.00001489	1636004	444.803	33.408	968.1759

Table 4.3: Calculated Δ_{tof} and Δ_ϕ values from the MHB to the RFQ relative to the reference particle.

particle, and v is the particle velocity. These values are summarized in Table 4.3, where γ represents the relativistic gamma of the particle given by

$$\gamma = \frac{E}{mc^2} + 1. \quad (4.3)$$

(As a reminder, E represents *kinetic* energy throughout this document.)

These calculated phase differences at various energy levels can now be compared with the measured data. While no information is available on the *absolute* time of flight from the MHB to the RFQ, we can compare the *relative* phases between the peak phase of each measurement, and the peak phase at the reference energy. These relative phases can then be compared with the predicted phase differences shown in Table 4.3. Since our observed phases are all in the range of 0 - 360, we can either compare relative phases modulo 360 or add or subtract integer multiples of 360 to correspond to the varying velocities. Figure 4.12 compares the predicted versus observed phase advances in the latter manner for ease of

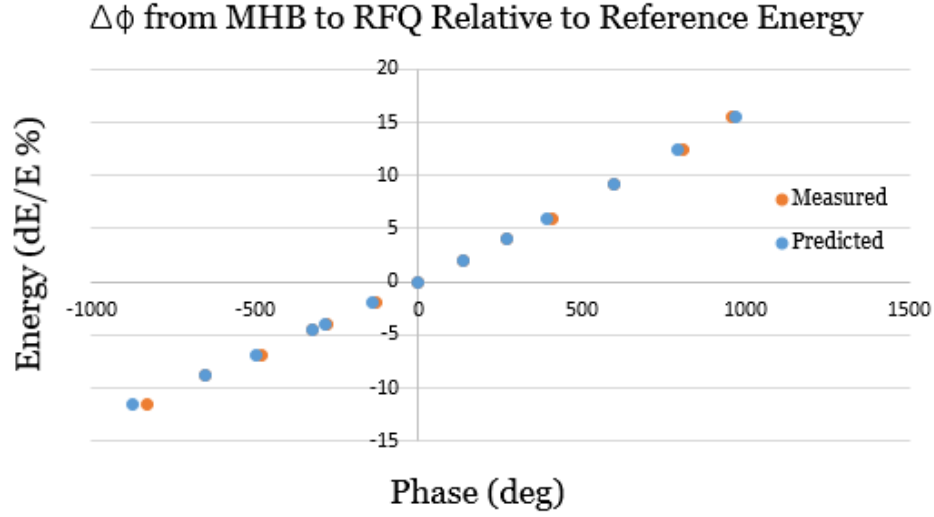


Figure 4.12: Measured and predicted phase advances relative to the reference energy for the phase measurements.

visualization.

For an ideal measurement, all of the difference between the calculated and measured phase advances can be attributed to a “tilting” of the acceptance window on the phase plot. If the peak transmission phase is observed to arrive earlier than predicted, the window is shifted earlier for that energy. If the peak transmission phase is observed later, the window is shifted later. For a perfect measurement with a window following an linear “tilt” the differences in phase advances would be expected to follow a linear trend with energy. Unfortunately, while this trend is *somewhat* visible in the actual data, there is a great deal of divergence from a smooth linear increase, as seen in Figure 4.13. There are a number of possible explanations for this divergence from ideality: the shape of the probe beam is certainly not a perfect vertical rectangle, and neither is the acceptance window. Further, the MHB was not reoptimized for best transmission at the various energy levels.

If the offsets shown in Figure 4.13 are applied to the phases of the measured transmission curves despite the noise, the resulting acceptance window is shown in Figure 4.14. Comparing

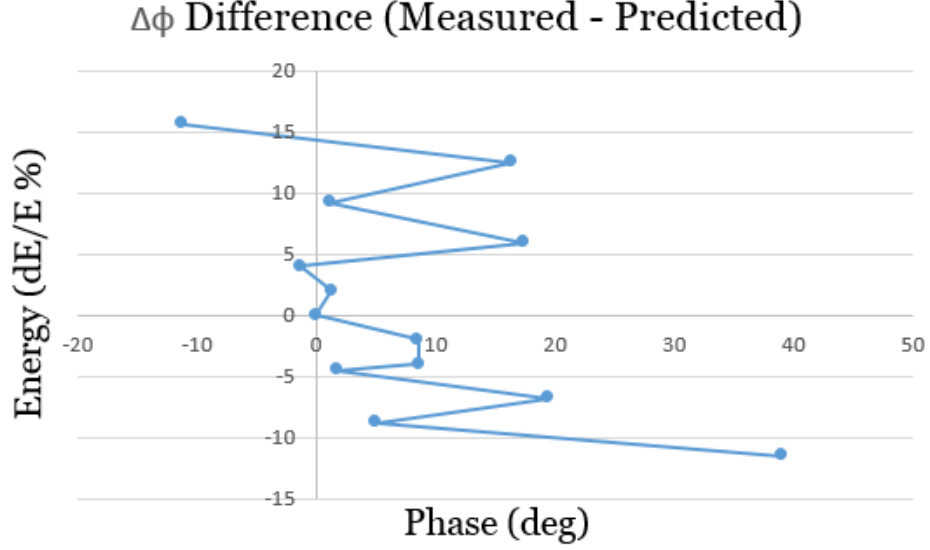


Figure 4.13: For each energy, the offset from the predicted relative phase advance.

this with the uncorrected contour plot in Figure 4.11 the window does display a tilt on the same order as the predicted tilt from simulation.

4.2.3 Error Estimates

For the measurement of transmission using the unbunched beam, it is unlikely that there was significant error in the beam energy. The power supply regulating the ion source has been measured to vary by less than $\approx 0.1\%$. A larger source of error in the transmission rate results from the tuning of the beam. While the beamline was retuned for maximum transmission at each energy level by an experienced operator, there is nonetheless a possibility that greater transmission would have been possible given further tuning. Therefore the measured transmission at each energy level should be regarded as a minimum.

A reasonable estimate would be that the transmission achieved at each level was within 5% of the maximum achievable. During the phase scans, the phase of the RFQ relative to the MHB was increased in increments of five degrees every two seconds. The error in




Figure 4.14: The contour plot of the measured acceptance surface with the time-of-flight correction applied, overlaid on the simulated acceptance.

phase for each measurement was thus estimated as \pm one phase increment or ± 5 degrees. However, as has already been noted, this error does not include the irregular shape of the probe beam, and is almost certainly a lower bound.

4.2.4 Implications for the Buncher

While it is disappointing that the time shape of the buncher could not be measured more precisely, the actual results of the measurement for this project were deemed to be highly encouraging. Both the time of flight corrected and uncorrected measurements showed an acceptance window somewhat larger than the hard-edged window of the simulation. In addition, the pure energy measurement showed at least the possibility for 100% transmission relative to the reference energy out to $\pm 5\%$ away from 12 keV/u, given an appropriate time structure of the injected beam.

As the DYNAC simulated acceptance window shown in the figures above is smaller than the measured acceptance window, these measurements validated using the simulation as a

conservative constraint on buncher parameters going forward. A real acceptance window does not have the hard cutoff of the simulation, but since the hard edge in this case is displaced inward from the measured transmission contour, it seems safe to conclude that most particles which pass through the simulation are also likely to pass through the real machine.

Chapter 5

Simulation and Selection of Design Parameters

From a beam optics standpoint, a buncher is a relatively simple device. In the ideal case, it has no transverse effect on the beam at all, and applies a simple sawtooth accelerating/decelerating voltage in the longitudinal direction. In this case the only output parameters which must be chosen are the acceptable time and energy spread of the pulses at the buncher focus, and the frequency of the bunched beam. The only input parameters which need to be chosen for this ideal buncher are the focal length and voltage needed to achieve the desired output. Further, those two parameters are not independent, as outlined below in Section 5.3.

For a realistic buncher, in addition to the focal length and bunching voltage, additional questions must be considered. What shape should be chosen for the electrodes applying the bunching voltage? Where is space available on the beamline to install the device? How can the appropriate frequencies be applied to the electrodes? What amount of power is practically available to drive the electrodes? What is the maximum amount of beam which can be allowed to fall outside of the acceptance of the accelerator?

This chapter will describe the process by which these design decisions were made for the ReA 16.1 MHz buncher.

5.1 ReA Beamline Simulation

Since the purpose of a particle accelerator is to supply beam to users, the final judgment on the effectiveness of the accelerator rests on the quality of that beam at the detector. Therefore, in order to evaluate any proposed change to an accelerator, it is important to be able to evaluate the effect of that change at the detector focus. (For simple changes, it would be sufficient to show only that the beam could be returned to its prior state at some point before the end of the line, but that is manifestly not the case here.) In order to ensure that the effect of the buncher on the beam at the detectors was well understood, it was necessary to simulate the entire ReA beamline from the injection of the beam from the EBIT to the detectors in the experimental hall.

This simulation proved challenging for several reasons. First, there are a wide range of devices along this line, including electrostatic and magnetic quadrupoles and dipoles, an RFQ, 15 two-gap accelerating cavities, as well as the buncher itself. The RFQ in particular posed a particularly thorny simulation challenge, as its complicated field geometry requires specialized software for simulation. Further, since bunchers by definition operate in the time domain, any simulation must properly account for longitudinal dynamics at multiple frequencies. Finally, while many of these devices can be approximated by simple analytical models such as those in Section 1.2, a more detailed numerical treatment is required in certain areas, such as the SRF cavities and for evaluating specific buncher electrode geometries.

A number of codes are available for beamline simulation, and several are compared in detail in Appendix A. In order to limit potential sources of error (as well as time and effort) it was preferable to use as few codes as possible. Where necessary, a simulated particle distribution can be extracted from one code and input to another, but one goal was to keep




Figure 5.1: A sample longitudinal model of ReA from the EBIT to the detector hall, showing x and y envelopes, particle count, reference particle energy, and the locations of beamline elements.

this type of code switching to a minimum.

For this project the primary code selected was DYNAC [41]. DYNAC has the advantage of being able to simulate the entirety of ReA, including all beamline elements, from the exit of the EBIT to the detectors. An end-to-end model of the ReA line was created using DYNAC, and used to test the result of various buncher configurations at the detector. A sample longitudinal plot of various beam parameters produced in this model is shown in Figure 5.1. Over the course of the project, a graphical front end to DYNAC was also developed using the MATLAB environment. This front end, DYNACGUI, is described in Appendix B.

Other codes were used for tasks unsuited for execution in DYNAC. The code TRACK, developed at Argonne National Laboratory [39], was used for modeling the effects of specific buncher electrode geometries. TRACK has the advantage of being able to use 3D numerical field models for beamline elements, in addition to basic analytic models. The use of TRACK to select electrode geometry is explored in Section 5.5. TRACK was not selected as the overall code of choice because while it does possess the ability to simulate RFQs, it proved too challenging to translate the RFQ specifications for the ReA RFQ into TRACK's format. Other

codes were used for tasks unsuited for execution in DYNAC. The code TRACK, developed at Argonne National Laboratory [39], was used for modeling the effects of specific buncher electrode geometries. TRACK has the advantage of being able to use 3D numerical field models for beamline elements, in addition to basic analytic models. The use of TRACK to select electrode geometry is explored in Section 5.5. TRACK was not selected as the overall code of choice because while it does possess the ability to simulate RFQs, it proved too challenging to translate the RFQ specifications for the ReA RFQ into TRACK's format.

One task that is often called for in optics simulation is the fitting of certain beamline settings to achieve a desired result, such as minimum transverse focus at a detector. DYNAC has no internal fitting capability, so the code COSY [23] was used in some cases to make preliminary fits. COSY has excellent internal fitting capabilities, but unfortunately the publicly available versions are not ideally suited to multiple frequency structures in the time/energy domain. Where refined fitting needed to be done using the DYNAC model, an extension to DYNACGUI was developed using MATLAB to do automated fitting using MATLAB's internal fitting tools to call DYNAC as an external function.

Finally, other programs were sometimes used for specialized tasks. These will be mentioned in the appropriate sections.

5.2 Buncher Frequency

As discussed in 2.3.2, it is possible to bunch the beam at any integer divisor of the frequency of the succeeding elements. If the reference frequency of the accelerator is f_{rf} , with a corresponding period of τ_{rf} , then bunching at subharmonic n entails a frequency of f_{rf}/n , and a period of $n * \tau_{rf}$. A first approach to frequency choice might then be to select the




Figure 5.2: Schematic illustration of bunching at several different subharmonics. Each plot shows energy vs. time. The first column shows the initial beam, the second column the beam immediately after bunching, and the third column shows the beam at the focus.

desired bunch separation τ and calculate the nearest appropriate integer n .

However, there are constraints on such a procedure. Larger values of n and τ correspond to bunching longer segments of the continuous beam. The longer the beam segment to be bunched, the greater the positive and negative voltages required to bring the fore- and rearmost particles to a focus in a given distance, corresponding to a greater energy spread in the final beam. If the value of n chosen is too high, it may not be possible to bring the entire beam to a focus within the energy acceptance of the accelerator at all.

This is illustrated in figure 5.2, which shows a simple model of energy vs. time for an originally continuous beam. The first line illustrates bunching with $n = 1$, at a frequency of f_{rf} . Column 1 shows the beam before bunching. Column 2 illustrates the beam immediately after bunching at f_{rf} . The particles with positive phase (i.e., behind the center of the bunch) have been increased in energy, and those with negative phase have been decreased in energy. The final column shows the beam at the focus of the buncher, where the trailing particles have caught up with the center of the bunch, and the leading particles have fallen back to

the center. The bunch separation is the same as the bunching period, τ .

The second and third rows show bunching at lower subharmonics, $n = 2$ and $n = 5$ respectively. In order to bunch more beam in the same focal length, the leading and trailing ends of the bunch must have their velocity altered by correspondingly larger amounts, leading to a larger energy spread. (The phase of the final diagram in the third row has been shifted by 180° relative to the others in order to show the separation between bunches.) In each case, the bunch separation at the focus is $n * \tau$.

A good way to make a simple estimate is via the longitudinal emittance. For a simple continuous beam, define the longitudinal emittance as $\epsilon = \tau * \Delta E$, where τ is the time extent of the beam to be bunched, and ΔE is the energy spread of the beam. If τ is chosen so that ϵ is larger than the acceptance window of the accelerator, there is by definition no way to bunch the entire beam into the acceptance window. The converse, however, is not necessarily true; it is possible to have a beam with an ϵ smaller than the acceptance window which nonetheless cannot be easily bunched in its entirety into the window, due to the particular shape of the window.

In order to minimize the required bunching voltage and beam energy spread, it is necessary to choose the lowest value of n which can still meet experimental demands, and then confirm that a beam bunched at this frequency can still fit within the acceptance of the accelerator (see Chapter 4). Other technical constraints, such as resonator geometry, may also mitigate against the construction of a buncher at certain frequencies.

For the case of ReA, a pulse separation of 100 ns was originally requested by some experimentalists. This would correspond to an n of about 8, and a frequency of about 9 MHz. Concerns about the energy spread and required voltage for bunching this amount of beam lead to a compromise choice of $n = 5$, corresponding to a period of 62 ns and a

bunching frequency of 16.1 MHz. It was agreed that this frequency would still present a substantial improvement in the potential for time-of-flight measurements at ReA, without posing insurmountable technical challenges.

5.3 Voltage and Focal Length

With the selection of 16.1 MHz as the fundamental frequency for the buncher, the next step is the determination of the location and bunching voltage for the device. The two parameters are not independent - for a given focal length, there is (to first order) only one optimal bunching voltage.

5.3.1 Ideal Case

The simplest approximation is to use an input beam with no energy spread. Let the incoming beam have velocity βc , the distance to the focus be L , and the time extent of beam to be bunched be τ_{bunch} , or 62 ns for the case at hand. The reference particle is located at the center of the bunch with an unchanged velocity, and will therefore reach the focal point at time t , where $t = 0$ corresponds to the front of the bunch reaching the buncher:

$$t = \frac{L}{\beta c} + \frac{\tau_{bunch}}{2} \quad (5.1)$$

For an ideal sawtooth bunching waveform with no reset time, an entire buncher period of the incoming beam will be bunched at the focus. As a result, both the foremost and rearmost particles must also reach the buncher at time t . The foremost particle must travel

a distance L in time t , and so must be decelerated to velocity:

$$v_- = \frac{L}{t} = \frac{2L\beta c}{2L + \beta c\tau_{bunch}} \quad (5.2)$$

Correspondingly, the rearmost particle must travel a distance of $L + \beta c\tau_{bunch}$ in time t , and so must be accelerated to velocity:

$$v_+ = \frac{L + \beta c\tau_{bunch}}{t} = \frac{2L\beta c}{2L + \beta c\tau_{bunch}} + \frac{2(\beta c)^2\tau_{bunch}}{2L + \beta c\tau_{bunch}} \quad (5.3)$$

If we use the non-relativistic approximation (appropriate in the case of this buncher, as $\beta = 0.00508$), then the change in kinetic energy for each of these particles is simply

$$\Delta KE_{\pm} = \frac{1}{2}m(v_{\pm}^2 - (\beta c)^2), \quad (5.4)$$

and the voltage required to produce that change of energy can be found simply from the Lorentz force law (where ΔS is the accelerating gap size):

$$qE_{\pm} = F_{\pm} \quad (5.5)$$

$$q\frac{V_{\pm}}{\Delta S} = F_{\pm} \quad (5.6)$$

$$qV_{\pm} = F_{\pm}\Delta S \quad (5.7)$$

$$qV_{\pm} = \Delta KE_{\pm} \quad (5.8)$$

$$V_{\pm} = \frac{1}{2}\frac{m}{q}(v_{\pm}^2 - (\beta c)^2) \quad (5.9)$$

As an example, if we use the original configuration of the ReA 80.5 MHz MHB, with $L = 0.73$ m, $\beta = 0.00508$, and $\tau_{bunch} = 12.4$ ns, for a helium beam ($Q/A = 1/4$), we find a

peak bunching voltage of ± 1230 V.

The most important fact to observe from this simple model is that for the low β case where the nonrelativistic approximation is appropriate, the velocity and hence the bunching voltage is inversely proportional to the focal length. This can also be understood intuitively in terms of phase space rotation; the longer the distance in which the rotation must occur, the slower it must take place, and the lower the necessary voltage to accomplish that rotation.

5.3.2 Numerical Simulation

Other than for the general insights just mentioned, this ideal case is of little use even for a first approximation, for two reasons. The first is that since the sawtooth wave is in practice built from its sinusoidal components, the ideal derived peak voltage will be well above the actual peak voltage applied to any individual mode. Of much more practical interest is the peak voltage which will need to be applied to the individual sinusoidal components of the wave, as described in Section 3.2.

Secondly, this simplistic approach fails to take into account the fact that the beam has a nonzero energy spread. The rotation in phase space shown in Figure 5.3 means that the wider the spread in energies of beam particles as they reach the buncher, the greater the time spread of those particles at the focus. Assuming the initial energy spread has no time dependence, the energy spread does not affect the level of the optimum voltage, but it does affect the efficiency with which particles can be bunched. The wider the initial energy spread, the less tightly compressed it is possible to make the final bunch in time. This can also be understood simply in terms of conservation of longitudinal emittance, and is illustrated in Figure 5.3.

For the actual determination of the optimum bunching voltage and focal length, a nu-

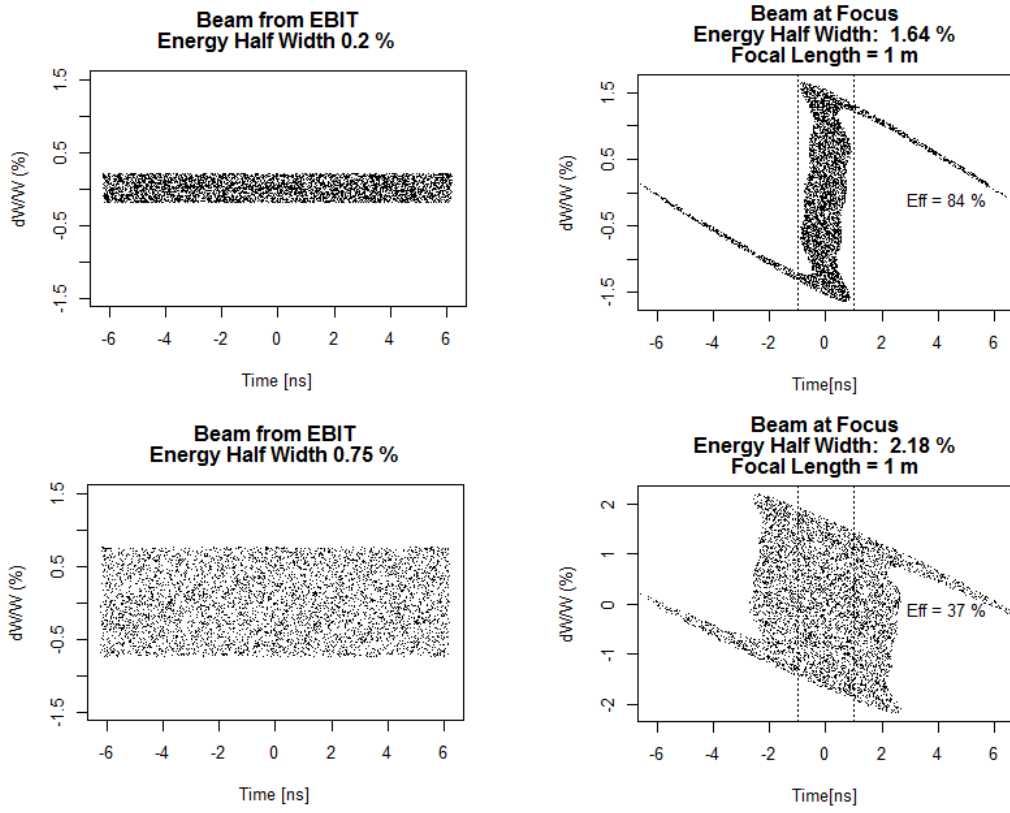


Figure 5.3: Simulated longitudinal beam phase space for two different values of initial energy spread.




Figure 5.4: Peak bunching voltage vs. focal length for a Monte Carlo model of a 16.1 MHz buncher.

merical Monte Carlo simulation was made in the statistics code R [42]. A large number of simulated particles was generated. These particles were uniformly distributed in time within the length of a single 16.1 MHz bunch and in energy within $\frac{\Delta E}{E} = \pm 0.2\%$, the design specification for the energy spread from the EBIT. The energies of those particles were then adjusted by applying a bunching waveform consisting of the sum of multiple sine waves. Finally, the particles were transported through a simulated drift space according to their new velocities. This procedure was repeated for a range of bunching voltages and focal lengths.

Using the final calculated time spread for the beam in each case, the final bunching efficiency was determined. Here, “bunching efficiency” is defined as the percentage of the original particles that reach the focal point within some acceptance time window. No par-

	$\Delta E = \pm 0.02\%$	$\Delta E = \pm 0.2\%$	$\Delta E = \pm 0.5\%$
$f_{rf} = 16.1$ MHz	Triangle	Triangle	Triangle, Lancosz, Gaussian
$f_{rf} = 80.5$ MHz	Lancosz	Gaussian, Lancosz, Triangle	Lancosz, Gaussian

Table 5.1: Window functions with the highest bunching efficiency by frequency and initial beam energy spread.

ticles are rejected for energies too far from the reference energy, but we do take note of the resulting energy spread for future reference. A plot of the bunching efficiency vs. focal length and peak voltage on the first bunching mode is shown in Figure 5.4. As the buncher is moved closer to the focal point, the bunching efficiency increases, but so does the required voltage. This suggests that a method for determining the optimal buncher placement is to select the maximum practical bunching voltage, and locate the buncher at the corresponding distance upstream from the focal point. In principle, the maximum acceptable energy spread would put a hard limit on the maximum voltage, but in this case the voltage becomes a limiting factor well before the energy acceptance measured in Chapter 4 is reached.

A number of different window functions (see 3.2.2) were tested at the same time. The results were found to be sensitive to both the energy spread of the initial beam and the bunching frequency. As the simulated beam energy spread increased, the differences between the various window functions decreased, as the larger overall longitudinal emittance of the beam became dominant over the small changes in bunching shape. For typical initial energy spreads for ReA, the best bunching efficiencies were observed using the Gaussian, Lancosz, and triangular window functions, with only small differences (see Table 5.1). In practice, these amplitudes would likely be used as starting points for the empirical tuning of a real buncher.

5.4 Beamline Reconfiguration

At this point, practical physical considerations come into play. The present project involves an upgrade to an existing facility, not simply determining the optimal focal length as part of a new design. It was determined that bunching voltages on the order of 1.5-2 kV would require amplifier power on the order of 300-500 W, which was seen as a reasonable upper limit. As such, the ideal location for the buncher was found to be somewhere between 1.7 and 2.5 meters upstream from the focal point at the entrance to the RFQ.

After consultation with the mechanical design staff at NSCL and the operators and engineers of ReA, a location was selected 1.96 m upstream of the RFQ entrance. In order to limit costs, rather than fabricating a new housing for the device, an existing diagnostic box (“L052”) was designated to receive the buncher. As originally installed, the box was located farther than 2 m upstream from the focal point. Additionally, the box was located after a relatively long 1 m drift space corresponding to the insertion box for the offline ion source, which could have made focusing the beam through the electrodes potentially problematic.

To put the buncher in the appropriate location, while reducing the distance the beam must travel without focusing, the L052 box was exchanged with the immediately following electrostatic quadrupole doublet. This change is shown in Figure 5.5. In this schematic drawing, the beam is traveling from left to right. The abbreviations in the schematic represent the following elements: “L052”, double width diagnostic box, “EQD”, electrostatic quadrupole doublet, “80”, 80.5 MHz multiharmonic buncher, “Box”, single width diagnostic box, “Sol”, room temperature solenoid, “16”, 16.1 MHz multiharmonic buncher.

This exchange of positions in principle solves multiple issues; the buncher can be installed at a position that allows for good focusing efficiency, the beam has to travel a shorter distance




Figure 5.5: Reconfiguration of the ReA LEBT section. The double width diagnostic box (L052) after the stable ion source has been moved downstream of the first quadrupole doublet, and divided into a single diagnostic box and the 16.1 MHz buncher.

without focusing, and focusing elements are placed immediately upstream of the 16.1 MHz buncher electrodes to allow for the beam to be focused at the bunching region.

Before this move was carried out, however, simulations needed to be conducted to verify that good beam focusing could be achieved, both through the electrodes of the 16.1 MHz buncher as well as through the even narrower aperture of the 80.5 MHz buncher. In the new configuration, only two quadrupoles were available between the desired focus points at the two bunchers, and so there was not a guarantee of a focal match between the center points. Using DYNAC, simulations were made of both the old beamline configuration and the proposed new one to determine if acceptable beam sizes could be achieved with the proposed reconfiguration.

Some results of these simulations are shown in in Figure 5.6. The x beam half width (1 RMS) is shown in red and the y width in green. The green bars indicate the positions of quadrupoles. The top graph is a simulation of the original beamline configuration, and the bottom shows the effects of moving one set of quadrupoles upstream and then retuning for the minimum beam size at the two bunchers (with a priority assigned to the second buncher, since it has the narrower aperture). Also visible in the figure is the move of the 80.5 MHz buncher 30 cm farther from the RFQ, which was carried out for unrelated reasons at the same time as the reconfiguration for the 16.1 MHz buncher.

The simulations showed that with the move of the quadrupole doublet upstream, it was




Figure 5.6: A before and after comparison of the simulated beam envelope for the reconfiguration of the ReA LEBT line. The top figure shows the simulated envelope for the original configuration, and the lower figure the fitted envelope after the quadrupole move. See text for further explanations.




Figure 5.7: Electrode design for the 16.1 MHz MHB

in principle possible to bring the beam to a tight focus at the 80.5 MHz buncher and still achieve a relatively good focus at the 16.1 MHz buncher. The simulation predicted a circular beam with a roughly 1 mm width at the 80.5 MHz buncher, and a somewhat more spread out beam with a 4 mm horizontal width and a 2 mm vertical width at the 16.1 MHz buncher. (All are widths given as 1 RMS full width values for the simulated particle distribution.)

5.5 Electrode Geometry

As discussed in Section 3.3, while an ideal beam buncher would have a perfectly uniform accelerating field, in practice this is not achievable because of the simple fact that there must be an opening for the beam to pass through. Choosing an electrode design with the profile of a pair of cones (Figure 5.7) allows for the majority of the accelerating field to be concentrated at the gap [43]. With this design, there are two main parameters of interest, the separation between the electrodes, and the width of the aperture.

Each of these parameters represents a tension between competing goals. A wide aperture radius makes focusing the beam through the aperture easier. The effect of a narrower aperture on the ease of beam tuning was seen in the beam transmission rates observed at




Figure 5.8: Cross section of the 16.1 MHz MHB electrodes modeled in CST STUDIO, showing the accelerating field in the $x = 0$ plane.

ReA after the original electrodes for the 80.5 MHz MHB were replaced with a narrower radius opening. On the other hand, the wider the aperture, the more the accelerating field will vary with radial distance, leading to undesired spreading effects.

The separation between the electrodes must also be balanced between two priorities. Arguing for a narrower separation is the Transit Time Factor described in Section 3.3.2. The TTF will fall as the electrodes are moved apart, requiring higher applied voltage for the same accelerating effect. This effect becomes more pronounced at higher frequencies. Balancing the desire for as narrow a separation as possible is the capacitance between the electrodes, which increases with decreasing separation and increases the needed power to achieve a desired accelerating voltage.

To select the optimum electrode geometry in this case, a three dimensional model of the electrode geometry was created in CST EM STUDIO (Figure 5.8). For a range of values of the aperture radius and electrode separation a three dimensional map of the static electric field was generated for a potential difference of one volt between the electrodes. These fields were then imported to the particle tracking code TRACK. Within TRACK the 3D field map could

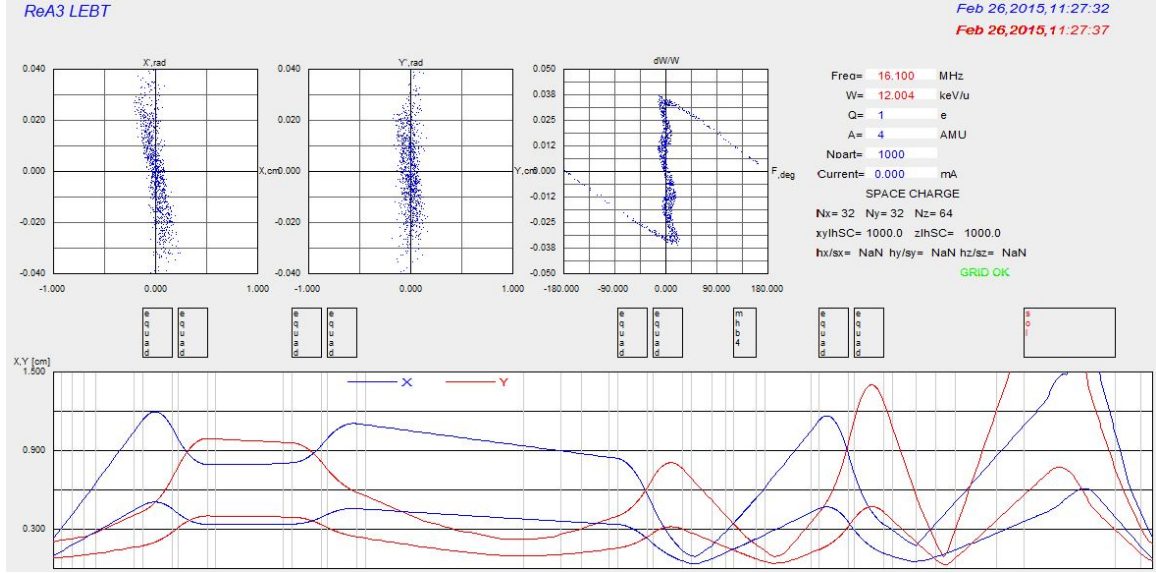


Figure 5.9: One simulation of the LEBT portion of the ReA beamline using a 3D field model for the 16 MHz buncher.

then be scaled up by the desired voltage applied to each mode on the buncher.

To determine the necessary per-mode voltage for each electrode geometry the following procedure was used: First, using the Monte Carlo simulation described in 5.3.2 combined with the buncher location determined in 5.4, the optimum bunching waveform was found for the ideal buncher case modeled with this simulation. (To reiterate, this ideal case does NOT consider transit time effects or radial field flatness.) Once this optimum waveform was found, the voltage required for each mode was recorded. Next, each mode with this ideal voltage was applied to the beam one at a time, and the resultant per mode energy spread of the beam after bunching was determined.

To determine the actual applied voltages needed for each 3D model in TRACK, this process was reversed. The beam was transported through the simulated field, and the voltage on each mode adjusted individually until the energy spread for that mode matched the final energy spread for that mode in simulation. Once the correct voltage for each mode was determined, all four modes were simulated using the calculated voltages, and correct focusing was verified.




Figure 5.10: Peak voltage for each mode for a range of electrode geometries.

Figure 5.9 shows a simulated passage through the buncher for a $Q/A = 1/4$ beam using TRACK. This procedure was repeated for a range of simulated electrode geometries.

The results, shown in Figure 5.10, indicate that while the voltage required for bunching is relatively insensitive to the gap between the electrodes, it is much more sensitive to the width of the aperture. However, even with a relatively wide aperture radius of 1 cm, the peak voltage required for bunching still resulted in a power requirement under the 500 W limit established as a reasonable value for the project. The final design values selected were 1 cm for the aperture radius and 6 mm for the gap between the electrodes.

It is worth noting that the gap for the 80.5 MHz MHB is only 4 mm, and could not be practically made wider. The reason for the narrower gap in this case is the much higher frequencies. The original design of the buncher called for a third bunching mode at a frequency of 321.5 MHz. The distance traveled by the beam ($\beta\lambda$) in one RF period for this mode is only 6.4 mm, giving an already low TTF of 0.45. By comparison, $\beta\lambda$ for the highest proposed frequency 64.4 MHz mode for the present buncher is 22 mm, meaning that for the case of the 6 mm gap, the TTF is a much higher 0.86.

Beam Energy [MeV/u]:		0.35	3	0.35	3
Freq.[MHz]	Rebunchers	dt [ns] ($\pm 3\sigma$)		dE/E ($\pm 3\sigma$)	
16.1	yes	6.8	0.6	0.5%	0.2%
	no	12.8	1.7	0.9%	0.3%
80.5	yes	3.2	0.5	0.2%	0.06%
	no	8.3	1.3	0.4%	0.2%

Table 5.2: Predicted minimum time and energy spread at target for beam energies from 0.3-3 MeV and 0.25-0.5 Q/A.

5.6 Effect on the Beam at the Detectors

Once the buncher location was determined, end-to-end simulations of the accelerator were run for a number of scenarios to estimate the effects on the final beam at the detector. The simulated beamline was tuned for the two most commonly requested experimental scenarios - minimum time spread on target, and minimum energy spread on target. Good transverse focusing is also required, and was included in the simulation, but is not discussed here.

The fundamental challenge in achieving a narrow time or energy focus at the detectors is that there are no longitudinally focusing elements in the ReA line after the final cryomodule of the linac, which is 30 m upstream from the target point. Depending on the beam energy, one or more cavities in the second or third cryomodule were used to rebunch the beam prior to transport, but especially at low energies, minimizing the longitudinal defocusing at the target proved extremely difficult.

As mentioned in Section 1.1.7.1, one possible future addition to ReA would be one or two rebunching cavities added just after the vertical “S” bend in the beamline. The addition of these cavities to the simulation resulted in a substantial improvement in the time and energy focus of the beam at the detector, especially at lower energies.

The results of these simulations are shown in Table 5.2. In each case, the time and energy spread is calculated as a $\pm 3\sigma$ deviation from the reference particle.

For the time spread, the results without rebunching cavities ranged from 6.8 ns for the lowest energy case, 350keV/u, to 0.6 ns for the 3 MeV/u case. The energy spread without rebunchers ranged from 0.9% for the low energy case to 0.3% for the 3MeV/u case. It is worth noting that the 350 keV/u scenario was not envisioned when the linac was designed, as it actually entails running the cavities at a *decelerating* phase to lower the beam energy from the fixed RFQ exit energy of 600 keV/u. At the more likely 3 MeV/u level, these energy spreads should be low enough to allow reasonable time-of-flight measurements to be made, in cases where the satellite bunches (see 5.7 below) are not an undue burden.

5.7 Satellite Bunches

5.7.1 Simulation of Satellite Bunches with No Cleaning

The same simulation was also used to estimate the number of particles which would remain in the undesired accelerator buckets and be transported all the way to the detector in these buckets. Simulating this question was not a simple matter using either TRACK or DYNAC. For internal reasons, both codes can only transport particles with phases $\pm 180^\circ$ away from the reference particle through a simulated RFQ. Particles outside of that phase range are either discarded, or shifted in phase modulo 360 until they are within one period.

This is acceptable when the simulated bunching frequency is the same as the RF frequency of the accelerator - any particles which are moved in phase can be assumed to be equivalent to particles from adjacent buckets. However, for the subharmonic case, a solution was developed using DYNACGUI to transport a particle distribution longer than a single RF period through an RFQ. The longer distribution was generated and transported to the entrance of the RFQ, where it was dumped to an external file and then divided into a number of distributions,




Figure 5.11: Simulation of the central bunch and satellite bunches at the detector position for the middle ReA beamline. This simulation was of a 2H^+ beam at 3MeV/u .

each one period in width. Those sub-distributions were transported through the simulated RFQ one at a time. The output distributions were rejoined into one large distribution, which was then transported through the remainder of the simulation.

For the case of the unmodified ReA3 with no bunch cleaning simulated, of the particles reaching the final detector position, roughly 96% were in the main bunch, leaving 4% in satellite bunches (Figure 5.11). This is not ideal, but should allow at least preliminary experiments to begin with higher current beams.

5.7.2 Simulation of Satellite Bunches with Cleaning

While the development of a comprehensive bunch cleaning solution was determined to be beyond the scope of this project, a few simulations were run to determine the effectiveness of several methods of cleaning satellite bunches. The most effective proved to be replacing




Figure 5.12: Normalized histogram showing particles in satellite bunches relative to the main bunch when cleaned with two 96 MHz cavities located at the position of the present CM1. Time axis is in ns.

the 80.5 Mhz accelerating cavity in the first cryomodule with two accelerating cavities at 96.6 MHz, which were able to reduce the relative sizes of the satellite bunches by a factor of 10^4 compared to the main bunch (See Figure 5.12). These off frequency cavities would only agree in phase with the RFQ and 16 MHz buncher at the zero phase of both. Any particles at the zero phase of the RFQ *not* in the main bunch would see a non-zero phase for the 96 MHz cavity and be deflected further in energy from the reference particle. These off-energy particles would then be dispersed in transverse position after bending elements and could be filtered using a slit.

5.8 Number of Buncher Modes

For the the present chapter, all calculations were made with the assumption that the buncher would be constructed to synthesize its waveform using four bunching modes. This assumption was made based on the four mode buncher in use at ANL, and as a reasonable choice given mitigation of some of the issues that were faced with the highest mode in the existing 80.5 MHz MHB by the lower overall frequency of this buncher. Some consideration was given to the potential effects if the buncher was eventually constructed using only three modes.

With fewer bunching modes, the linear portion of the bunching waveform would be shorter, resulting in a lower bunching efficiency. The estimated drop in bunching efficiency for 3 vs. 4 modes was approximately 4%. However, the lowered efficiency is balanced by lower power requirements - with three modes, the voltage for all modes is lower than in the four mode case.

Chapter 6

Fabrication and Installation

A project of this scope requires contributions from a number of specialties. The expertise required to design every element of this device, from the low level radio frequency (LLRF) module to the physical electrodes, is likely beyond the capabilities of any one person. Similarly, an exhaustive description of the design and operation of each subsystem is beyond the scope of this document.

This section will provide an overview of the design and construction of all the major subsystems of this project in a reasonable amount of detail, but will not attempt to encompass all of the design choices in the same level of specificity as Chapter 5. The broad areas addressed are: the electrodes and electrode mounting assembly, the RF tuners and resonators, and the LLRF and amplifier electronics.

6.1 Mechanical Design

Once the physical dimensions of the buncher electrodes were determined, (as outlined in 5.5) these specifications were transmitted to the NSCL / FRIB mechanical design department. In addition to providing the appropriate electrode geometry as described above, a number of other requirements also needed to be met. The final design of the electrode assembly is shown in Figure 6.1.

First, after installation the centerline of the installed electrodes must precisely match the




Figure 6.1: The completed electrode assembly, showing the adjustment screws for the x and y axes and the alignment monuments. The center plate can also rotate to ensure the centerline of the electrodes is parallel to the beam axis.

reference axis of the beamline to avoid steering or focusing effects on the beam. The two possible approaches to this constraint were to either precisely design the electrode mount to exactly match the box into which it would be mounted, or else to make the electrode mounting adjustable. The second option was determined to be more desirable, both because the exact dimensions of the L052 box were not precisely known, but more importantly to provide flexibility should the electrodes be moved to another location in the future.

Three axes of adjustment were incorporated in the design - horizontal adjustment along the x axis, vertical adjustment along the y axis, and rotational adjustment around the center of the mounting flange. No adjustment was incorporated in the z axis to reduce the expense and complexity of the design, after simulation showed minimal effect on the beam parameters from small changes (< 1 cm) to the z position of the electrodes. To allow for adjustment of electrode position relative to the beamline after installation into the box, the electrodes were rigidly attached to a 6" vacuum conflat which was then connected via a bellows to the 8" flange on the outside of the L052 box. Small changes to electrode position will therefore be absorbed into the shape of the bellows.

The 6" conflat was mounted to a "collar" which mates to a second collar supported by threaded rods connected to the 8" flange on the box. The height of the collar connected to the 8" flange is adjustable using nuts on either side of this collar. These nuts provide the horizontal (x axis) adjustment for the electrodes. The two collars can slide against one another, allowing their precise relative position to be adjusted by rounded screws and then locked in place with three screws on the sides. This mechanism provides the vertical (y axis) adjustment. Finally, the 6" conflat is itself rotatable within the mounting flange, which allows the rotational orientation of the electrodes to be adjusted. To facilitate alignment once the assembly is mounted in the beamline, a set of three monuments were installed on

the outside of the 6" conflat.

Secondly, the electrodes needed to be held at a precise relative alignment while being electrically separated. This issue was resolved by mounting the electrodes on a precisely machined ceramic block. As the NSCL machine shop does not have the tools necessary for such machining, this block was the only customized component of the design which had to be fabricated off site.

The third major requirement was to provide a path to transfer the electrical power from the exterior of the beamline to the electrodes. To this end, a pair of ceramic feedthroughs were mounted in 1" flanges on the outer 6" conflat. The connection from the inside of the feedthroughs to the copper electrodes included a length of flexible copper strap, allowing for a small amount of mechanical tolerance during assembly.

6.2 RF Design

6.2.1 Resonators and Tuners

To apply the RF power to the electrodes, some form of resonator is required. The tank circuit design used at the ANL buncher and described in Section 3.5.2 was not selected in this case. The existing 80.5 MHz MHB uses rigid coaxial structures as resonators, but this is not feasible at the lower wavelengths of the present device. To first order, the required resonator length is simply given by:

$$L = \frac{\lambda}{4} = \frac{c}{4f} \tag{6.1}$$

(The factor of 4 is due to the fact that this is a quarter wave resonator.) As such, for an 80.5

MHz mode, the resonator must be approximately 93 cm long. By contrast, for a 16.1 MHz mode, the required length for a $\lambda/4$ resonator is approximately 4.66 m, which is impractical for a rigid resonating structure. Instead, two lengths of 1" diameter air-core coaxial cable with a dielectric insert were used as resonators. The longer cable acts as a quarter wave resonator for the lowest, 16.1 MHz mode, and the shorter one acts in quarter wave mode for the second, 32.2 MHz mode. In addition, each resonator can be excited in the $\frac{3\lambda}{4}$ mode to provide the 48.3 MHz and 64.4 MHz modes respectively.

The chosen cables have a dielectric material (HDPE) separating the inner and outer conductors, which alters the capacitance of the cable. This lowers the propagation speed of the electromagnetic wave in the cable by 7%, which in turn shortens the required resonator length by the same amount.

The actual length of cable used is still shorter, however, because part of the resonator length is provided by a rigid copper structure with a sliding short which is used for precise tuning of the resonator frequency. This rigid structure is 80 cm long (with 70 cm of usable length) for the longer resonator and 70 cm (with 50 cm usable) for the shorter one. All of these lengths are summarized in Table 6.1. The final cable length is not simply the HDPE $\lambda/4$ length minus the tuner length because there is no dielectric material in the tuner itself.

To facilitate ease of adjustment, the tuners were located on the roof of the ReA3 clean-room enclosure with the resonating cables passing through a sealed hole in the roof to connect to the feedthroughs for the electrodes. This allows the tuners to be adjusted without the full gowning required to enter the clean room itself. Additionally, since approaching the buncher feedtroughs requires turning off the beam to the cleanroom, placing the tuners on the roof allows for adjustment of the bunching frequency without disabling the beam.

Frequency [MHz]	$\frac{\lambda}{4}$ (vacuum) [m]	$\frac{\lambda}{4}$ (w/HDPE) [m]	Tuner Length [m]	Cable Length [m]
16.1	4.66	4.33	0.70	3.68
32.2	2.33	2.17	0.50	1.70

Table 6.1: Lengths of the resonators and tuners for the 16.1 MHz buncher.

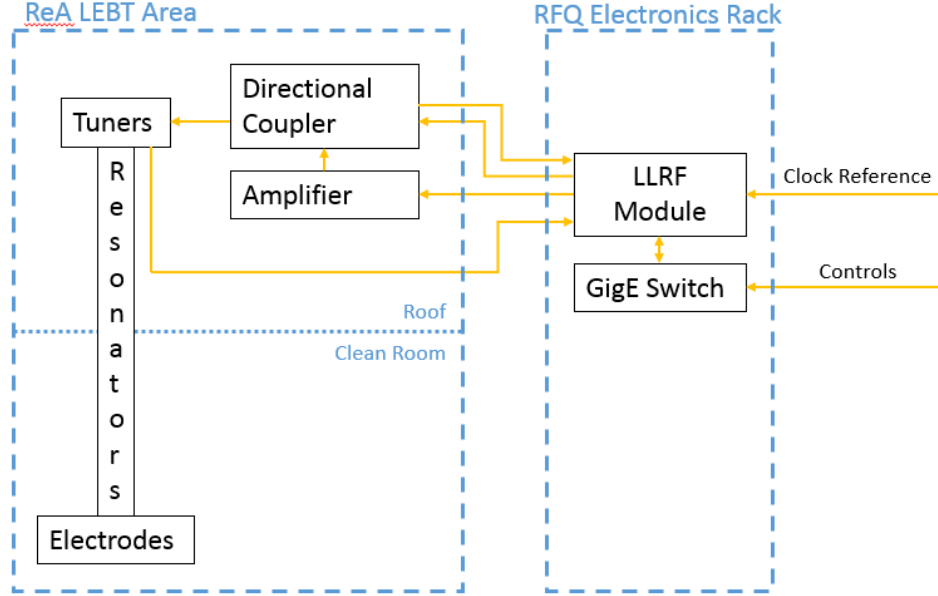


Figure 6.2: A schematic drawing of the connections for the RF system for the 16.1 MHz buncher.

6.2.2 Low Level RF

The low level radio frequency (LLRF) module used to drive the system requires a number of interconnections, shown in Figure 6.2. The LLRF module receives a timing signal from the master clock circuit for ReA. It then sends a drive signal to the amplifier which is fed to the tuner/resonator assembly via a directional coupler. The return signal from the directional coupler is used by the LLRF module to actively adjust the quality of the transmitted wave.

The LLRF module uses active cancellation techniques to produce a pure signal at each mode using only one amplifier and control module when operated in closed loop mode. This module is the first time these active cancellation techniques have been used in the ReAccelerator project. The operation of the LLRF module is described in more detail in

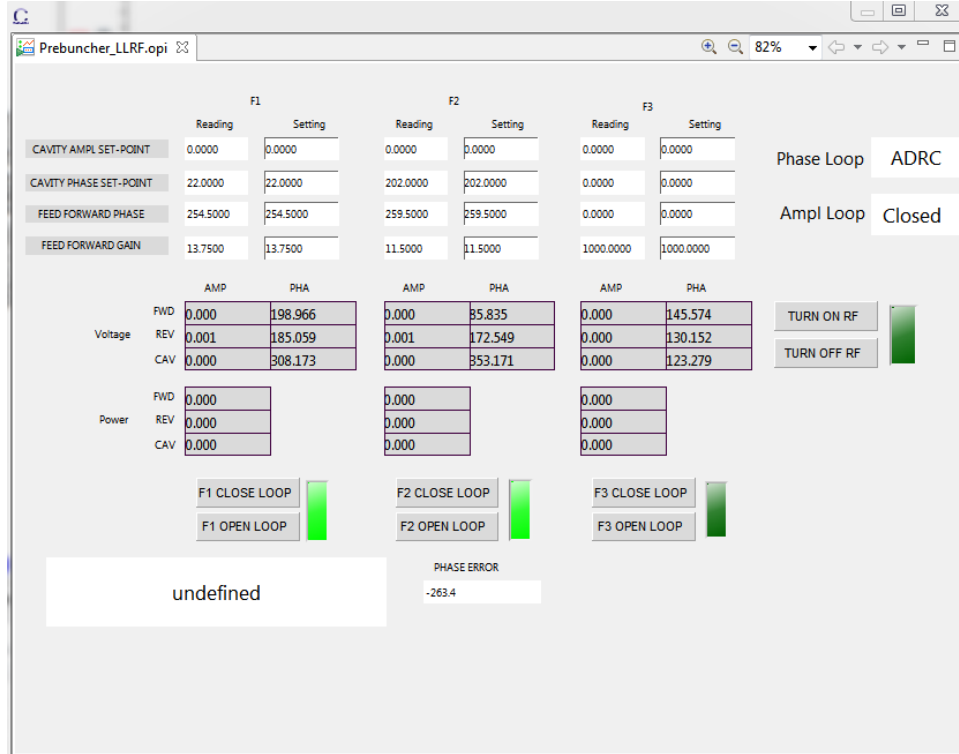


Figure 6.3: The control screen for the LLRF module, showing the controls for each of the three modes.

Section 7.4.

The module is controlled via Ethernet, and is connected to the laboratory control system through a gigabit Ethernet switch located in the same rack with the LLRF module. During normal operations, all control of the system is via the LLRF module. A Control System Studio (CSS) [44] application can be used to set the amplitude and power for each mode, and monitor the forward and reflected power for each. The operating screen for this application is shown in Figure 6.3.

6.3 Electrode Installation

As a first step, a “dirty” (i.e., not under clean room conditions) assembly of the electrodes was performed to verify the presence of all necessary parts, correct fits, etc. Several small

changes were required at this point, including shortening several screws, adding spacer nuts to the barrel connectors for the feedthroughs, etc. Once the dirty assembly was complete, the electrodes were disassembled and the parts cleaned in preparation for the final assembly.

Final assembly was made in clean room conditions, and the electrodes were fiducialized by the NSCL / FRIB alignment group with respect to the monuments on the outside of the 6" conflat to allow for alignment of the electrodes from outside the vacuum chamber after installation. The electrodes as installed in the beamline are shown in Figure 6.4. In this image, the resonators can be seen passing through the hole in the clean room roof.

The final connections from the resonators to the feedthroughs were made with copper straps. The entire assembly was then shielded with a protective cover, as the connecting straps support a high voltage connection during device operation.




Figure 6.4: The electrode assembly installed in the beamline, with the resonators visible .

Chapter 7

Commissioning: Methodology and Results

Testing the output of the 16.1 MHz multiharmonic buncher posed very specific challenges. Only one type of diagnostic available at ReA, the timing wire, operates in the time domain on the scale required for measuring the amount of bunching produced by the device. An indirect bunching measurement was also possible by measuring the transmitted beam current after the RFQ and linac, since less of the unbunched beam is accelerated to the full velocity of the reference particle.

This chapter outlines the diagnostic equipment, testing methodology, and results from the commissioning of the buncher.

7.1 Timing Wire Detectors

The primary diagnostic device used for measuring the time distribution of the beam was a timing wire detector. The design of the timing wires used at ReA is comparable to the device described by Verzilov [45] for use at the heavy ion Linac at TRIUMF. The detector consists of a thin wire stretched along the axis of a metal cylinder, coupled to a multi-channel plate (MCP). It is shown in Figure 7.1.

While the cylinder is held at ground potential, the wire is charged to a potential on the




Figure 7.1: Two views of the timing wire detector. The MCP is indicated by number 16, and the wire itself by 18. In the right hand view, the direction of beam travel is into the page, striking the wire before exiting on the far side of the “can.”

order of a few kilovolts, thus creating a radial electric field. When the beam strikes the wire, secondary electrons emitted from the wire are accelerated by the radial field. A hole in the cylinder permits some of these secondary electrons to pass through and strike a Hamamatsu micro-channel plate detector (model F4655-12).

The data acquisition path is shown in Figure 7.2. The center of the data acquisition is an Ortec 566 Time-to-Amplitude Converter (TAC). This device “listens” for signals on its start and stop terminals, and generates an output pulse with amplitude proportional to the time delay between those signals. The Start input is connected via a constant fraction discriminator to the output of the timing wire. Each detected secondary electron pulse from the timing wire thus triggers the start input.

The stop input requires a constant frequency signal synchronized with the bunching signal to be analyzed. For closed loop measurements, where the buncher is synchronized to

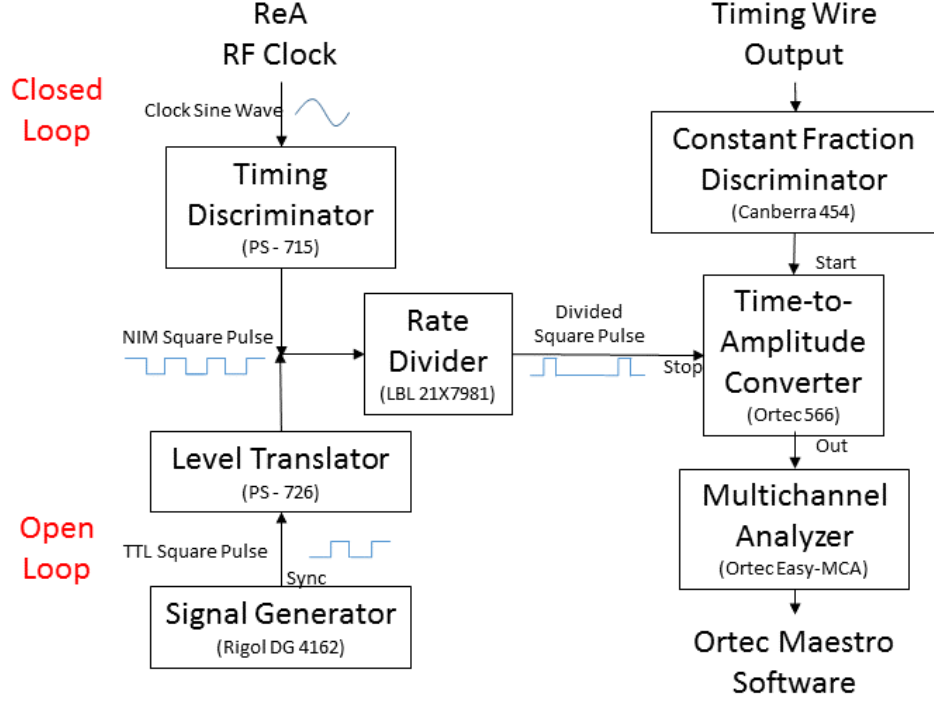


Figure 7.2: The data acquisition path for the timing wire measurements.

the ReA master clock signal, this clock signal is passed through a timing discriminator to transform the sinusoidal reference to a square pulse. For open loop measurements where the buncher is driven by a signal generator, the stop signal is taken from the sync output of the signal generator driving the buncher, and then passed through a level translator to provide the correct voltage. Since this gives a maximum start to stop delay of one RF period in either case, the signal is then passed through a rate divider [46]. This lowers the frequency of the stop signal by an integer divisor, and allows for multiple pulses to be visible at once in the data. Finally, the output from the TAC is passed to a Multi Channel Analyzer (MCA) for display and analysis on a laptop computer using the MAESTRO software from Ortec [47].

The recorded spectrum consists of a histogram showing total counts for each possible time separation between the start and stop signals detected at the TAC. This histogram is then extracted for analysis in MATLAB or another code. An example histogram is shown in




Figure 7.3: A example of a histogram from the TAC. This example is for beam bunched with the second (161.0 MHz) mode of the 80.5 MHz MHB at a control system setting of 11.

Figure 7.3.

It is important to note that in these histograms, time is increasing to the *left*. The reason for this is that the stop signal for the TAC is coming from a fixed clock, while the start signal comes at a variable time. Earlier signals have higher time separation from the reference clock than later ones. Since a higher time separation corresponds to a higher amplitude pulse from the TAC, such pulses are graphed farther to the right. Conversely, later signals have shorter separations and lower TAC pulse amplitudes, and are thus plotted at left. So the progression from earlier to later pulses in the histograms is from right to left.

7.2 Timing Wire Testing

Prior to testing the 16.1 MHz buncher, the timing wire detector located just before the RFQ (designated “L061”) was tested using beam bunched by the existing 80.5 MHz buncher. For this test, a beam of $^{14}\text{N}^{5+}$ was supplied by the EBIT. The EBIT was used to ensure

a smaller transverse beam at the locations of the buncher and the timing wire than could be achieved with the test source used for the measurements described in Section 4.1. Since the beam from the test source is bent at a ninety degree angle just before the 16.1 MHz buncher, its transverse extent is relatively large at that point. The beam from the EBIT travels in a straight line for four meters before the buncher, making it easier to achieve a smaller transverse size at both bunchers and at the timing wire.

For these measurements, the beam energy was 12.9 keV/u. This is not the reference injection energy for ReA, as the EBIT was kept at the same voltage level from a prior experiment, which used a beam with a slightly different mass to charge ratio. With the EBIT held to the higher voltage, the bending dipole in the Q/A separator was adjusted to select for a $^{14}\text{N}^{5+}$ beam.

For all measurements using the L061 timing wire, the room temperature solenoid located just upstream of the detector was turned off to avoid interference between the solenoid fringe field and the paths of electrons from the wire to the MCP of the detector. However, without the additional focusing provided by the solenoid, the beamline must be retuned to bring the beam to as tight a focus as possible on the wire without the use of the solenoid. Once the beam was tuned to provide adequate current on the timing wire, the amplitude of each mode of the 80.5 MHz buncher was individually ramped, and the timing wire data for each power setting was recorded.

The primary purpose of these measurements was to test the data acquisition system in preparation for the measurements of the 16.1 MHz buncher. Fortuitously, it is also possible to gain additional information about the beam and the existing buncher by comparing the overall results of these measurements with the predictions of theory. In particular, based solely on these measurements, it is possible to estimate both the energy spread of the incipient

beam, and the applied voltage at the buncher. Since the control system for the 80.5 MHz buncher was not well calibrated, this information is potentially useful. At the time of these measurements, the amplitude setting for the two modes of the bunchers could be set from 0 to 80 in the control system. However, while the units on the display are marked as “MV/m”, these numbers were essentially uncalibrated, although they were adjusted to scale linearly with electrode voltage, rather than amplifier power.

To analyze the data, each histogram was fitted using the PEAKFIT.M curve fitting module for MATLAB developed by O’Haver [48]. For each control system setting, the mean full width at half maximum (FWHM) of the fitted curves and the mean separation between the peaks was tabulated. Since the frequency of the buncher is known, the separation between the peaks is just a single period of the mode under measurement. This allows the width of the histogram bins to be calibrated:

$$\tau_{bin} = \frac{10^3}{f_{mode} * n_{bins}} \quad (7.1)$$

where τ_{bin} is the length of the bin in ns, f_{mode} is the RF frequency in MHz, and n_{bins} is the mean number of bins between peaks. Given this calibration, the FWHM can also be converted to an equivalent width in ns. Figure 7.4 shows the FWHM plotted against the control system setting for each mode. The approximately parabolic shape represents the transition from an underfocused beam to optimum focusing, and then to overfocusing.

In order to compare these measurements with theory, a series of simulated bunched beam particle distributions were generated in MATLAB using the procedure outlined in 5.3.2. For each mode, the mean FWHM of the bunched peaks was calculated for a range of applied voltages and initial beam energy spreads. The result of these simulations for the first mode of




Figure 7.4: The mean width of the bunched beam at a range of control system settings using the 80.5 MHz buncher. The 80.5 MHz mode is shown at left and the 161.0 MHz mode at right.

the 80.5 MHz buncher is shown in Figure 7.5. The minimum possible mean width is related to the energy spread by simple conservation of emittance; the greater the initial longitudinal emittance, the less it is possible to compress the beam by rotation in phase space. In the figure, this is apparent in the increasing height of the flat portion of the FWHM curves with initial energy - the greater the initial energy spread, the higher this “plateau.”

Comparing these simulated curves to the data from Figure 7.4 allows for an immediate estimate of the initial energy spread of the beam - the FWHM of the “plateau” in the measured curve corresponds to a simulated initial energy spread of $\pm 0.1\%$. Adjusting the scaling of x-axis until the measured data aligns with simulation also allows for an estimated correspondence between the voltage seen by the beam and the control system setting.

Figure 7.6 shows the results of combining the simulated width measurements with the measured data for the first and second mode of the 80.5 MHz buncher. The measured data is multiplied by a scaling factor to give the estimated on-axis buncher voltage that corresponds to the observed bunch widths. The estimated scaling factor for the first mode is




Figure 7.5: Comparison of the FWHM curve for simulated bunched beams with a range of initial energy spreads for the first mode of the 80.5 MHz multiharmonic buncher.




Figure 7.6: Measured beam FWHM from Figure 7.4 overlaid on simulated bunch widths. The X-axis of the measured data has been scaled to match the simulated beam voltage.

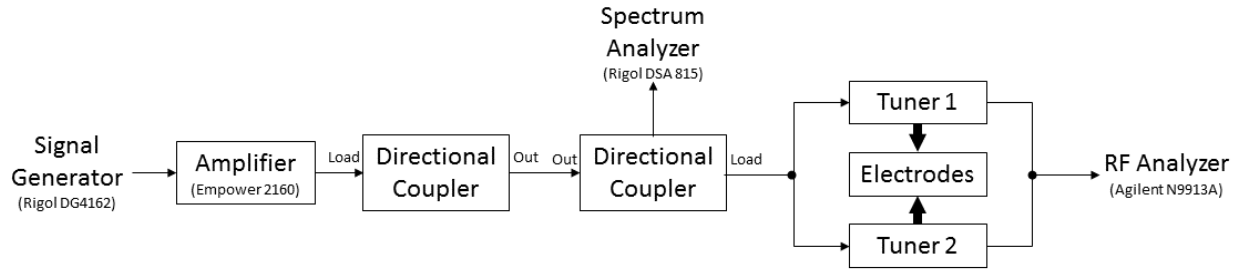


Figure 7.7: The drive power setup for open loop testing of the 16.1 MHz buncher.

approximately 18, and for the second mode is approximately 10. In both cases, the minimum observed FWHM corresponds to a simulated beam with initial energy spread of $\pm 0.1\%$.

7.3 Open Loop Bunching

7.3.1 Test Setup

For the first tests of the 16.1 MHz buncher, the device was run in “open loop” mode (defined below). The drive power configuration for these tests is shown in Figure 7.7. A sinusoidal drive signal was generated by a Rigol DG4162 signal generator at the frequency to be tested, and fed to a 100 watt Empower 2160 amplifier. (The power of this amplifier was predicted to be insufficient for best bunching of the highest rigidity beams, but adequate for testing with the $^{14}\text{N}^{5+}$ beam used here.) The amplified signal was then fed through two three-port couplers as seen in Figure 7.7. (The incident terminal of the first coupler was not used for this test.) The power from the second coupler was fed to the electrodes via the tuners and resonators. The forward power was monitored via a Rigol DSA 815 Spectrum Analyzer connected to the incident terminal of the second coupler. Cavity power was monitored via an Agilent N9913A RF Analyzer connected to the tuner pickups.

Open loop in this case refers to the fact that there is no feedback between the signal generator and the amplifier or the resonators. Without such feedback, adjusting the phase or amplitude of any of the three modes can have an impact on the other phases and amplitudes. As such, any attempt to tune all three modes together in this mode is difficult and unlikely to produce consistently reproducible results. Open loop mode *is* suitable for studying the effects of the modes individually. As described below, a brief attempt was made to combine modes one and two as a proof of concept.

The same beam and tuning (with minor retouching) were used for these tests as for the 80.5 MHz timing wire tests described in Section 7.2: a beam of $^{14}\text{N}^{5+}$ with energy 12.9 keV/u was generated by the EBIT and transported to the L061 timing wire with the L060 solenoid turned off. The data acquisition for this test was via the “Open Loop” path shown in Figure 7.2. The sync output of the signal generator was used for the stop input of the TAC, via a level translator and a rate divider. The rate divider and TAC were set to sample a maximum length of 200 ns, meaning that if a stop signal did not arrive within 200 ns of the signal from the timing wire, no output pulse would be sent by the TAC. This limited the maximum number of peaks visible for the first mode to at most four.

7.3.2 Results

The raw data for the single mode scans for the first three modes of the 16.1 MHz buncher are shown in Tables 7.1, 7.2, and 7.3. From left to right, each table shows the displayed amplitude of the input signal on the signal generator in dBm, the forward power read from the spectrum analyzer in both dBm and watts, the cavity power read from the RF analyzer in dBm, the mean fitted widths of the peaks in ns, and the estimated on-axis peak voltage for that input power in volts. The fitted FWHM were calculated with MATLAB using the

same procedure outlined in Section 7.2 above.

Since the same beam was used for this test as for the prior tests, the energy spread was assumed to be the same as the spread fitted to the 80.5 MHz multiharmonic buncher data, $\pm 0.1\%$. A simulated peak width vs. buncher voltage curve was generated for each mode using this energy spread, this time using the frequency and focal length for the 16.1 MHz buncher. The peak voltage for each mode was then estimated by overlaying the measured FWHM curve over the simulated data. In this case, since the measured input value was power, the scaling factor was relative to the square root of the forward power in watts (since power scales as voltage squared). On axis peak voltage is thus given by the equation:

$$V_p \approx s_p \sqrt{P[W]} \quad (7.2)$$

The values of the scaling factor s_p were found to be 195, 70, and 50 for modes 1, 2 and 3 respectively. These results are shown in Figures 7.8, 7.9, and 7.10.

In addition to the single mode results tabulated here, an attempt was made to combine the first two modes to produce improved bunching over the first mode individually. The signal generator was set to “harmonic” mode, which allowed it to produce a synthesized waveform combining sine waves of varying phase and amplitude. However, although the relative phases and amplitudes at the signal generator amplitude were well controlled, without feedback, the final synthesized waveform at the electrodes was less predictable due to the excitation of higher harmonics in the resonators. In particular, shifting the phase or amplitude of either mode could result in unpredictable shifts in the other parameters. This was verified by viewing the pickup signal from the tuners on an oscilloscope.

Further complicating the attempt to produce a two-mode bunching signal on the beam

SG Amplitude dBm	Forward Power		Cavity Power dBm	Mean FWHM ns	Est. V_p V
	dBm	W			
5	-12.2	4.265795188	-6.8	3.206705	103.2690078
6	-11.3	5.248074602	-5.7	2.297505	114.5433826
7	-10.3	6.60693448	-4.7	2.075214	128.5197891
8	-9.3	8.317637711	-3.7	1.609026	144.2015752
9	-8.25	10.59253725	-2.65	1.507151	162.7308917
10	-7.2	13.48962883	-1.6	1.304313	183.6411502
10	-7.44	12.76438809	-2	1.426882	178.6364191
10.4	-7	14.12537545	-1.53	1.364091	187.9187021
11	-6	17.7827941	-0.4	1.231745	210.8482517
11	-6.4	16.21810097	-1	1.301527	201.3585172
11.9	-5.4	20.41737945	0	1.209429	225.9279722
12	-4.9	22.90867653	0.66	1.235012	239.3150462
12.2	-5	22.38721139	0.48	1.199475	236.5756295
12.6	-4.45	25.40972706	1	1.243016	252.0403096
13	-4	28.18382931	1.47	1.509026	265.4422221
14	-3.25	33.49654392	2.2	1.593702	289.3809942
15	-2.5	39.81071706	2.96	2.46934	315.4786722
16	-1.74	47.42419853	3.65	3.183149	344.3261482

Table 7.1: Open loop testing results for the first (16.1 MHz) mode of the buncher.




Figure 7.8: Measured FWHM for the first (16.1 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match.

SG Amplitude dBm	Forward Power dBm W	Cavity Power dBm	Mean FWHM ns	Est. V_p V
3	-11.75 9.66050879	-3.66	3.738695	147.8
4	-10.55 12.73503081	-2.45	2.560171	101.21
5	-9.6 15.84893192	-1.53	1.864187	73.696
6	-8.95 18.40772001	-0.88	1.231645	48.69
6.1	-9.1 17.7827941	-1.05	1.502713	59.406
6.7	-8.6 19.95262315	-0.51	1.246518	49.278
7	-8.35 21.1348904	-0.28	1.181762	46.718
7.5	-8.07 22.54239212	0.01	1.229494	48.605
8	-7.7 24.54708916	0.36	1.150648	45.488
8.4	-7.47 25.88212915	0.61	1.196484	47.3
8.7	-7.27 27.10191632	0.78	1.212799	47.945
9	-7.07 28.37919028	0.99	1.222816	48.341
9.6	-6.77 30.40885026	1.29	1.358097	53.689
10	-6.52 32.21068791	1.54	1.421362	56.19
11	-5.63 39.53666201	2.4	2.129286	84.176
12	-4.52 51.0505	3.49	3.322329	131.34
13	-3.9 58.88436554	4.1	4.003288	158.26

Table 7.2: Open loop testing results for the second (32.2 MHz) mode of the buncher.




Figure 7.9: Measured FWHM for the second (32.2 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match.

SG Amplitude dBm	Forward Power dBm W	Cavity Power dBm	Mean FWHM ns	Est. V_p V
5	-18.65 6.99841996	-8	4.130606	7.999
6	-17.77 8.570378452	-7.1	2.257019	6.052
6.7	-17.4 9.332543008	-6.8	1.842162	5.598
7	-17 10.23292992	-6.37	1.392666	5.031
7	-17 10.23292992	-6.48	1.599765	4.542
7.6	-16.6 11.22018454	-6	1.370627	4.204
8	-16.15 12.44514612	-5.5	1.232451	3.775
8	-16.3 12.02264435	-5.7	1.313748	3.888
8.6	-15.9 13.18256739	-5.28	1.226708	4.188
9	-15.4 14.79108388	-4.77	2.068946	4.987
9	-15.63 14.02813705	-5	1.338013	3.731
10	-14.2 19.498446	-3.55	3.067402	8.695
11	-13.65 22.1309471	-3	4.201705	6.65
12	-12.89 26.36331386	-2.25	5.00303	6.889

Table 7.3: Open loop testing results for the third (48.3 MHz) mode of the buncher.




Figure 7.10: Measured FWHM for the third (48.3 MHz) mode of the buncher overlaid on a simulated curve, with the x -axis scaled to match.




Figure 7.11: Fitted peaks for manually tuned open loop bunching using modes 1 and 2.

were delays inherent in the signal generator, which required several seconds to readjust each time a waveform was changed, and the data acquisition system, which required 30-60 seconds to display the results of each new setting. Given all of these issues, no systematic attempt was made to explore the parameter space, but rather an empirical tuning was attempted as proof-of-concept.

The results of this attempt are shown in Figure 7.11. While the resultant peaks were not particularly symmetrical, making the fit somewhat suspect, the fitting routine estimated a mean FWHM for the peaks of 2.47 ns. This is of the same order of magnitude as the bunching achieved with the first mode alone, and indicates that the second mode can at least in principle be combined at the appropriate relative phase with the first mode.

7.4 Closed Loop Bunching

7.4.1 Setup

The final operational mode of the 16.1 MHz buncher is “closed loop.” As opposed to open loop bunching, closed loop bunching uses the feedback from the resonators and the directional couplers to analyze the amount of power in the undesired upper harmonics for each

bunching mode. It then actively adjusts the transmitted power at those frequencies to cancel the unwanted harmonics. This requires more power than open loop mode, because the cancellation is actively driven by the same amplifier. It has the advantage of providing a very stable signal with precisely controllable amplitude and phase for each mode. This active control is provided by a low level RF (LLRF) module designed and built by the NSCL / FRIB RF department.

For the closed loop measurements, a 12.0 keV/u beam of $^{40}\text{Ar}^{13+}$ was used. This was selected in place of the previously used nitrogen beam on the theory that since the accelerator had been tuned for this beam for the immediately prior experiment, it would therefore require minimal retuning. This assumption proved somewhat optimistic, as the lower beam current available meant that more careful tuning was needed to place sufficient beam onto the relatively small area of the timing wire. Two timing wires were used, the wire at “L061” described above, located immediately before the RFQ, and the “L110” timing wire, located at the end of the linac.

Unlike the open loop measurements, which were directly controlled from a signal generator located near the device, once the LLRF module was in place the closed loop measurements could be controlled and monitored from the ReA control room area using the Control System Studio application described in Chapter 6.

The other major difference in the test setup for the closed loop measurements was the source of the synchronization signal for the TAC connected to the timing wire detector. In the case of the open loop tests, the 16.1 MHz synchronization signal was taken directly from the signal generator driving the buncher. In the closed loop case, the clock signal driving the LLRF module was an 80.5 MHz feed from the ReA master clock, divided internally to 16.1 MHz within the LLRF module. Unfortunately, the LLRF module has no output for

this divided signal. Additionally, no rate divider was available which was specified to work at the 80.5 MHz clock frequency.

Since a synchronized, lower frequency signal is a requirement of this data acquisition method, a synchronization signal was instead taken from the 10 MHz output of the ReA clock module. This output was originally intended for RF test equipment, but seems never to have been used until this experiment. This signal was used to phase lock a 16.1 MHz signal from a Rhode & Schwartz SML 03 Signal Generator, which was then fed to the same rate divider as in the open loop case.

For the first set of measurements, a capacitor was added between the longer resonator and the feedthrough in an attempt to help stabilize the frequency of the third mode. While initial measurements showed that the third mode was stable in this configuration, it had the drawback of increasing the required power for the first mode to the point that the beam could not be brought to a longitudinal focus at the detector. As a result, for the final testing, this capacitor was removed, and the buncher was operated with two modes only.

Finally, the 100 watt amplifier which was used during the open loop testing was replaced with a 300 watt (E & I model A-300) amplifier to provide additional headroom and long term support for higher rigidity beams.

7.4.2 Single Mode Results

Since limited time was available for testing, fewer data points were taken for the closed loop measurements than for the open loop case. For the first and second mode individually, a range of voltages were set in the control system, and the power consumption recorded and a histogram captured from the timing wire. At a later date, the histograms were analyzed to determine the average FWHM per peak at each bunching power.

CSS Setting	Forward Power	Reflected Power	Mean FWHM	Est. V_p
V	W	W	ns $\pm 1\sigma$	V
100	1.4	0.4	20.8 ± 0.7	170
150	3	0.7	17.0 ± 1.2	255
200	5.5	1.4	14.6 ± 0.6	340
250	8.3	2	9.6 ± 0.3	425
300	12.1	3	3.2 ± 0.1	510
350	16.5	4.1	1.9 ± 0.1	595
400	21.4	5.3	2.4 ± 0.04	680

Table 7.4: Closed loop testing results for the first (16.1 MHz) mode of the buncher.

CSS Setting	Forward Power	Reflected Power	Mean FWHM	Est. V_p
V	W	W	ns $\pm 1\sigma$	V
50	0.18	-	11.9 ± 3.7	60
100	0.75	-	10.5 ± 1.3	120
150	1.7	-	9.1 ± 1.0	180
175	2.3	-	3.9 ± 0.9	210
200	3	-	3.7 ± 0.9	240
225	3.8	-	2.7 ± 0.6	270
250	4.7	-	2.2 ± 0.4	300
300	6.8	-	2.0 ± 0.6	360

Table 7.5: Closed loop testing results for the second (32.2 MHz) mode of the buncher. No reflected power was observed at any of the chosen voltage settings.

The raw data is shown in Tables 7.4 and 7.5. No reflected power was observed during the mode 2 observations, indicating the cavity resonant frequency was tuned quite accurately to the desired 32.2 MHz resonant frequency. The errors in the fitted FWHM are derived from the “bootstrap” method used in the peak fitting routine [49], and are reported at the one sigma level.

It can be seen from Figures 7.12 and 7.13 that the single mode data taken in closed loop mode is not as comprehensive as that in open loop mode. This is unfortunate, but is a result of two main factors. The first was simply time pressure to finish the measurements in the time allotted. The other was that at the high end of the voltage range that was scanned for each




Figure 7.12: Measured FWHM for the first (16.1 MHz) mode of the buncher in closed loop mode, overlaid on a set of simulated curves, with the x -axis scaled to match.




Figure 7.13: Measured FWHM for the second (32.2 MHz) mode of the buncher in closed loop mode, overlaid on a set of simulated curves, with the x -axis scaled to match.




Figure 7.14: An illustration of an overfocused beam, showing the expected double peaked structure of the histogram.

mode, the peak in the timing wire histogram was seen to bifurcate. A bifurcation at higher voltages is an expected consequence of the mechanics of bunching. While an underfocused beam appears as a wide peak, an overfocused beam will appear as two separate peaks, as the most accelerated and retarded portions of the bunched beam are overlaid on the tails of the longitudinal distribution. This is illustrated in Figure 7.14.

Unfortunately, after analysis of the sizes of the peaks, it would appear that the bifurcation in this case was likely *not* the result of overfocusing, but rather the result of a second, narrow peak which was hidden by noise at the lower bunching energies. The most likely explanation for this peak is that some small fraction of the beam was hitting the “can” of the timing wire detector, resulting in a time signal slightly offset from the desired signal from the wire itself. The fits shown in the tables and figures for closed loop bunching are made to the larger of the two peaks, and continue to match the simulated shape of the bunching curve reasonably well. This would also tend to support the theory that the second peak was not the result of overfocusing.

On a more positive note, however, the data that was collected still indicates a reasonably good fit for the simulated curves shown, and the two modes give consistent results for the initial beam energy spread, in this case about $\pm 0.15\%$. Since the control system was

calibrated for voltage, not power, the relationship between the control system setting (V_{cs}) and the actual peak centerline voltage (V_p) is the simple linear relation:

$$V_p \approx s_v V_{cs} \tag{7.3}$$

The constants s_v relating the control system setting to the actual on axis beam voltage are estimated at 1.7 for mode 1 and 1.2 for mode 2. Since the maximum amplifier power used for either mode was less than 30 watts, there would seem to be plenty of leeway available to reach focusing, even at the higher requirements imposed by multi-mode bunching and higher rigidity beams.

7.4.3 Multiple Mode Results

The most critical result for the performance of the buncher was to find the characteristics of the beam after the RFQ and linac. These structures have the greatest impact on the longitudinal structure of the beam, and the entire motivation for this project was to have an impact on those characteristics.

Once the approximate shape of the FWHM to voltage curves had been plotted, this information was used to make a rough estimate of the bunching voltage needed when the two modes were combined. Using this estimate as a starting point, both modes of the buncher were turned on and the beam observed at L061 again. The two modes were then adjusted empirically to produce the best possible focus in a reasonable amount of time. The final control system settings used were 305 volts for the first mode and 55 for the second. Based on the more thorough calibration performed after the fact, this was likely slightly underfocused. These values were then scaled down (by the ratio of the focal lengths) to

move the focus 25 cm forward to the entrance of the RFQ. The experiment was then paused to allow the RFQ to warm up to its stable operating temperature.

In principle, once the beam has been brought to a longitudinal focus at the RFQ, there is only one remaining buncher parameter to optimize - the relative phase between the buncher and the RFQ. The relative phases between the modes were expected to be correct within the LLRF module and the CSS control application - i.e. if the first and second mode phases were both set to “0”, both would have their ascending zero crossings at the same time. However, the relative phase between the 16.1 MHz buncher and the RFQ is determined by the cable lengths between each device and the master clock, as well as time of flight for particles from one to the other, and any delays due to electronics. While these could all in principle be calculated, it would seem to be far easier to simply find the correct relative phase empirically.

In order to accomplish this, the L110 timing wire, located at the end of the linac was used. The control system settings used for the cavities in the first two cryomodules are given in Table 7.6. None of the cavities in the third cryomodule were used for this tune. As a reminder, the choice of $^{40}\text{Ar}^{13+}$ for these measurements was specifically motivated by the fact that the linac had already been tuned for this beam, so minimal retuning was needed, other than some adjustment of the focusing solenoids to refocus the beam at the location of the timing wire.

Once the beam was accelerated to the end of the linac and brought to a focus, the L110 timing wire was inserted and the longitudinal beam profile observed. Even with no adjustment whatsoever, the predicted large peaks separated by four satellite bunches were immediately apparent on the histogram, as seen in Figure 7.15. This was attributable largely to luck as it was possible that the bulk of the beam could have been bunched entirely outside of the acceptance of the RFQ, but was a promising sign that the buncher was having its

Cavity Name	Control System Setting
L077	3.3
L082	15.78
L084	23.17
L085	22.68
L088	22.62
L089	27.61
L091	0

Table 7.6: Cavity settings for testing of the buncher with accelerated beam. All of the cavities in the third cryomodule were set off.

desired effect.

Next, the two modes of the buncher were gradually moved upwards in phase until the maximum amount of counts were present in the main peaks, corresponding to the buncher focus being aligned in phase with the RFQ acceptance window. Finally, some experimentation was performed to see if the bunching performance could be improved still further with different amplitude settings on the two modes. Given that the RFQ acceptance is not a perfectly upright rectangle, as outlined in Chapter 4, it would not have been surprising if other shapes than perfectly vertical longitudinal focusing at the RFQ entrance could produce improved bunching results. Indeed, some improvement was seen. The best result achieved is shown in Figure 7.16. The buncher settings used for this case were amplitudes of 335 V and 35 V for modes one and two respectively, and phases of 22 and 202.

One note about phases - the phase of the second mode was actually set erroneously once it was moved away from its starting value. It was determined after the conclusion of the experiment that the units for the phase of the second (and third) mode are degrees relative to the frequency *of that mode*. So if the zero crossings of mode 1 and mode 2 are aligned, and the phase of mode 1 is shifted by twenty degrees, in order to keep the two crossings aligned, the second mode must be shifted by *forty* degrees. Since the second mode was shifted during




Figure 7.15: Initial beam time structure measurement taken at L110 with no phase adjustment.




Figure 7.16: Best case bunching achieved at L110 with amplitudes 335 V and 35 V, and phases 22 and 202 degrees.

the actual observations by half of the required amount, but that amount was relatively small, some improvement in bunching is conceivably possible just from correcting that error.

Once the timing wire measurements were concluded, two other measurements were made. First, the beam was steered through the dipole just after the linac and the dipole field measured to give an empirical measurement of the beam energy. Secondly, the beam current was measured before and after the RFQ and linac to give a measurement of total transmission in the unbunched, 80.5 MHz bunched, and 16.1 MHz bunched cases. This data is summarized in Table 7.7. Approximately 30% of the beam was accelerated to the end of the linac even with no bunching before the RFQ at all. This is consistent with the values observed during

Summary	
Beam	$^{40}\text{Ar}^{13+}$
Energy	Initial: 12 keV/u After Acceleration: 1.6 MeV/u
Transmission	No bunching: 30% 16.1 MHz bunching: 50% 80.5 MHz bunching: 75%
Bunching Efficiency	85%

Table 7.7: Summary of bunched beam properties. “Transmission” gives the percentage of the beam current measured after the linac relative to the beam current measured just before the buncher. “Bunching Efficiency” gives the percentage of the total beam detected at L110 in the main pulse.

the RFQ acceptance measurements.

7.5 Comparison with Simulation

As a verification, after the conclusion of the experiment, a DYNAC simulation was run using the tune settings used for both the 80.5 MHz bunching case and the 16.1 MHz bunching case. This also allows inferences to be made about beam characteristics that were not directly observed during the measurement. This simulation was run for a period substantially longer than the 62 ns to ensure that the simulated satellite bunches included any particles from main peaks on either side. For clarity, only the satellites from a single buncher period are shown in the figures.

The aspects of the simulation which can be reasonably assumed to be in excellent agreement with the reality of the beamline are: the physical geometry and placement of the beamline elements and the amplitude of the applied voltage to the RFQ and the electrostatic elements in the low energy section of the beamline. Somewhat less certain, but still consistent with observation, are the current/field curves for the magnetic elements and the

accelerating fields of the superconducting cavities. The largest uncertainties in the simulation come from the initial conditions out of the EBIT and the exact behavior of the beam through the RFQ. Additionally, the settings for the buncher itself are tuned in simulation to provide maximum transmission through the RFQ, but the real buncher was only tuned empirically, so there are likely some discrepancies between the simulated and actual buncher fields.

Despite these sources of error, the simulated and measured bunched beams were quite similar in characteristics. The measured beam energy from the dipole field after the linac was 1.60 MeV/u. The simulation predicted a final beam energy of 1.64 MeV/u (65.55 MeV total for this beam), a difference of only 2.5%. A side by side comparison of the time structure of the observed beam with the simulation at L110 is shown in Figure 7.17. In each case, the number of counts in the main peak was compared with those of the sum of the counts in the four nearest satellite bunches (two before and two after the main peak).

The best case bunching efficiency achieved in simulation was 88%, slightly higher than the 85% achieved with the real beam. It is likely that had more time been available for tuning and calibration, at least some improvement in the real efficiency could have been achieved.

The simulation also allows for the estimation of a quantity not directly measured during the experiment, namely the energy spread of the beam after acceleration in the linac. While the predicted energy spread immediately following the buncher is on the order of \pm several percent, the acceleration and longitudinal focusing effects of the RFQ and linac reduce this relative deviation. The predicted absolute energy spread is ± 0.71 MeV (3σ), or a relative difference of about ± 1 %.




Figure 7.17: Side-by-side comparison of the time structure of the beam at L110. The left hand image is the actual histogram measured on the timing wire, and the right hand image is a simulation of the same conditions.

Chapter 8

Conclusions

Based on the stated needs of a number of experimentalists at the National Superconducting Cyclotron Laboratory, a 16.1 MHz multiharmonic buncher has been designed, installed, and tested in the ReAccelerator linac. In addition, valuable tools have been developed for modeling the beam through the ReAccelerator, and important information gathered about the existing facility that should inform future beam development.

8.1 Summary

There are two significant research focuses in this document - the empirical measurement of the longitudinal acceptance of the ReA RFQ, and the design and construction of the 16.1 MHz multiharmonic buncher. While the buncher will likely see more widespread use, the RFQ measurements also help increase our understanding of the performance of this complex device.

The measurements of the RFQ longitudinal acceptance were carried out using a test beam which was used as a probe of the energy and time dimensions of the phase space area in which particles will be successfully accelerated through the RFQ. The probe beam was first made as narrow as possible in energy to probe the ranges of energy which would be accepted, and then as narrow as possible in time to probe the phase acceptance. As it is possible to make the beam much narrower in energy than in phase, the predicted accuracy

of the measurement is much better in the vertical dimension on a time vs. energy plot.

The final result found that the measured acceptance window was at least as wide, if not wider than the predicted window. This result has significant implications, not just for the buncher constructed through this project, but for any further modifications to the ReA beamline upstream of the RFQ.

The multiharmonic buncher proved to be a complex design project, drawing on the expertise of a large number of faculty and staff at the NSCL. Once the acceptance of the RFQ was determined, studies were conducted to determine appropriate values for the location and power required to bunch the beam, and the eventual effects on the longitudinal beam properties at the end of the linac. These studies required the use of a number of different accelerator modeling codes, and the development of a new software frontend, DYNACGUI.

The physical geometry for the bunching electrodes was simulated using 3D numerical modeling software, and further simulations conducted to predict the behavior of the beam through various electrode geometries. The final selection of electrode geometry was based on tradeoffs between ease of tuning and amplifier power, and between electrode capacitance and transit time factors. Once selected, the electrode and RF specifications were generated, and the device was constructed, installed, and aligned utilizing appropriate NSCL technical staff.

Finally, the buncher was tested in open and closed loop modes, and was found to perform in a manner consistent with simulation. While satellite bunches remain a serious issue, the beam is sufficiently bunched even with only two modes to allow for time of flight measurements for beams with reasonably high beam currents in some cases.

One other outcome of this experiment which may prove useful in the future was the gain in experience of using the timing wire detectors in the ReAccelerator beamline. While these

detectors had been used in the past, the TAC/MCA data acquisition path used here has proven far easier to use than the somewhat challenging software used previously.

8.2 Future Plans

At the time of this writing the 16.1 MHz buncher has been demonstrated to effectively bunch the beam with a 62 ns separation between primary pulses. However, there are a number of areas which could be profitably improved and explored to improve the efficiency of the device and its usefulness to experimentalists.

- Stabilization of the Third Mode

During the open loop testing of the buncher, the 48.8 MHz mode was shown to function properly when driven independently by a signal generator. However, during the closed loop testing, it was not possible in the time allotted to both stabilize the third mode and keep the power consumption of the first mode to a reasonable level. As such, the third mode was jettisoned for the final measurements.

If the third mode can be stabilized in closed loop mode, it is predicted to add up to another 5% bunching efficiency, which will greatly reduce the statistical noise faced by experimentalists doing time of flight measurements. This is illustrated in Figure 8.1 which shows a simulated beam time structure at L110 using three bunching modes. While this figure shows the best bunching efficiency achieved so far in simulation, it is possible that this represents only a local maximum in parameter space, and that further improvement may be possible.

- Improved Calibration Measurements




Figure 8.1: A simulated time distribution using all three bunching modes at L110. The calculated efficiency is 92.96%.

The single mode measurements in closed loop mode were encouragingly consistent, but were hampered by the low beam current of the argon beam and by the undetected presence of a secondary peak in the time spectrum. As such, while the existing calibration data is a good start, improved data, possibly collected with a higher current beam, could improve our understanding of the power to voltage relationship in the device. In addition, any future activation of the third mode would require a set of calibration measurements which have not yet been taken.

Further profitable measurements could also be made of the performance of the beam later in the beamline than the L110 timing wire. In particular, the evolution of the time structure of the beam to the target area would be valuable information for experimentalists potentially wishing to use this bunch structure. Also remaining to be measured are the effects on the time structure of using later cavities in the linac at

rebunching phase.

- Control Application Improvements

The control application proved very useful during the closed loop testing, but at the time of the measurements, the phase offset relative to the master clock had to be set manually each time the RF was turned on. As this process involves simple observations and computations, it could easily be automated. It also remains to be seen whether controls for this device can or should be integrated directly into the existing ReA control system, ROCS, (ReAccelerator Operations and Control Software) or into a potential future CSS based control scheme.

- Documentation

For the final goal of making this device accessible to users, the beam specifications must be properly documented in a simple form and made available as part of the ReA Service Level Description [50]. At a more detailed level, procedures for operating and tuning the device should be documented and made available to the ReAccelerator operators for use in operating the buncher in a production environment.

While the buncher as it stands at the time of this writing is already capable of providing beams with a wider average separation to users, these additional steps will move it from an experimental device to another reliable tool in the hands of the NSCL users.

8.3 Satellite Bunches

As predicted in Sections 3.4 and 5.7, the final results from the buncher did display satellite bunches as a result of particles seeing the “tails” of the bunching waveform. Section 3.4 also

outlined a number of possible avenues for cleaning these satellite bunches. This section will give a brief outline of two possible approaches to addressing these bunches in the specific case of ReA.

8.3.1 Integer Ratio Cavities

As described in Section 3.5.2, one way to clean satellite bunches is to use a second buncher. This buncher should have a frequency which is at an integer ratio to that of the main accelerator, such that it will also transport the primary bunches from the subharmonic buncher, but not the satellite bunches. The satellite bunches will instead be placed at incorrect energies to transport through the rest of the accelerator.

In practice, this would actually likely mean more than one cavity, as a single sinusoidal buncher with enough amplitude to deflect the satellite bunches away from the reference energy would also have a significant impact on the structure of the main bunch, and a multiharmonic buncher would be impractical at the required voltages. A series of such cavities would allow for the deflecting “kicks” to be applied at each cavity passage to the satellite bunches without excessive distortion of the primary bunch.

In the case of ReA, these cavities would likely need to be at either a frequency of 64.4 MHz or 96.6 MHz, representing either a 4:5 or a 6:5 ratio with the main RF frequency of the accelerator. The most promising location for these cavities would be in place of the existing first cryomodule. After the linac, the higher beam energy would require much more deflecting voltage to filter the satellite bunches effectively.

At present, the first cryomodule consists of two superconducting solenoids and a quarter wave superconducting cavity. The cavity is used exclusively as a buncher, to longitudinally focus the beam for more efficient acceleration through the linac. Because the required voltage

for this bunching is very low, the cavity operates at a substantially lower voltage than other devices in the line, within the range of what could reasonably be accomplished with a normally conducting cavity [51]. The fixed 600 keV/u beam energy at this point is also in a regime where the beam could be focused using normally conducting quadrupoles.

There are two significant drawbacks to this approach. The first is the large expense of removing the first cryomodule and replacing it. In addition to the two cleaning cavities, the focusing capabilities of the two superconducting solenoids located in the that cryomodule would need to be replaced. While neither the cavities nor the new focusing elements would need to be superconducting, this would still represent a substantial investment of resources and effort.

Since these proposed integer ratio cavities would only align in phase with the RFQ and the linac when operated in 16.1 MHz mode, they would need to be turned off when ReA was operated in normal 80.5 MHz mode. As such, the first cavity of the *second* cryomodule would have to be dedicated to longitudinally focusing the beam. This new rebunching location would give the beam more time to spread out before being rebunched, and less time to come to a focus before entering the accelerating cavities, both of which would require an increase in the beam energy spread for effective longitudinal focusing. Secondly, dedicating this cavity to rebunching would reduce by one the number of cavities available for acceleration, lowering the maximum possible energy of the final beam. This effect would be permanent, even for beams operated in 80.5 MHz mode.

8.3.2 Beam Chopper

The more conventional approach to removing unwanted bunches is a beam chopper, also described in Section 3.4. A pair of deflecting plates apply an electric field in order to

“kick” unwanted beam to the side. The voltage on the plates is modulated so that the deflecting voltage is zero for the desired portions of the beam. The transverse deflection of the unwanted particles must be sufficient such that downstream of the chopper, the deflected and undeflected beams do not overlap. This allows a slit to be inserted to allow passage of the desired beam particles and block the satellites.

In the specific case of ReA, the limitation on this method is the need for a kicker which can cycle fast enough to either deflect all the particles destined for the tails before they reach the buncher, or else kick all of the unwanted particles aside afterward. Both of these would require a square pulse with approximately 12 - 15 ns of combined rise, hold, and fall time, with a repetition frequency of 16.1 MHz, a capability which is not outside of the realm of possibility, but quite difficult to achieve with off the shelf equipment [52].

The other specific consideration for ReA would be the location of the chopper and its corresponding slit. The more space allocated for the chopper, the longer the beam would see the deflecting field, and the less voltage would be required for the kick. Correspondingly, the longer the distance between the chopper and the slit, the less voltage would be required. As low a voltage as possible is critical, given the aforementioned speed with which it would need to be switched.

The specific locations in ReA where a chopper could be located are upstream of the buncher, between the buncher and the RFQ, between the RFQ and the linac, or after the linac. Locations immediately before and after the RFQ are problematic for reasons of limited space - the distance from the RFQ to the first cryomodule is only 80 cm, and the space between the 16.1 MHz buncher and the RFQ is completely occupied by necessary focusing and bunching elements.

Locating the chopper after the linac would require a very high voltage. For example, for

a beam with a final energy of 4500 keV/u and a charge to mass ratio of 1/5, separating two beams by 5 mm in a chopping distance of 8 m would require approximately 30kV applied to a 7 cm long sine wave chopper, or 2 kV applied to a 50 cm traveling wave chopper. Neither of these would seem practical at a frequency of 16.1 MHz.

The remaining location would be upstream of the 16.1 MHz buncher, and this would appear to be the most promising possibility. The beam there has a low energy, and there are a few potential locations for a chopper which could be used to separate the beam before entering the 16.1 MHz buncher. Although a number of substantial engineering hurdles would need to be overcome, this would seem to be the best starting point for producing completely clean 16.1 MHz bunching for users.

8.3.3 EBIT Switching

The other approach outlined in section 3.4 which bears consideration is fast switching of the trap electrodes on the EBIT. The electrodes would need to be switched with a square pulse with a 50 ns combined rise / hold / fall time with a voltage shift on the order of 500 - 1000 V. However, it remains a tantalizing possibility, given the extreme desirability of the result.

If the emitted pulse from the EBIT could be shortened to 50 ns or less, that entire pulse could be compressed by the 16.1 MHz buncher into one bunch. In that case, the final bunch separation at the detector would be determined not by the frequency of the buncher, but by the frequency with which the trap was opened, giving the operators complete control of the separation between bunches. In practice, there would likely be a *minimum* rather than a maximum separation, based on the maximum switching speed of the electrodes. In this mode, no beam would be lost to the satellite bunches or unaccelerated “tails” because all of the transmitted beam would be compressed by the buncher.

This approach faces substantial hurdles - heating of the superconducting EBIT due to the fast switching being the most obvious. However, because it neatly solves so many problems, it is still worthy of serious consideration.

APPENDICES

Appendix A

Comparison of Accelerator Codes

At first glance, simulating the motion of a single particle in a linear accelerator may appear to be a relatively simple task. The Lorentz force law that determines the motion of a charged particle in an electromagnetic field is a simple equation, and the fields produced by various steering and accelerating devices are well understood. This motion becomes more complicated when higher order effects such as particle interactions are included, but to first order the motion of particles in electromagnetic fields would seem to be a simple calculation.

As always, the devil is in the details. A large number of codes have been written to simulate accelerators, and each has advantages and limitations. This overview of codes is by no means exhaustive, covering only the codes that were used in the preparation of this dissertation and a few of the major alternatives that were not used.

In general, accelerator simulation codes can be classified into two types: matrix codes and numerical field models [53]. Matrix codes calculate the analytical transport matrix for each element up to a specified order and combine them to produce an overall mathematical description of the accelerator, through which particles may then optionally be tracked. Numerical field codes simulate the three dimensional electromagnetic field of each element and numerically track the motion of particles through those fields with finite element modelling. Some codes use a combination of both approaches, which will be noted below.

A.1 Dynac

- Authors: Lapostolle, Valero, and Tanke
- Institution: CERN
- Type: Matrix code with particle tracking
- Website: <http://dynac.web.cern.ch/dynac/dynac.html>
- Open Source: Yes

A.1.1 Overview

DYNAC was the primary code used for this dissertation. It was selected for its ability to model each element of the ReA beamline in a single simulation run. This includes the low energy electrostatic elements at the start of the ReA line, the RFQ, the accelerating cavities, as well as the buncher itself. Another advantage of DYNAC is that it is open source, so it is relatively straightforward to look at the source code to determine the inner workings of an element or routine, or add needed functionality. (Like any program fork, this should be done with caution.) It is also well-documented, and still being actively maintained and updated at the time of this writing.

Appendix B of this dissertation describes DYNACGUI, a graphical front end for DYNAC written in the MATLAB environment by the author.

A.1.2 Simulation Method and Capabilities

DYNAC simulates a beamline by calculating the first order transport matrix for each element and then transporting the particle distribution through the matrix for that element. (Second order matrices may be specified for some, but not all elements.)

Certain elements are treated quasi-numerically. The electric field in an accelerating cavity may be specified by an external file which gives the on-axis accelerating field as a 1D function of distance, or by the Fourier components of the on-axis field. The 3D field components are then extrapolated away from the axis and the particles transported through this field distribution based on the size and length of the cavity. Solenoids may be treated either by specifying an arbitrary 1D field distribution, or by a simple hard edge model.

RFQs are also treated quasi-numerically in a manner similar to the RFQ design code PARMTEQ [54]. The geometry of the RFQ is specified by an external file which divides the device into a number of cells. The 3D field components of each cell are calculated based on the cell type and parameters, and the particles transported through each cell individually. Bunchers are treated as zero length elements which provide an instantaneous kick to the beam. Each individual bunching mode is specified as an independent device.

DYNAC provides support for a number of capabilities not utilized in this project, including beams consisting of multiple charge states, space charge, modeling of electron guns, synchrotron radiation, and IH type devices. DYNAC also provides routines for the simulation of misalignments.

A.1.3 Input and Output Files

The accelerator lattice for DYNAC is specified by an input file, which gives the parameters for each device on the beamline as well as auxiliary information, such as changes in the reference frequency or the locations to generate output graphs. DYNAC can generate an initial particle distribution using a variety of models, or can import an arbitrary particle distribution from an external file. Optional external input files can also be used to specify field distributions for accelerating cavities, RFQs, and solenoids.

On execution, DYNAC produces a number of output files. Each run will produce:

- *dynac.short* - a summary of basic information for each beamline element and the number of particles remaining at that point.
- *dynac.long* - a more verbose element-by-element description of the beamline.
- *dynac.print* - a tabular summary of the beam properties at each element.
- *emit.plot* - the raw data for all requested plots. (see A.1.4)
- *dynac.in_pr.dst* - the initial particle distribution
- *rfq_list.data,rfq_listmid.data,rfq_lost.data,rfq_coef.data* - various data about the RFQ.
(if used)

These files take the same names after each run, so care must be taken in managing output data. In addition, optional files can be produced to store particle distributions at various points on the beamline.

Sample code fragment:

```
;ReA3_46Ar15_L016_to_JANUS_3MeV.in
GEBEAM
5 1
8.05e+07
1000
0 0 0 0 0 0
0 0.4 10
0 0.4 10
0.0001097 0 0
INPUT
931.494 46 15
0.552 0
DRIFT
14.2968
QUAELEC
15. 1.1316 5.
```




Figure A.1: A sample emittance plot produced using DYNAC

A.1.4 Data Visualization

When the main input file for DYNAC is created, the user may specify locations for various types of plots, such as emittance plots, beam profile graphs, phase and energy width as a function of beamline position, etc. All of the data for the requested plots is stored in one large file, *emit.plot*.

The DYNAC package is distributed with DYNPLOT, a set of routines for plotting graph data stored in the *emit.plot* file. DYNPLOT converts the stored plot data to GNUPLOT format and displays the data using that program, which must also be installed locally. While DYNPLOT allows plots to be displayed or saved, it is not particularly flexible, in that the plots may only be viewed sequentially, and may only be saved in GNUPLOT format. DYNACGUI addresses a number of these issues by using a menu-based approach to the generated plots, and MATLAB's own built in plotting functions. Figure A.1 shows an example of an emittance plot generated

with DYNPLOT.

A.1.5 Limitations

The main drawbacks of DYNAC for this project were its lack of any sort of fitting capability and its inability to transport particle distributions longer than one period through an RFQ model. Both of these issues were addressed with DYNACGUI. DYNACGUI was also used to improve on the limitations of the built in visualization routines. DYNAC is also not the ideal choice for users who need to use detailed three dimensional field models for elements rather than analytical approximations. DYNAC does not include any built in functionality for analyzing periodic systems or rings.

A.2 COSY INFINITY / *cosy.for*

- Authors: Berz and Makino
- Institution: Michigan State University
- Type: Matrix Code (see below)
- Website: http://www.bt.pa.msu.edu/index_cosy.htm
- Open Source: Beam Libraries Only

A.2.1 Overview

COSY is not exclusively a beam physics code. It is a programming environment designed for scientific computing. It includes a library of routines (*cosy.for*) specifically designed for beam physics applications. While COSY itself is not open source, the beam physics libraries are.

What is unique about COSY is that it uses Taylor series and differential algebra data types to store much of its data. This has the advantage of allowing accurate, high order computations of transfer matrices to be calculated in reasonable amounts of time. However, this approach also has drawbacks. For example, modeling a periodic signal consisting of multiple frequency components using Taylor series can require unreasonably high order computation. This limitation makes it impractical for simulation of multiharmonic bunchers. The COSY environment is sufficiently powerful that this limitation could presumably be overcome, but this was deemed not to be the most efficient approach for the project.

COSY is still under active development at the time of writing, and features extensive documentation of the underlying language and of the beam physics routines.

A.2.2 Simulation Method and Capabilities

Since COSY is a programming environment, beamline models in COSY take the form of full programs written using the *cosy.fox* library of beamline elements. This allows for extensive programming and analysis capabilities within the beamline model itself, as opposed to requiring postprocessing of results as with DYNAC and most other codes. The standard beamline elements include a wide range of magnetic, electrostatic, and optical elements, which can be simulated to arbitrarily high order. Further, COSY allows detailed treatment and specification of fringe field models. Finally, *cosy.fox* includes a number of routines for the analysis of periodic structures, spectrographs, and aberrations.

By default, COSY does not track ensembles of particles through a beamline. Rather the overall transfer matrix of the line is calculated to arbitrarily high order, and then rays are traced through the matrix. This approach is useful when a geometrical optical approach to the beam dynamics is warranted, but is less helpful when one wishes to specify an input

particle distribution as a series of particles or by its Courant Snyder parameters.

One area in which COSY excels is fitting using its built in routines. Using COSY, it is very easy to specify objectives to be achieved and parameters to be varied, and then use a variety of algorithms to match the desired objectives. The fitting routines in COSY are among the most powerful and easy to use of any beam dynamics code.

The other advantage of the fact that COSY is a programming environment, rather than simply a dynamics code, is that desired capabilities that are not provided by default can be coded by the end user. The initial modeling of ReA was done by Mauricio Portillo using COSY[18]. Since COSY does not implement tracking of particle ensembles by default, Portillo’s model includes a large number of routines to generate, track, and analyze particle distributions through the transfer maps produced by the beamline model. These routines allow for either the import of arbitrary particle distributions, or to randomly generate such distributions using initial CS parameters. In turn, these distributions can be saved and then analyzed, and fitting performed based on the properties of these final distributions. This type of distribution based fitting is not as easily accomplished with “stock” COSY.

A.2.3 Input and Output Files

As COSY is a programming environment, rather than a beam physics code *per se*, the only fixed file requirement is for the beamline code itself. As with a program written in any other language, the user may specify input and output files with any desired names or formats. The only input file requirement is that the compiled binary of the *cosy.fox* beam library be available to the beamline code.

Sample code fragment:

```

INCLUDE 'COSY';
PROCEDURE RUN;
  VARIABLE Q1 1; VARIABLE Q2 1; VARIABLE OBJ 1;
  Q1 := .1; Q2 := -.1;
  OV 5 2 0 ;
  RP 10 4 2 ;
  SB .05 .05 0 .05 .05 0 0 0 0 0 0;
FIT Q1 Q2;
  UM ; CR ;
  ER 1 3 1 3 1 1 1 1 ;
  BP;
  DL 2 ;
  MQ .4 Q1 .15 ;
  DL .2 ;
  MQ .4 Q2 .15 ;

```

A.2.4 Data Visualization

As indicated above, natively COSY is not set up to track large ensembles of particles, but rather individual rays or small ensembles of rays. It does, however, have quite simple to use tools for displaying the behavior of those rays. In order to access those tools, one must first successfully set up one of the supported graphics interfaces (PGPLOT, GRWIN, AQUATERM) none of which are particularly intuitive or contemporary. With that accomplished, COSY produces very good plots of rays being traced through the system, with the locations and relative sizes of beamline elements clearly indicated. Furthermore, due to the programmatic nature of COSY it can update these plots as it tries various solutions for a fitting routine, so the overall suitability of the fitting strategy can be visually evaluated.

COSY does have the ability to display the envelope of the beam, consisting of the maximum transverse extent of any initial rays at a given point. It can also display curved reference trajectories, a capability not shared by most other beamline codes. Figure A.2 shows an envelope plot generated using COSY.




Figure A.2: An example of a longitudinal plot produced with COSY

A.2.5 Limitations

In general, the *cosy.fox* library is less flexible in the longitudinal dimension than in the transverse direction. Partially, this is due to the systematic limitations noted above, and partially due to the fact that there are simply very few accelerating elements available in the library. Further programming could certainly be done by an interested end user to overcome this second limitation. One possible way to overcome the systematic issues might be to further expand the particle tracking routines constructed by Portillo.

One other factor which prevented use of COSY for modeling the entire ReA beamline is that while RFQs have been successfully modeled in this environment, those routines were not publicly available at the time of writing.

A.3 TRACK

- Authors: Ostroumov, Aseev, and Mustapha
- Institution: Argonne National Laboratory
- Type: Numerical Tracking Code

- Website: <http://www.phy.anl.gov/atlas/TRACK/>
- Open Source: No

A.3.1 Overview

TRACK is designed primarily as a numerical simulation code for tracking distributions of particles through 3D numerical simulations of electric and magnetic fields generated by beamline elements. An arbitrary number of field files can be specified, and these files can be reused throughout the beamline. TRACK can also generate field models for idealized versions of many common elements such as quadrupoles and solenoids.

TRACK was still in active use at Argonne National Laboratory at the time of this writing, but the documentation had not been updated to reflect the most recent versions of the code.

A.3.2 Simulation Method and Capabilities

TRACK is a numerical particle tracking code intended for ion beams in linear accelerators. It will either generate a starting particle distribution given a set of parameters or import one from a distribution file. When generating the distributions within TRACK, the only initial distributions supported are 4D waterbag with a DC longitudinal beam or 6D waterbag. Although it will track particles through idealized models of some elements, such as the textbook thin lens quadrupole, the primary focus of the code is on tracking particles through 3D simulated field models.

Of specific interest to the present project is the ability of TRACK to simulate bunchers using a 3D field model. Given a static 3D electromagnetic field distribution for a buncher, TRACK can scale the amplitude of that field with a given frequency and phase for each

mode of the buncher. This allows for testing of the effects of specific electrode geometries, something which is not possible in either DYNAC or COSY. TRACK also includes the ability to simulate an idealized buncher without a field file for comparison purposes.

TRACK does support simulation of RFQs *via* a specific model file format which gives the expansion coefficients of the EM field of each cell of the device, in a similar fashion to DYNAC and PARMTEQ. However, the file format is not directly compatible with either, and it can prove difficult to make a direct conversion or to verify one code against the other.

A.3.3 Input and Output Files

Input to TRACK is given by two text files. The first, *track.dat*, specifies the options for running the program, including the initial conditions of the beam. The other, *sclinac.dat* specifies the beamline itself, including the layout and settings for each element. A 3D field distribution file must also be specified for each element for which a field model is used. The same model file may be used more than once in the same beamline for identical elements. Because the input file names are always the same, care must be taken to separate and identify each simulated case.

When 3D field files are used for elements, it is expected these files will be generated using the EM STUDIO component of the CST STUDIO [55] suite of modeling codes. (This is potentially an issue for some users, for although TRACK itself is free, the same is decidedly not the case for CST STUDIO.) The input files are simple ASCII tables of field values on a 3D lattice, so it should in principle be possible to generate them using other 3D field codes such as OPERA. TRACK has hard coded limits on the maximum resolution of the input lattice, although certain versions of the program may support different resolutions. (These exceptions are not well documented.) The conversion from CST STUDIO to a format usable

by TRACK is *via* a set of provided utility programs.

Sample code fragment:

```
1 drift 4.45 3.0 3.0
0 equad 6472 15.0 15.0 5.0 60
1 drift 3.0 3.0 3.0
0 equad -8414 15.0 15.0 5.0 60
1 drift 4.45 3.0 3.0

1 drift 10.5 3.0 3.0
1 drift 7.1 3.0 3.0

1 drift 4.3 3.0 3.0
!1 drift 4.1 3.0 3.0
!1 drift 7.6 3.0 3.0
!1 drift 1.8 3.0 3.0

4 mhb4 11.4 1 143.15 1 1485 0 2 -748 0 3 343 0 4 -150 0
```

A.3.4 Data Visualization

One of TRACK's greatest strengths is its ability to display graphical emittance plots for each step as the beam is tracked along the beamline. By comparison, DYNAC only produces emittance plots that can be viewed after the entire line is calculated, and COSY does not natively produce emittance plots at all.

Interaction with TRACK during the simulation is *via* a graphical interface which displays x , y , and z emittance plots, as well as a longitudinal profile plot during the run. The interface also displays important beam characteristics such as mass and charge, as well as providing options to advance by a selected number of steps at once. While the interface is extremely useful for visualizing the beam, there is unfortunately no simple way to entirely disable it for automated runs of the program. Although there is an option which will disable the display of the emittance plots, once the simulation is completed, the user is still required to manually

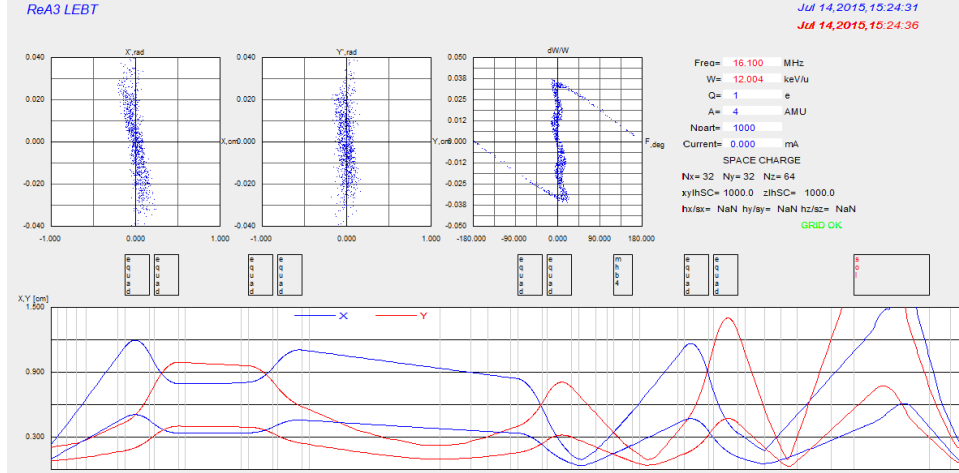


Figure A.3: An example of TRACK's output screen. (Phase vs. z plot not shown)

confirm that they wish to exit. Figure A.3 shows an example of the graphical interface for TRACK.

A.3.5 Limitations

TRACK does include some fitting functionality, but this functionality highlights one of the biggest problems with TRACK: the incompleteness of the documentation. Although there is an extensive (68 page) documentation file available for the code, at the time of this writing it was two versions out of date and not being maintained. The fitting functionality is not even mentioned in the documentation, and while a few sample files are available, it is a non-trivial undertaking to extrapolate how to do fitting using the code. This issue extends to a number of other undocumented changes to existing features or undocumented new features. And as TRACK is not open source, it is not possible to simply examine the source code to reverse engineer the behavior of the code.

Also, as TRACK is intended for linear heavy ion accelerators, it does not have any functionality for simulating synchrotron radiation or for analyzing periodic structures.

A.3.6 Errata and Comments

Provided here are a few corrections and comments on version 37 of the TRACK manual.

- The energy spread of the input beam is given as $\frac{\Delta\beta}{\beta}$, not $\frac{\Delta E}{E}$ as in most other codes.

$$\frac{\Delta\beta}{\beta} = \frac{1}{\gamma(\gamma+1)} \frac{\Delta E}{E} \text{ [24].}$$

- Undocumented: A four harmonic ideal buncher is given by:

```
n mhb4 d_elem rap MHB_ph0 1 Ef1 ph1 2 Ef2 ph2 3 Ef3 ph3 4 Ef4 ph4
```

where the meanings of these parameters are the same as for the two harmonic buncher.

- The value used for voltage for electrostatic quadrupoles is the voltage *difference* the two plates. This is *double* the value used by COSY and DYNAC which use the value of the voltage on a single plate relative to ground.
- igrph=0 is the command in *track.dat* to run the code without the GUI.
- The file *out.dat* is referenced in the v37 documentation. This has been replaced with the file *beam.out*.
- The DW/W [rel.u.] column in the *beam.out* file is the 1 RMS value, not the 100% value as indicated in section 8 of the manual.
- Undocumented: iread_dis=2 and iwrite_dis=2 can be used to read and write particle distributions in ASCII, rather than binary format.

A.4 MAD

- Authors: Grote, Schmidt, Deniau, and Roy
- Institution: CERN
- Type: Matrix Code
- Website: madx.web.cern.ch

- Open Source: Yes

A.4.1 Overview

MAD and its variants have a long history in the design of large scale accelerators and storage rings. Developed at CERN, Stanford, and Fermilab, MAD was the primary code used in the design of the LEP and LHC, among other machines. MAD is primarily geared toward high energy, periodic systems, but has a number of modules that expand its capabilities in other ways.

Like COSY, MAD executes scripts rather than merely interpreting a description of a beam-line, so programmatic control of fitting or other routines from within MAD is possible. The code is still in active use at CERN and at other laboratories around the world, and is still under active development. In addition to extensive documentation, there are also a wide range of sample files and introductory guides available.

There are a number of MAD descendents and variants, including DIMAD and BMAD which also extend or differ from the parent program in various ways.

A.4.2 Simulation Method and Capabilities

MAD is primarily concerned with the analysis of the machine lattice. To that end, the user defines “sequences” using a simple scripting language. The language is deliberately constructed with accelerator simulation in mind, so many commands are specifically tailored for this purpose. For example, element positions can be defined relative to other elements with a single command.

Once the lattice is defined, MAD computes the matrix representation of the machine and

then performs any desired operations on that matrix. Built in modules allow for plotting of CS and dispersion functions, matching for desired conditions subject to constraints, conversion of coordinates to Cartesian coordinates for machine alignment, and a wide array of other accelerator tasks. While the default behavior is not to track ensembles of particles, this task can be accomplished either within the scripting language, or by the use of external modules or macros.

A.4.3 Input and Output Files

Like COSY, since MAD is essentially a scripting language, the basic unit of input is the program file, usually designated by extension **.madx*. Since all elements are represented by their matrix definitions, no auxiliary input files for field definitions are required. Some MAD modules do use input or output files to store data. For example, the orbit correction module uses “before” and “after” files for the values of the correctors and beam position monitors used in the process. It is also possible to break beamline segments, called “sequences,” out into external files, and then import them from the primary *madx* file.

The input file format used by MAD is used by a range of other codes in addition to MAD itself.

Sample code fragment:

```
while (n <= ncell-5) {
  qfsps: qfsps,    at=(n-1)*lcell;
  mbsps: mbsps,    at=(n-1)*lcell+16.0;
  qdsps: qdsps,    at=(n-1)*lcell+32.00;
  mbsps: mbsps,    at=(n-1)*lcell+48.00;
  n = n + 1;
}
m1    : m1    ,    at=(ncell-5)*lcell;
qf1   : qf1   ,    at=(ncell-5)*lcell;
```





Figure A.4: An example of a figure generated with MAD

```

mbsps: mbsps,    at=(ncell-5)*lcell+16.0;
qd1   : qd1 ,    at=(ncell-5)*lcell+32.00;
mbsps: mbsps,    at=(ncell-5)*lcell+48.00;

```

A.4.4 Data Visualization

MAD has built in support for writing graphical output files using Postscript. For plots relative to the axis of the machine, there is also the ability to add a cartoon schematic of the elements along the displayed portion of the beamline. Plotting is *via* basic commands within MAD, so no editing of plots is possible once they are generated. Figure A.4 shows a plot of the CS beta function generated using MAD.

A.4.5 Limitations

MAD was designed with the intention of simulating high energy beams. As such, it is important to make sure that if it is used to simulate beams far from $\beta = 1$ that no assumptions are being made which would compromise the results. As a result of this high energy focus, MAD does not include natively any electrostatic focusing elements, which precluded its con-

sideration for the present dissertation. However, should the need arise, MAD does allow for user defined elements. MAD also has no support for RFQs at the time of this writing.

A.5 Trace-3D

- Authors: Crandall, K.R. and Rusthoi, D.P.
- Institution: Los Alamos National Laboratory
- Type: Matrix Code
- Website: <http://laacg.lanl.gov>
- Open Source: No

A.5.1 Overview

TRACE3D is a pure matrix simulation code. The transfer matrix for each beamline element is calculated analytically from the parameters of the element. It is primarily intended for design and tuning of a simulated beamline to match the beam to some desired parameters. The beam itself is represented purely by a 6x6 σ -matrix and cannot consist of an arbitrary particle distribution.

TRACE3D is an extremely mature code. The original TRACE dated from 1973, and was developed for a number of different systems until the most recent version, released in 1997. It shows its age in the simple character based interface, but is nonetheless a very efficient tool for certain types of beam dynamics problems. It also features an extensive manual which is an excellent reference for transfer matrices for the elements supported by the program.

A.5.2 Simulation Method and Capabilities

As a pure matrix code, TRACE3D can only transport particle distributions that can be described in terms of a 6x6 σ -matrix. This rules out, for example, unbunched beams, or beams with non-ellipsoidal phase space projections. Once the initial beam is expressed in terms of this matrix, transporting that beam through the beamline is a simple matter of matrix multiplication. The only nonlinear effect that is simulated in this code is a simplified model of space charge, which is treated as an instantaneous electrostatic “kick” to the shape of the ellipsoid.

TRACE3D has a limited library of available beamline elements which it can simulate. Although the documentation refers to the ability to add user-defined elements to the code, since the source code is not actually available, the user is in practice limited to the provided elements.

One other built in capability of interest is the ability to calculate the emittance of a beam given measurements of the beam size at three points along the beamline. This is a well known technique, but the presence of a built in utility could potentially prove useful.

A.5.3 Input and Output Files

Each TRACE3D model consists of a single input file which contains all of the information on beam description, beamline elements, and simulation parameters in a single file. In addition, if any matching is to be performed, the parameters of this fit are also included in the file. All of this information is contained within the file in a set of name-value pairs. When this file is loaded into TRACE3D, the user has a series of one-letter commands for executing the simulation, editing the parameters from within TRACE3D, and visualizing the results in

various ways.

There are only two means to produce an output file. If the “Hard Copy” switch is enabled, TRACE3D will immediately store all output to a buffer which is dumped to the default printer when the switch is turned off. (Fortunately, there is the option to select an output file as the printer.) There is also a command option that will produce a summary table of the beam envelope size for a given run, although it will not include CS parameters or matching information.

Sample code fragment:

```
&DATA
ER= 1875.60000 Q= 1. W= 2.00000 XI= 100.000
EMITI = 60.000000 60.000000 1000.000000
BEAMI = 3.11400 0.75360 -2.62020 0.57890 0.12100 0.56870
BEAMF = 0.00000 0.41540 0.00000 1.32890 0.00000 0.35700
BEAMCI= 0.00000 0.00000 0.00000 0.00000 0.00000 0.00000
FREQ= 80.000 PQEXT= 2.50 ICHROM= 0 IBS= 0 XC= 0.0000
XM= 14.0000 XPM= 60.0000 YM= 20.00 DPM= 30.00 DWM= 100.00 DPP= 30.00
XMI= 14.0000 XPMI= 60.0000 XMF= 14.0000 XPMF= 60.0000
DPMI= 30.0000 DPMF= 30.0000 DDMI= 100.0000 DDMF= 100.0000
N1= 1 N2= 20 SMAX= 5.0 PQSMAX= 2.0 NEL1= 1 NEL2= 20 NP1= 1 NP2= 20
NPRIN= 6
IJPRIN=1,001 1,002 1,006 1,010 1,014 1,018
MT= 8 NC= 4 IPLANE= 0 0 0 MP=1,004 1,008 1,012 1,016 0,000 0,000
MVC=0,000,0 0,000,0 0,000,0 0,000,0 0,000,0 0,000,0
VAL= 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000 0.0000000
CMT(001)=’RFQ ’ NT(001)= 11 A(1,001)=-.71000000 110.4 86.44 -28.7 0.0
CMT(002)=’RFQ ’ NT(002)= 11 A(1,002)=0.71000000 110.4 86.44 -180.0 0.0
CMT(003)=’ ’ NT(003)= 1 A(1,003)= 173.00000
CMT(004)=’Quad #1 ’ NT(004)= 3 A(1,004)=-26.200000 96.0 0.0 0.0 0.0
CMT(005)=’ ’ NT(005)= 1 A(1,005)= 39.750000
CMT(006)=’RF Gap ’ NT(006)= 10 A(1,006)=0.18155000 -40.0 0.0 1.0 0.0
CMT(007)=’ ’ NT(007)= 1 A(1,007)= 39.750000
CMT(008)=’Quad #2 ’ NT(008)= 3 A(1,008)= 26.200000 96.0 0.0 0.0 0.0
```

A.5.4 Data Visualization

The main screen of TRACE3D (shown in Figure A.5) includes several available plots for visualizing the result of the beam calculation. The primary one of these consists of a pair of “before” and “after” emittance plots in the transverse (x and y overlaid) and longitudinal planes. In addition, this view shows a schematic layout of the beamline with the size of the

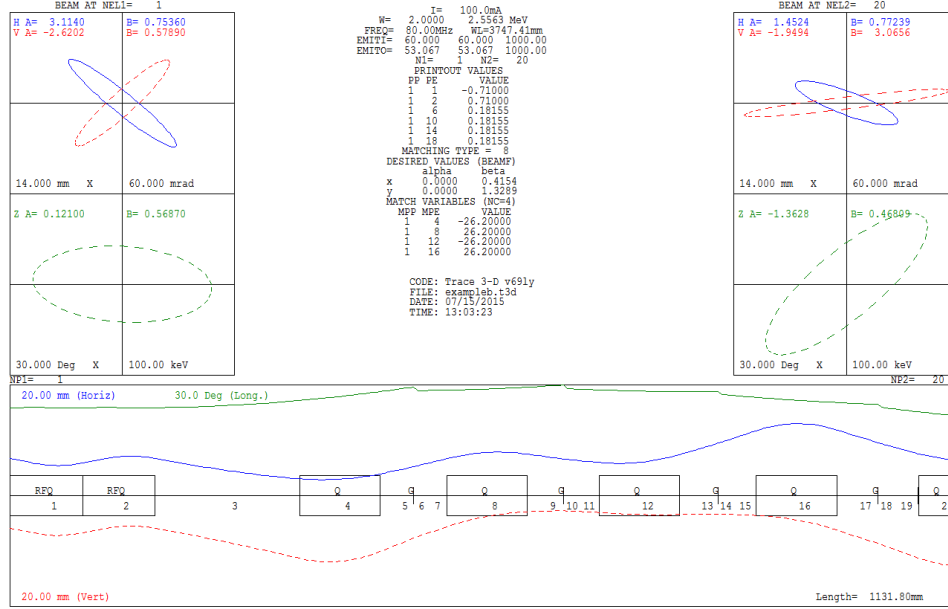


Figure A.5: The main screen of TRACE3D showing a profile graph and before and after emittance plots.

x , y and z envelopes overlaid on the plot. A useful feature is that subsequent simulations can be overlaid on the existing plot, showing the effect of changing certain beam parameters. The other available visualizations is a set of realspace projections of the beam.

A.5.5 Limitations

Because this code produces few output files, a decent analogy would be a “graphing calculator” for beamline simulation. The program allows for quick display of beamline envelopes, and will overlay the results of manual changes for easy visualization. It will also perform basic fitting of beamline parameters to achieve a desired beam shape. However, while the settings file can be modified and resaved, any other information must be either copied off the screen or printed, as not even rudimentary cut-and-paste functionality is supported.

A.6 TRANSPORT

- Authors: Brown, Carey, and Rothacker
- Institutions: Fermilab, SLAC, CERN
- Type: Matrix Code
- Website: <https://cdcv.s.fnal.gov/redmine/projects/fermitools/wiki/Transport>
- Open Source: Yes
- Graphical Transport Framework
 - Author: Rohrer
 - Institution: PSI
 - Website: http://aea.web.psi.ch/Urs_Rohrer/MyWeb/trans.htm

A.6.1 Overview

TRANSPORT is the predecessor and spiritual ancestor of virtually all modern accelerator codes. First written in ALGOL in 1963, it has been expanded repeatedly over the years, and is routinely cited by authors of other codes as a primary influence. The code is still in use today by some scientists, and has been fitted with a relatively robust graphical user interface, a rarity in this primarily text-focused category. The manual for the code is extensive and includes a great deal of theory, but care should be taken that the version of the manual used matches the version of the program being run.

A.6.2 Simulation Method and Capabilities

Like the other matrix codes described here, TRANSPORT calculates the transfer matrix elements for each item on the beamline and then aggregates these elements to produce an

analytic transport model for the entire line. The most recent official version of TRANSPORT can calculate matrix elements up to third order for most device types, which is a higher order than any other code described here besides COSY. Although TRANSPORT does not natively support particle tracking, a related code called TURTLE can be used for situations which require tracking, such as space charge.

Beamline elements supported by TRANSPORT comprise a relatively limited list, including multipole magnets, solenoids, and simple accelerating gaps, but not electrostatic elements, or any elements with a time dependent component such as bunchers. TRANSPORT does include the ability to use user-defined transport matrices.

In addition to simulation of beamlines to third order, TRANSPORT also includes a strong set of commands for beamline fitting. One of the strengths of the fitting routines in this program in particular is the ability to place explicit constraints on the independent variables to be adjusted.

A.6.3 Input and Output Files

While earlier versions of TRANSPORT used a native specification for the format of the input files, in 1984 the program was altered to also accept the input format used by MAD (see A.4). This does not include the more complex scripting commands used by MAD itself, but only the specific descriptions of beamline segments.

Upon execution, TRANSPORT generates a potentially lengthy single output file. This file includes a complete listing of the input line, the transport matrix for each element of the line, the beam matrix at each point in the line, and information on any other commands which may have been executed during the run. The “print” command may be inserted in the input file to request the output of a number of types of information about the beam or

the transport system at that point in the beamline. The requested data will be added to the master output file.

In addition, for each “plot” command issued in the input script, a separate output file will be generated with the data for that plot. Available plot types include plots of any variable vs. accumulated length, schematic layouts of the beamline geometry, plots of the beam ellipse, and plots of the effects of varying a specific beamline element on another variable. This final type of plot can also include fitting steps between each iteration.

Sample code fragment (original TRANSPORT notation):

```
0
13. 48.0 /ANG/ ;
15. 1.0 /MM/ 0.1 ;
16. 14.0 0.0 /SLIT/ ;
16. 3.0 1836.69 /MASS/ ;
-17. /SEC/ ;
-1. 2.5 3.5 2.5 4.2 0.0 0.0 1.2048 /BOT/ ;
1. 2.5 8.5 2.5 8.5 0.0 0.0 1.198 /BEAM/ ;
-1. 2.5 8.5 2.5 8.5 0.0 0.0 1.198 /NARR/ ;
-1. 2.5 8.5 2.5 8.5 0.0 1.0 1.198 /WIDE/ ;
-1. 2.5 8.5 2.5 8.5 0.0 0.0 1.198 /BEA2/ ;
3. 0.03 ;
3. 0.19 ;
6. 1.0 13.44 3.0 13.44 ;
3. 0.25 ;
6. 1.0 28.71 3.0 28.71 /KHE0/ ;
3. 1.51 /MHB5/ ;
3. 0.02 ;
6. 1.0 40.0 3.0 40.0 /KHE1/ ;
```

A.6.4 Data Visualization

While TRANSPORT has the ability to generate data for plots as outlined above, it does not have any built in mechanism for plotting this data directly, so it must be imported

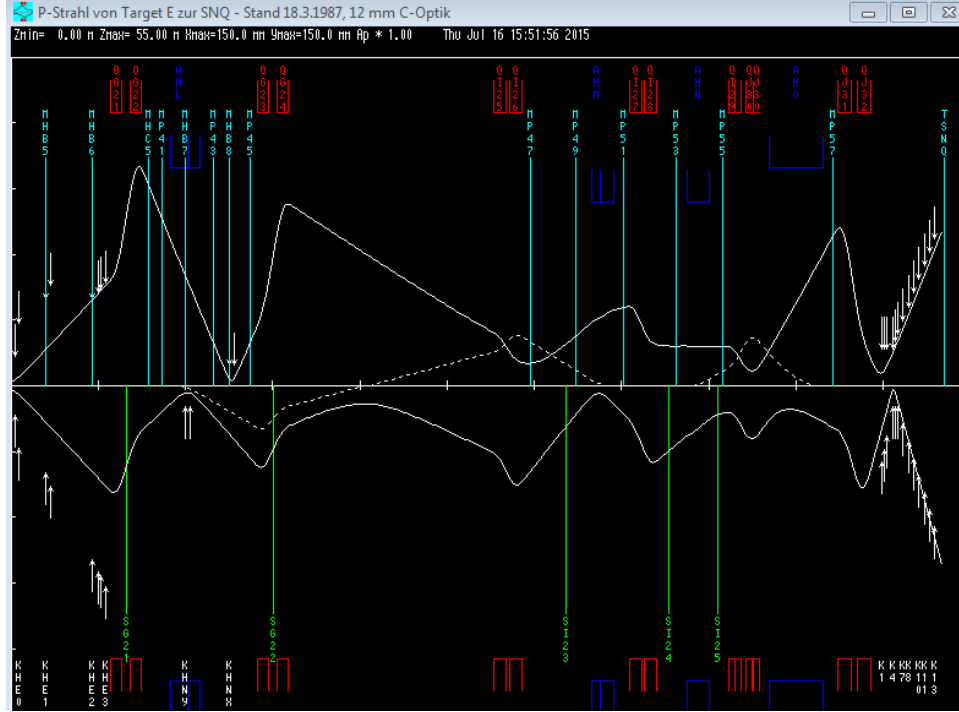


Figure A.6: A sample plot generated by the GRAPHIC TRANSPORT FRAMEWORK interface to TRANSPORT.

into another program for graphical visualization. However, a GUI for TRANSPORT, titled GRAPHICAL TRANSPORT FRAMEWORK has been developed by Urs Rohrer at PSI. In addition to allowing the direct plotting of the the plots specified above, this interface also allows for visualization of most of the data in the output file. A further capability added by this software is support for a point and click approach to building the beamline in the first place; elements can be defined, arranged, edited, and deleted using a menu based system rather than by editing text files. Figure A.6 shows a sample plot produced by this GUI.

A.6.5 Limitations

TRANSPORT is no longer in wide use - the official CERN/SLAC/Fermilab version has not been updated since 1999. The GUI version available from PSI is more recent (2006), but it is unclear if either has ongoing support. Neither version appears to offer support for

electrostatic elements, rendering it unsuitable for applications which require these elements. As a pure matrix code, TRANSPORT does not support the simulation of arbitrary particle distributions, and so cannot simulate particle distributions with complex time domain distributions.

A.7 IMPACT

- Authors: Qiang, Ryne, Habib, and Decyk
- Institution: Lawrence Berkeley National Laboratory
- Type: Particle Tracking Code
- Website: http://blast.lbl.gov/BLASTcodes_Impact.html
- Open Source: No

A.7.1 Overview

IMPACT is a particle tracking code specifically designed to handle high current situations where space charge effects are significant. It is technically a pair of codes: IMPACT-Z, which uses longitudinal position as the independent variable, and IMPACT-T, which uses time. The code is specifically designed with parallel processing in mind, as this sort of highly numerical particle tracking is extremely computationally intensive.

While IMPACT is still in active use, there are limited examples available. A documentation file does exist for IMPACT-T, but is not available on the project website, and there is less information available for IMPACT-Z.

A.7.2 Simulation Method and Capabilities

IMPACT uses direct numerical integration through 3 dimensional grids of fields for all calculations. This method has the potential to be extremely accurate, because the fields of beamline elements are calculated in a more precise way, and space charge can be calculated in far more detail. This method is also extremely computationally intensive, and execution requires a powerful computer or ample time to wait for results.

Starting particle distributions can be either imported from an external file or generated from a wide range of starting configurations. In addition to using arbitrary field distributions supplied from external calculations, IMPACT can also automatically calculate fields for a small library of standard elements such as quadrupoles and solenoids.

A.7.3 Input and Output Files

The basic input file for IMPACT consists of a series of statements defining the conditions of the simulation followed by a list of elements contained in the accelerator lattice. IMPACT may have the least human readable input file of any accelerator code, since every statement in this file consists entirely of strings of numbers. For example, a line defining a quadrupole element might look like:

```
0.30 4 20 1 -5.67 0. 0.014 0. 0. 0. 0. 0.
```

IMPACT will generate a range of output files, in particular either complete particle distributions or summaries of particle distribution properties at selected points on the beamline. The output file names are fixed by the program, and are a mix of descriptive names (for example, *Xprof.data*) and FORTRAN default output files. (such as *fort.8*) The values in the output files are often in normalized units which need to be combined with values in other

files to produce meaningful results.

Sample code fragment:

```

2 1
6 20000 2 0 2
65 65 129 4 0.140000 0.140000 0.1025446
16 0 0 1
20000
0.0
2.6838596289025144e-10

0.1400979E-02 0.1198306E-03 -0.9950372E-01 1.000 1.000 0.000 0.000
0.1400979E-02 0.1198306E-03 -0.9950372E-01 1.000 1.000 0.000 0.000
0.2073764E-01 0.1029707E-05 0.1337864E+00 1.000 1.000 0.000 0.000

0.0 0.6e6 931.49432e6 0.250 80.50e6 0.0

0.0 0 40 -2 1 0.4 /

0.7500000E+00 75 20 0 0.400000E-01 /

0.2239500E+00 22 20 0 0.400000E-01 /
0.2000000E+00 1 20 3 0.200000E+01 0.000000E+00 0.200000E-01 /
0.1230000E+00 12 20 0 0.400000E-01 /
0.2400000E+00 48 20 103 0.200000E+00 0.805000E+08 0.142956E+03 0.1000E+01 0.150000E-01 /
0.1230000E+00 12 20 0 0.400000E-01 /
0.2000000E+00 1 20 3 0.245000E+01 0.000000E+00 0.200000E-01 /
0.2239500E+00 22 20 0 0.400000E-01 /

```

A.7.4 Data Visualization

The IMPACT package has no built in data visualization tools.

A.7.5 Limitations

IMPACT is an extremely accurate and powerful code to address the classes of problems it is suited for. However, it is extremely difficult to use. Field models must be prepared for all but the simplest elements, and the process for converting those models for use by the code is not trivial. Generating lattice files requires adaptation to a purely numeric scripting language, and there is no simple way to visualize the output data once it is produced. Very limited documentation and examples exist for either version of the code.

A.8 Other Codes

There are an enormous number of other accelerator simulation codes available. While only DYNAC, COSY, and TRACK were directly used for this project, a number of other codes were tested and found to be undesirable for one reason or another. Often there was simply insufficient documentation or support for a particular code to be usable without a great deal of effort. Other codes did not include support for low energy electrostatic components. Some simply did provide any additional functionality not represented by the three codes used.

No single code is a “kitchen sink” capable of solving all problems put to it, which is probably the reason there are so many different incomplete codes in existence. Each new code was written to be useful in a specific instance, and no single author has an incentive to extend their code to cover the missing use cases of others. Also, documentation maintenance, let alone user interface design, is time consuming and often not regarded as the best investment of limited time or resources.

For reference, Table A.1 is a table of units used by a number of accelerator codes.

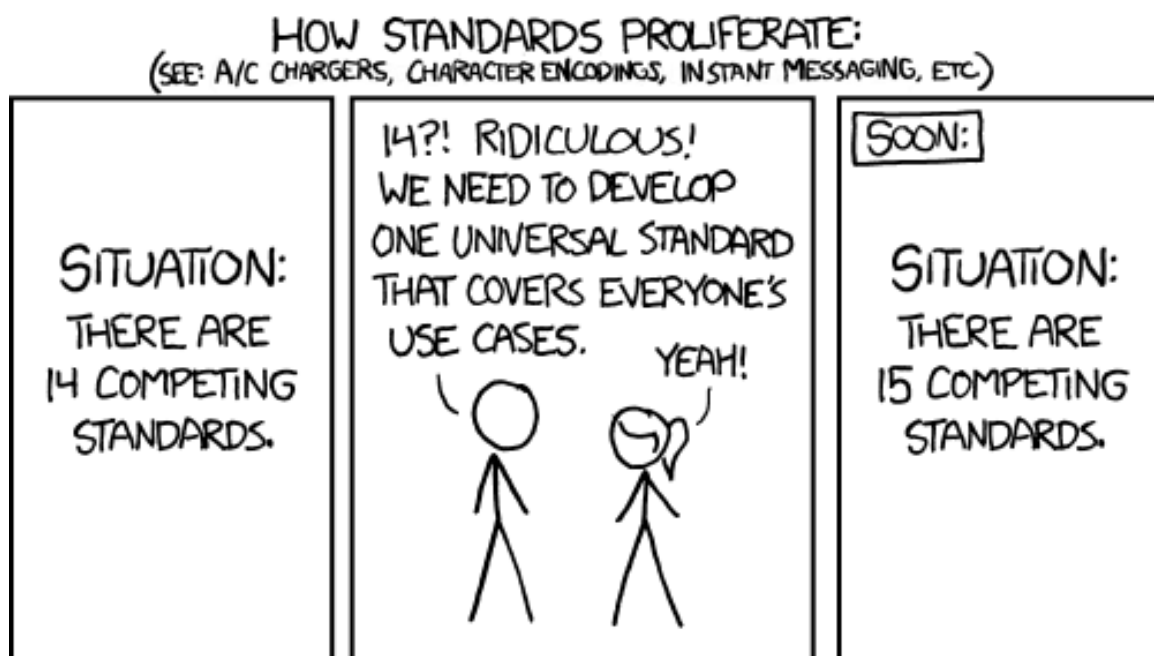


Figure A.7: Image Credit: xkcd.com. Used under Creative Commons Attribution - Non-commercial 2.5 License [2].

	DYNAC	COSY	TRACK	TRACE-3D	IMPACT	MAD	PARMTEQ
Particles							
x / y	cm	m	cm	mm	m	m	cm
x' / y'	rad	rad	rad	mrad	rad	$\frac{\Delta p_x}{p}, \frac{\Delta p_y}{p}$	mrad
z	$\Delta\phi$ (rad)	$-(t - t_o)v_o * \frac{\gamma}{1+\gamma}$	ns	$\Delta\varphi$ (deg)	m	p_z	deg
z'	MeV	$\frac{\Delta K}{K_o}$	$\frac{\Delta\beta}{\beta_o}$	ΔW (keV)	rad	$\frac{\Delta p_z}{p}$	MeV
Mass (output files)	–	$\frac{\Delta m}{m_o}$	amu	–	eV/c ²	GeV/c ²	–
Charge	q	$\frac{\Delta q}{q_o}$	q	q	q	q	–
Beam Parameters							
Alpha (x/y)	none	none	none	none	none	none	none
Beta (x/y)	mm / mrad	m / rad	cm / rad	m / rad	m / rad	m / rad	cm / rad
Emittance (x/y)	mm * mrad (4 RMS)	m * rad (1 RMS)	cm * rad	m * rad (1 RMS)	mm * mrad (RMS)	m * rad (1 RMS)	cm * rad
Alpha (z)	none	none	none	none	none	none	–
Beta (z)	deg / keV	m	deg / ($\Delta W/W$)(%)	deg / keV	deg / %	m / rad	–
Emittance (z)	keV * deg (4 RMS)	m	deg * ($\Delta W/W$)(%)	deg * keV	ns * keV/u (RMS)	m * rad (1 RMS)	–
Other Parameters							
Particle Mass	amu	amu	amu	MeV/c ²	eV/c ²	GeV/c ²	amu
Length	cm	m	cm	mm	m	m	cm
Energy	MeV	MeV	MeV / u	MeV	eV	GeV	MeV
Voltage	kV	kV	V	kV	V	MV	V
B-Field	kG	T	G	G	T	T	G
B-Field Gradient	dimensionless	T/m	–	T/m	T/m	T/m	G / cm
Frequency	Hz / MHz	Hz	Hz	MHz	Hz	MHz	–

Table A.1: Comparison of Units Used in Several Accelerator Codes

Appendix B

DynacGUI

As discussed in Appendix A, DYNAC is a particle tracking code developed originally by Pierre Lapostolle, Salby Valero, and Eugene Tanke, and maintained by Valero and Tanke. While it is an extremely powerful and flexible code, it has a few limitations both in terms of technical capabilities and ease of use. Over the course of this project, DYNACGUI [56] was developed as a graphical front end to the DYNAC code to address some of these limitations. DYNACGUI was written in the MATLAB computing environment, and requires the user to have an installed and licensed copy of that program.

B.1 Basic Dynac Operation

When a basic installation of DYNAC without DYNACGUI is used, the software consists of two executable files, the primary DYNAC executable which performs all of the particle tracking and simulation, and an auxiliary program which handles plotting. In order to run a simulation, an input file is created which specifies the starting conditions of the beam and then the parameters of each element on the beamline sequentially. This input file (sometimes anachronistically referred to as a “deck”) contains all of the information describing the beamline itself. Certain beamline elements, in particular accelerating cavities, solenoids, and RFQs, may refer to external field distribution files. In addition, if the user wishes to specify a starting particle distribution on a particle-by-particle basis, rather than in terms

of CS parameters, a file containing that distribution must be specified as well. The location of each of these supporting files is specified within the input deck.

Once the input deck is prepared, the DYNAC executable is run with the input deck as an argument and, assuming no errors occur, the particles are transported through the simulated beamline. DYNAC then generates a number of output files, of which the following are of interest here:

- *dynac.short* - a short element by element summary of the beamline, giving the basic parameters of each element and the number of particles remaining at the end of that element. If the “EMIT” or “EMITGR” commands are encountered in the input deck, a short summary of beam parameters at that point is also added to this file.
- *dynac.long* - A more detailed element by element listing of beamline objects. This file also reports transport matrices for quadrupoles, dipoles, and solenoids.
- *dynac.print* - A tabular summary of beam parameters with one or more line(s) per beamline element.
- *emit.plot* - This file contains the data for every plot generated by the simulation. Plots are generated each time an appropriate command such as “EMITGR” or “ENVEL” is encountered in the input deck. All plot data are appended into this single file.
- *Output Distribution Files* - Whenever a “WRBEAM” command is encountered in the input deck, DYNAC will write a file containing the complete particle distribution at that point.

Once the file is run, any plots generated can be viewed by running the auxiliary plotting program. This program sequentially displays each plot in the *emit.plot* file, with the option

of saving the plot to a separate graphics file. All other data about the simulation is contained in the text of the output files.

This all-text approach is typical of many accelerator codes, and reflects a desire to expend limited available coding time on improving the underlying model rather than focusing on user interface design. However, in the long run, a more streamlined interface can actually save more time than required for its creation, and has certainly done so in this case. The goals of DYNACGUI are to provide a simple way to quickly alter the parameters of a beamline, enhanced tools for viewing results, both graphical and textual, to facilitate saving and restoring prior simulations, and to provide other utilities that enhance the basic functionality of DYNAC. DYNACGUI is NOT itself an accelerator simulation code, merely a tool to facilitate the use of DYNAC.

B.2 File Structure of DynacGUI

Rather than having one master file containing all of the information for a particular simulation, DYNACGUI divides the information contained in the DYNAC input file into three files: the device file, which contains the physical parameters of the devices used to build the beamline, the layout file, which describes the arrangement of those devices within the accelerator, and the tune settings file, which summarizes the settings of the tunable parameters for each device for a given simulation.

These divisions correspond roughly to the frequency with which this information is expected to change. The physical dimensions of a magnet or cavity are unlikely to be altered once they are built. The layout of a beamline may change occasionally, but for the most part is also likely to remain static. On the other hand, the settings for the individual devices are

likely to change for each new simulation run. When DYNACGUI is used, it combines these three files to synthesize a valid DYNAC input deck containing all of these parameters.

Auxiliary device files, such as field distributions, must still be specified as part of the device definition in the devices file. External files for starting particle distributions can also be identified.

Since DYNACGUI does not change the actual operation of DYNAC in any way, once the input deck is constructed, and DYNAC is run, the same output files as described above are produced.

B.3 Using DynacGUI

Once the input files have been prepared, DYNACGUI is run from within MATLAB and displays a control panel to the user which allows them to select which layout, device, and tune files they wish to combine (Figure B.1). The option is also present to select an input particle distribution file, although the CS parameters will be used by default.

The user generates the input file, which can be viewed directly from within DYNACGUI to check for errors, and then selects “Run Deck” to call the DYNAC executable to process the input file. Once the program completes execution, the DYNAC’s execution return message is displayed in the output pane, and all of the graphs in the *emit.plot* file are summarized in the “Generated Graphs” pane.

B.3.1 Data Visualization

The graphical display of generated data is the first area in which DynacGUI substantially improves on the usability of bare DYNAC. Regardless of the specific plots selected within the

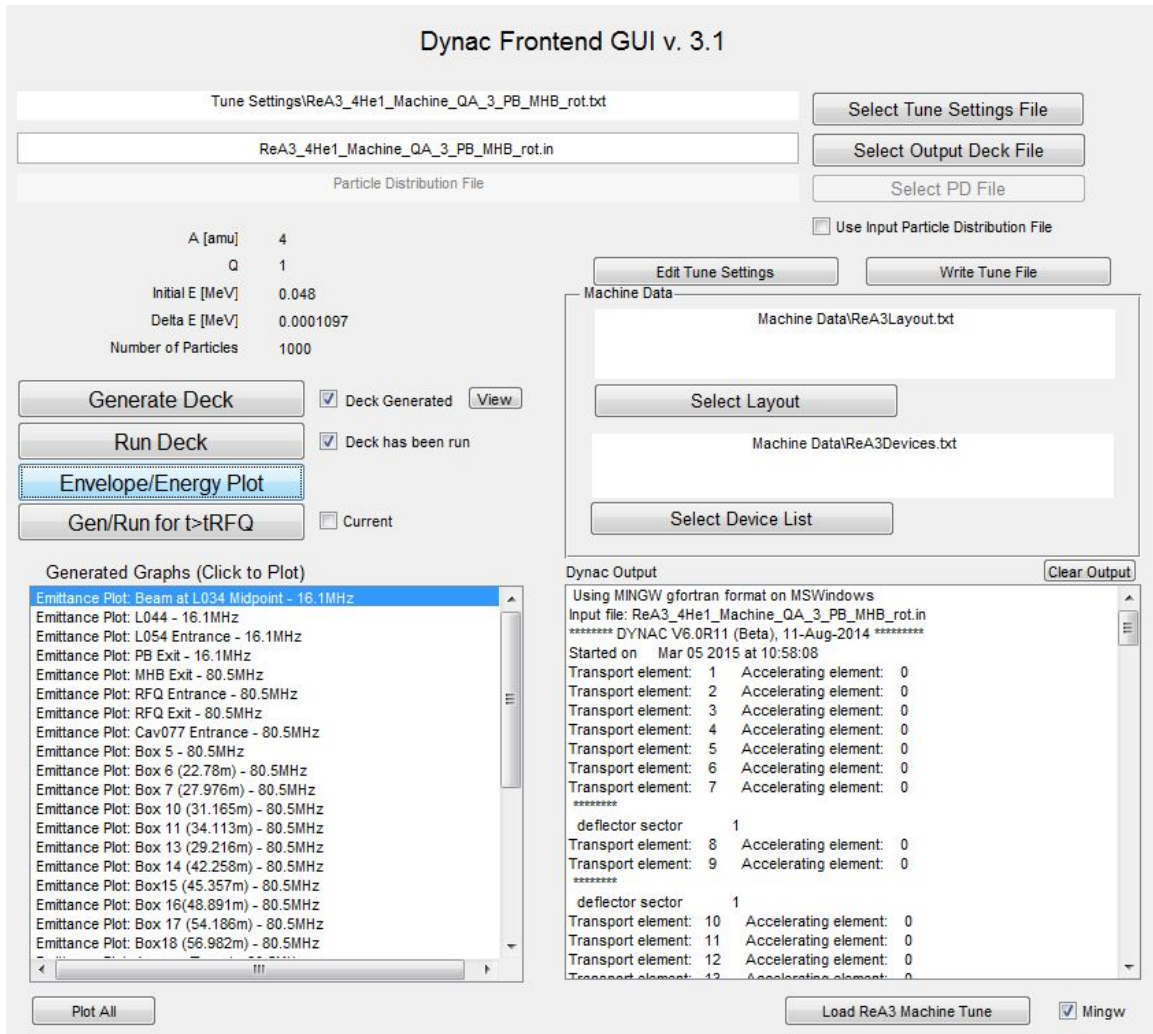


Figure B.1: DYNACGUI Main Control Panel




Figure B.2: DYNACGUI Main Z-Axis Plot

input file, the “Envelope/Energy Plot” button is always available to produce a plot of the beam parameters in the *dynac.print* file relative to the position on the z axis (Figure B.2).

These parameters include:

- x and y RMS beam profiles and envelopes
- Longitudinal emittance
- Reference particle energy
- Number of particles remaining
- x and y beta functions
- Time and energy spread
- x and y dispersion functions

The locations of the generated emittance graphs (see below) can also be overlaid on the plot, and the plot narrowed to only show specific portions of the beamline.

For emittance plots, DYNACGUI replaces the plotting routine supplied with DYNAC with a routine built around MATLAB’s internal plotting functions. Rather than requiring users to cycle through each generated plot, a list of all plots is generated at the time the simulation is run, and the user may select an individual plot from the list.

When an emittance plot (Figure B.3) is selected, DYNACGUI generates the standard four-panel layout of graphs - x vs. x' , y vs. y' , x vs. y , and time vs. energy. In addition, histograms and RMS widths are shown for the axes of the last two graphs. While in this view, a “Tools” menu is also available to provide other functionalities to the user. Tools include the ability to display beam data at this point, save this particle distribution to an

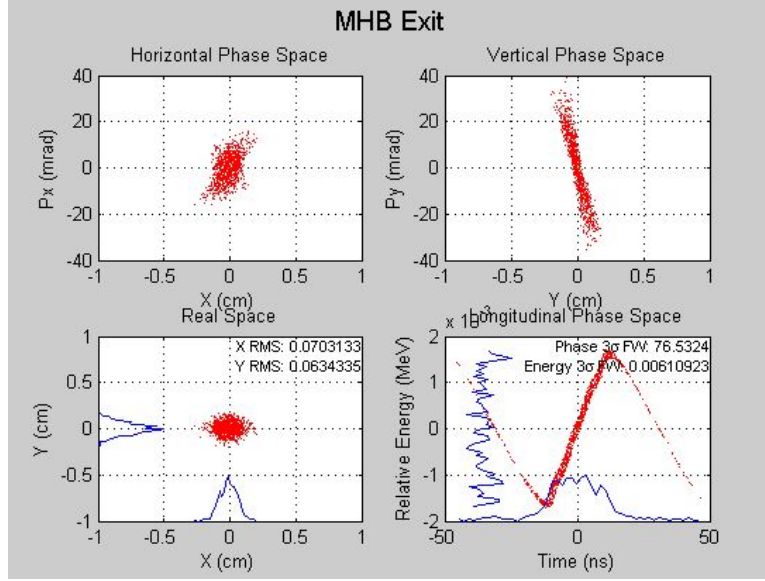


Figure B.3: DYNACGUI Emittance Plot

external file in DYNAC, COSY¹, or TRACK format, and display a group of less frequently used plots, such as x vs. z or dp/p vs. x . In particular, the ability to quickly save a particle distribution file at a given point which can then be reimported into DYNACGUI as a starting distribution is quite useful.

A further advantage of using MATLAB's plotting tools is that in addition to the tools explicitly defined in DYNACGUI, MATLAB's own plot editing tools are still available. These can be used for tasks such as manual scaling of axes and saving images in various formats.

B.3.2 Editing, Loading, and Saving Information

In addition to the ability to save and load particle distributions described so far, DYNACGUI has several other ways to modify and save data. First, using the “Edit Tune Settings” button, the user can alter any of the tunable parameters located in the tune settings file

¹COSY is not natively a particle tracking code, so there is no set format for a particle distribution file for that code. DYNACGUI exports to the COSY particle distribution file format developed by Mauricio Portillo at NSCL.

from within DYNACGUI. This allows the user to quickly alter one or more parameters and rerun the simulation without manually editing text files. Tunes so altered can also be saved to new tune settings files for future use without altering the existing file.

In addition, DYNACGUI can save the results of an entire run, including all plots, input and output files, and beam information to a specified directory, where it can be viewed later without re-running the entire simulation.

B.3.3 Fitting

One major drawback of DYNAC is its lack of any built in capability for automated fitting of beamline parameters. This capability, present in a number of other accelerator codes, is extremely useful when the beamline must be tuned for a desired condition at a certain point. DYNACGUI compensates for this absence by using the capabilities of MATLAB's Optimization Toolbox.

After generating an input deck, the user selects the "Fitting Tool" from the "Tools" menu within DYNACGUI (Figure B.4). They are presented with a list of every parameter within the deck, including non-tunable parameters such as element dimensions and drift lengths. The user selects which parameters are to be used as independent variables, and optionally may select a point on the beamline at which a desired fitting condition is to be met. (The end of the beamline is used by default.) Finally, a desired final dependent variable, such as beam size or final energy, is selected for optimization. The fitting tool can attempt to maximize or minimize the dependent variable, or fit it to a desired value. Instructions are also included for setting up customized combinations of dependent variables, although this requires the user to edit the DYNACGUI source code themselves.

Once the conditions are set up, the MATLAB optimizer is called. For each iteration of the

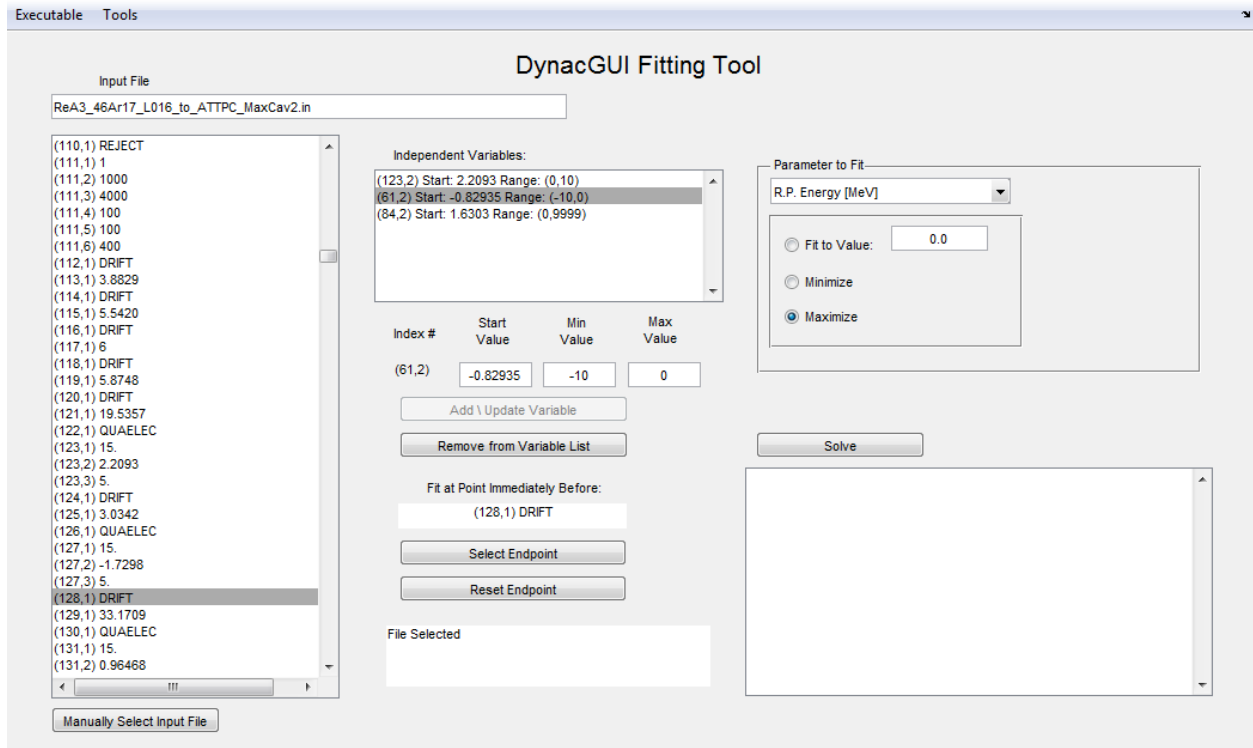


Figure B.4: DYNACGUI's fitting tool screen.

optimizer, DYNACGUI generates a new deck file based on the trial values of the independent variable(s), runs DYNAC on the deck file, and returns the value of the dependent variable at the selected point. From the point of view of the optimizer, Dynfunc is simply a black box - input values are fed in and output values are returned. This allows standard optimization methods to be used to tune the simulation for a desired condition. One drawback of this approach is that since it is entirely based on reading and writing DYNAC input and output files, it is inherently slower than a corresponding fitting strategy which does not use repeated disk access.

B.3.4 Other Capabilities

As discussed in Section 5.7, one major challenge in simulating ReA from end to end is that none of the codes available, including DYNAC, will simulate a distribution of particles in an

RFQ longer than one RF period. DYNACGUI incorporates a set of routines to overcome this limitation. When the “Gen/Run for $t > \text{RFQ}$ ” option is selected, the particle distribution is flown from the start of the simulation to the entrance of the RFQ. At that point, the distribution is saved to a file, and DYNACGUI splits that file into a series of smaller particle distribution files, each containing only one RF period of the original distribution at the entrance.

Those distributions are flown one at a time through the simulated RFQ, and the resulting output distributions rejoin into a single large distribution which is then transported through the remainder of the simulated beamline. Most data is retained through this process, so the emittance and envelope plots still function properly, although running totals within the RFQ itself may not plot properly.

In addition to this capability, DYNACGUI can generate a scaled version of a tune for a different mass to charge ratio or energy setting. This feature is useful when rescaling a beamline for a species with a different mass to charge ratio. It should still be used with caution, however, as DYNACGUI cannot determine ahead of time what the actual energy of the particle will be at any point of the beamline. As such, any scaling that takes place in the middle of an accelerating section will necessarily be an approximation.

Finally, DYNACGUI has the ability to export a rudimentary version of a COSY deck for a given DYNAC simulation. This capability is extremely limited, as some elements in DYNAC are difficult to represent equivalently in COSY. Further, since COSY is a ray-tracing code and DYNAC is a particle tracking code, it is difficult to ensure a one to one comparison between the initial conditions present in both simulations.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] D. Edwards and M. Syphers. *An Introduction to the Physics of High Energy Accelerators*. Wiley-VCH Verlag GmbH & Co., Weinheim, Germany, 1993.
- [2] Randall Munroe. xkcd.com, xkcd927 - Standards. <http://imgs.xkcd.com/comics/standards.png>, July 2011, cited 2014-12-02.
- [3] H. Koivisto, D. Cole, A. Fredell, C. Lyneis, P. Miller, J. Moskalik, B. Nurnberger, J. Ottarson, A. Zeller, J. DeKamp, P.A. Vondrasek, P. A. Zavodszky, and Z.Q. Xie. Artemis - The New ECR Ion Source for the Coupled Cyclotron Facility at NSCL/MSU. In *Proceedings of the Workshop on the Production of Intense Beams of Highly Charged Ions*, volume 72, page 135, Catania, Italy, 24-27 Sep. 2000. Italian Physical Society Conference Proceedings.
- [4] P. A. Zavodszky, B. Arend, D. Cole, J. DeKamp, G. Machicoane, F. Marti, P. Miller, J. Moskalik, J. Ottarson, J. Vincent, and A. Zeller. Design of SuSI — Superconducting Source for Ions at NSCL/MSU — I. The Magnet System. In *AIP Conference Proceedings*, volume 749, pages 131–134, Berkeley, CA, 2005. AIP.
- [5] P. A. Zavodszky, B. Arend, D. Cole, J. DeKamp, G. Machicoane, F. Marti, P. Miller, J. Moskalik, J. Ottarson, J. Vincent, and A. Zeller. Design of SuSI – Superconducting Source for Ions at NSCL/MSU – II. The Conventional Parts. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 241(1–4):959–964, December 2005.
- [6] H.G. Blosser. Application of Superconductivity in Cyclotron Construction. In *Proceedings of the 9th International Conference on Cyclotrons and their Applications*, pages 147–157, Caen, France, September 1981.
- [7] P. Miller, F. Marti, D. Poe, M. Steiner, J. Stetson, and X.Y. Wu. Commissioning of the Coupled Cyclotron Facility at NSCL. In *Particle Accelerator Conference, 2001. PAC 2001. Proceedings of the 2001*, volume 4, pages 2557–2559 vol.4, Chicago, IL, USA, 2001.
- [8] X. Wu, H.G. Blosser, D. Johnson, F. Marti, and R. C. York. The K500-to-K1200 Coupling Line for the Coupled Cyclotron Facility at the NSCL. In *Proceedings of the 1999 Particle Accelerator Conference*, pages 1318–1320, New York, NY, USA, 1999.

- [9] Available Beams — National Superconducting Cyclotron Laboratory — Michigan State University. <http://nscl.msu.edu/users/beams.html>, cited 2014-12-16.
- [10] D. J. Morrissey. Planning for a Standard in-Flight RNB Experiment. https://groups.nscl.msu.edu/a1900/experimentplanning/a1900_flowCa.pdf, July 2003, cited 2014-12-16.
- [11] X. Wu, D. J. Morrissey, B. M. Sherrill, R. C. York, and A. F. Zeller. Tracking studies and performance simulations of the NSCL A1900 fragment separator. In *Proceedings of the Particle Accelerator Conference, 1997*, volume 1, pages 198–200, Vancouver, CA, 1997. IEEE.
- [12] D. J. Morrissey, B. M. Sherrill, M. Steiner, A. Stolz, and I. Wiedenhoever. Commissioning the A1900 projectile fragment separator. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 204:90–96, May 2003.
- [13] L. Weissman, P. A. Lofy, D. A. Davies, D. J. Morrissey, P. Schury, S. Schwarz, T. Sun, and G. Bollen. First extraction tests of the NSCL gas cell. *Nuclear Physics A*, 746:655–658, December 2004.
- [14] G. K. Pang, G. Bollen, S. Chouhan, C. Guénaut, F. Marti, D. J. Morrissey, D. Lawton, J. Ottarson, S. Schwarz, A. F. Zeller, and others. The cyclotron gas stopper project at the NSCL. In *Particle Accelerator Conference, 2007. PAC. IEEE*, pages 3588–3590. IEEE, 2007.
- [15] S. Schwarz, G. Bollen, J. R. Crespo López-Urrutia, O. Kester, K. Kittimanapun, A. Lapiere, D. Leitner, J. Ottarson, and M. Portillo. Initial commissioning results with the NSCL Electron Beam Ion Trap. *Review of Scientific Instruments*, 83(2):02A908, 2012.
- [16] Thomas M. Baumann, Alain Lapiere, Stefan Schwarz, Kritsada Kittimanapun, and Georg Bollen. Energy spread and time structure of ion beams extracted from the ReA-EBIT rare isotope charge breeder. In *Proceedings of the XII International Symposium on Electron Beam Ion Sources and Traps*, volume 1640, pages 80–87. AIP Publishing, January 2015.
- [17] O. Kester, D. Bazin, C. Benatti, J. Bierwagen, G. Bollen, S. Bricker, A. C. Crawford, S. Chouhan, C. Compton, K. Davidson, and others. Rea3—the Rare Isotope Reaccelerator at MSU. *LINAC2010, Tsukuba, Japan*, 2010.

- [18] M. Portillo, G. Bollen, S. Chouhan, O. Kester, G. Machicoane, J. Ottarson, S. Schwarz, A. Zeller, and others. An achromatic mass separator design for ions from the EBIT charge breeder at the NSCL. In *Proceedings of the 2009 Particle Accelerator Conference*, pages 4341–4343, Vancouver, CA, 2009.
- [19] Q. Zhao, V. Andreev, F. Marti, S. O. Schriber, X. Wu, and R. C. York. Design Studies of the Reaccelerator RFQ at NSCL. In *Proceedings of the Particle Accelerator Conference, 2007*, pages 1772–1774, Albuquerque, NM, 2007. IEEE.
- [20] R. C. York, X. Wu, Q. Zhao, M. Doleans, F. Marti, and E. Pozdeyev. Beam Dynamics in the FRIB Linac. In *Proceedings of HB 2010*, pages 319–323, 2010.
- [21] D. Bazin. Summary of meeting on time structure of beams for ISLA, November 2009.
- [22] D. Bazin and W. Mittag. ISLA: An Isochronous Spectrometer with Large Acceptances. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 317:319–322, 2013.
- [23] M. Berz and K. Makino. COSY INFINITY 9.0 Beam Physics Manual. Technical Report MSUHEP 060804, Michigan State University, August 2006.
- [24] Pierre M. Lapostolle and M. Weiss. Formulae and procedures useful for the design of linear accelerators. Technical Report CERN-PS-2000-001 (DR), CERN, Geneva, Switzerland, January 2000.
- [25] K. R. Crandall and D.P. Rusthoi. TRACE3D.PDF. Technical Report LA-UR-97-886, LANL, Los Alamos, NM, May 1997.
- [26] Mario Conte and William McKay. *An Introduction to the Physics of Particle Accelerators*. World Scientific, Singapore, 2008.
- [27] M. Leitner, B. Bird, S. Casagrande, S. Chouhan, C. Compton, J. Crisp, K. Elliot, A. Facco, A. Fox, M. Hodek, M. Johnson, I. Kiupel, D. Malloch, S. Miller, D. Morris, D. Norton, R. Oweiss, J. Ozelis, J. Popielarski, L. Popielarski, A. Rauch, R. Rose, K. Saito, M. Shuptar, N. Usher, D. Velianoff, D. Victory, J. Wei, J. Whitaker, K. Witgen, T. Xu, Y. Xu, O. Yair, and S. Zhao. The FRIB Project at MSU. In *Proceedings of SRF 2013*, pages 1–10, Paris, France, 2013.
- [28] F.J. Lynch, R.N. Lewis, L.M. Bollinger, W. Henning, and O.D. Despe. Beam buncher for heavy ions. *Nuclear Instruments and Methods*, 159(2–3):245–263, February 1979.

- [29] K. R. Crandall. RFQ radial matching sections and fringe fields. In *Proceedings of the 1984 Linac Conference, GSI-84-11*, 1984.
- [30] John Staples. Reducing RFQ Output Emittance by External Bunching. *Particle Accelerators*, 47:101, 1994.
- [31] Z. Kohley. Experimental difficulties with the time structure of the ReA beams, 2014.
- [32] Behlke FSWP 91-01 Specification Sheet. http://http://www.behlke.com/pdf/fswp_91-01_rs.pdf, cited 2015-02-12.
- [33] Mauro Paoluzzi. A fast 650V chopper driver. Technical Report CERN-ATS-2011-026, CERN, Geneva, Switzerland, August 2011.
- [34] P. N. Ostroumov, J. D. Fuerst, M. P. Kelly, B. Mustapha, K. W. Shepard, and J. Xu. A new ATLAS efficiency and intensity upgrade project. In *Proceedings of the 14th International Conference on RF Superconductivity, Berlin, Germany*, Berlin, Germany, 2009.
- [35] J.W. Gibbs. Fourier’s Series. *Nature*, 59(1522):200, December 1898.
- [36] Fredric J. Harris. On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.
- [37] R. C. Pardo, J. M. Bogaty, and B. E. Clift. A study of beam chopping options for the ATLAS Positive Ion Linac. In *LINAC 1996*, pages 65–67, 1996.
- [38] P. N. Ostroumov, Z.A. Conway, C. Dickerson, S. Gerbick, M. Kedzie, M. P. Kelly, S. H. Kim, Y. Luo, S. MacDonald, R. Murphy, B. Mustapha, R. C. Pardo, A. Perry, T. Reid, S. I. Sharamentov, K. W. Shepard, J. Specht, and G. Zinkann. Completion of Efficiency and Intensity Upgrade of the ATLAS Facility. In *Proceedings of LINAC 2014*, pages 449–451, Geneva, Switzerland, 2014. JACoW.
- [39] P.N. Ostroumov, V.N. Aseev, and B. Mustapha. TRACKv37.doc, March 2007, cited 2015-02-26.
- [40] D. Alt, S. Krause, A. LaPierre, D. Leitner, S. Nash, R. Rencsok, J. A. Rodriguez, M. Syphers, and W. Wittmer. Measurement of the Longitudinal Acceptance of the ReA RFQ. In *Proceedings of IPAC 2014*, pages 3346–3348, Dresden, Germany, 2014. IPAC.

- [41] E. Tanke, P. LaPostolle, and S. Valero. Program Dynac User Guide, September 2014, cited 2015-02-25.
- [42] R Core Team. R: A Language and Environment for Statistical Computing. Technical report, R Foundataion for Statistical Computing, Vienna, Austria, 2014.
- [43] J. Holzbauer, W. Hartung, F. Marti, E. Pozdeyev, and Q. Zhao. Electromagnetic Design of a Multi-Harmonic Buncher for the FRIB Driver Linac. In *Proceedings of the 2011 Particle Accelerator Conference*, pages 1000–1002, New York, NY, USA, 2011.
- [44] K. Kasemir and G. Carcassi. Control System Studio Guide. <http://cs-studio.sourceforge.net/docbook/>, October 2015, cited 2016-02-18.
- [45] V. A. Verzilov, R. E. Laxdal, M. Marchetto, and W. R. Rawnsley. Time Domain Diagnostics for the ISACII Superconducting Heavy Ion Linac. *DIPAC, Venice, Italy*, 2007.
- [46] R.J. McDonald, D.A. Landis, and G.J. Wozniak. Fast Dual Scale-Down Module (LBL#21x7981P-1). Technical Report LBL-21930, Lawrence Berkeley Laboratory, Berkeley, CA, August 1986.
- [47] ORTEC. MAESTRO MCA Emulator for Microsoft Windows XP Professional SP3 and Windows 7 x32 and x64, 2012.
- [48] Tom O’Haver. Curve fitting C. Non-linear Iterative Curve Fitting. <https://terpconnect.umd.edu/~toh/spectrum/CurveFittingC.html>, December 2015, cited 2015-12-09.
- [49] Kesar Singh and Minge Xie. Bootstrap: a statistical method. <http://www.stat.rutgers.edu/home/mxie/stat586/handout/Bootstrap1.pdf>, 2008, cited 2016-03-11.
- [50] NSCL. ReA3 Service Level Description. <http://nscl.msu.edu/users/ReA3-SLD.pdf>, 2015, cited 2016-02-20.
- [51] Shin, Ki. Double-Gap Rebuncher Cavity Design of SNS MEBT. In *IPAC 2012*, New Orleans, LA, USA, 2012. IEEE.
- [52] Friedhelm Caspers, T. Kroyer, and M. Paoluzzi. CERN Chopper final report. Technical Report CERN-AB-Note-2008-040 RF, CERN, 2008.
- [53] T. Kalvas. Beam Extraction and Transport. *arXiv preprint arXiv:1401.3951*, 2014.

- [54] Kenneth R. Crandall and Thomas P. Wangler. PARMTEQ—A beam-dynamics code for the RFQ linear accelerator. In *Proceedings of AIP 1988*, volume 177, pages 22–28, San Diego, CA, December 1988.
- [55] CST GmbH. CST Microwave Studio 2015 - Workflow and Solver Overview, 2015.
- [56] D. Alt. DynacGUI User Manual. <https://github.com/NSCLAlt/DynacGUI/blob/master/DynacGUI%20Manual.docx>, January 2016, cited 2016-02-22.