

A Control System for the E-Linac View Screen System

by

Jason Matthias Abernathy
B.Sc., University of Victoria, 2010

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

MASTER OF SCIENCE

in the Department of Physics and Astronomy

© Jason M. Abernathy, 2015
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by photocopying or other means, without the permission of the author.

A Control System for the E-Linac View Screen System

by

Jason Matthias Abernathy
B.Sc., University of Victoria, 2010

Supervisory Committee

Dr. D. Karlen, Supervisor
(Department of Physics and Astronomy)

Dr. J. Albert, Departmental Member
(Department of Physics and Astronomy)

Supervisory Committee

Dr. D. Karlen, Supervisor
(Department of Physics and Astronomy)

Dr. J. Albert, Departmental Member
(Department of Physics and Astronomy)

ABSTRACT

The TRIUMF view screen system encompasses a set of devices which individually image, and produce measurements of, the transverse profile of an accelerated electron beam. A control system is an essential component of the overall diagnostic device. The system requirements were compiled from those produced by the TRIUMF laboratory and from those based on the needs of the individual diagnostic devices. Based on the requirements, a control system was designed and implemented with a combination of industrial electrical and mechanical hardware, and a variety of software components. One component of the image reconstruction algorithm was validated with experimental data; the accuracy and precision of beam profile measurements was evaluated through simulation studies. Although it was not possible to demonstrate the satisfaction of requirements relating to alignment, it was shown that all other requirements were satisfied.

Table of Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
List of Tables	vi
List of Figures	vii
Acknowledgements	xiii
Dedication	xiv
1 Introduction	1
2 ARIEL, e-Linac and the View Screen System	3
2.1 ARIEL	3
2.1.1 Rare Isotope Beam Production	3
2.1.2 The e-Linac	4
2.2 The View Screen System	5
2.3 Control System Requirements	6
2.3.1 External Requirements	7
2.3.2 Internal Requirements	9
3 A Control System for the e-Linac View Screen System	11
3.1 High-Level System Design	11
3.2 Service Implementation	12
3.2.1 Iris Control	13
3.2.2 Camera Control and Triggering	14

3.2.3	Computing	16
3.2.4	Electric Power Distribution	18
3.2.5	Lighting	19
3.2.6	Environmental Monitoring	19
3.3	Service Enclosures	20
3.4	Cabling	21
3.5	Image Processing	23
3.5.1	Coordinate Systems	24
3.5.2	Geometry Correction	26
3.5.3	Magnification Correction	37
3.5.4	Collection Efficiency Correction	39
3.5.5	Calculation of Beam Profile Statistics	43
3.5.6	Image Styling	45
4	Experiments	48
4.1	Collection Efficiency Validation	48
4.1.1	Data Collection	48
4.1.2	Analysis	49
4.2	Evaluation of the Quality of Beam Profile Measurements	56
4.2.1	A Measurement of the Systematic Bias Present in Ideal Optical Configurations	56
4.2.2	Improving the Quality of Beam Profile Measurements by Accounting for the Systematic Bias	62
4.2.3	Precision of the Systematic Bias Corrections	65
4.2.4	Satisfaction of Beam Profile Measurement Requirements	68
5	Conclusions	74
	Bibliography	76
	A Supplementary Information	77
A.1	Optical Simulation	77
A.2	External Software Dependencies	79
A.3	Data	80
	Acronyms	82

List of Tables

Table 2.1	Host Computer Hardware Components	9
Table 3.1	Host Computer Hardware Components	18
Table 3.2	Additional View Screen Computing Resources	18
Table 4.1	Beam Parameters and Imaging Conditions for the Ideal Electron Low energy Beam Transport (ELBT) Optical Configuration. The beam centroid was constrained to fall within the boundary of the target foil: -17.68 to 17.68 mm in the x-dimension and -25 to 25 mm in the y-dimension.	57
Table 4.2	Beam Parameters and Imaging Conditions for the Ideal Elec- tron High energy Beam Transport (EHBT) Optical Configura- tion. The beam centroid was constrained to fall within the bound- ary of the target foil: -8.84 to 8.84 mm in the x-dimension and -12.5 to 12.5 mm in the y-dimension. The range of the iris diam- eter was increased to include 36 mm because the diameter of the iris is expected to be larger to collect more light when imaging targets in the higher energy sections of the beamline.	57
Table 4.3	Function Diagnostic Requirements of the View Screen System.	68
Table A.1	Experimental Physics Industrial Control System (EPICS) Soft- ware Dependencies	80
Table A.2	Non-EPICS Software Dependencies	81
Table A.3	Cabinet Temperature Transmittor Calibration Data	81

List of Figures

Figure 2.1	Components of a view screen unit (a) and of a camera box (b).	6
Figure 3.1	System Level Diagram of Control System Interfaces. A majority of the necessary interfaces are implemented as services by the view screen control system. Hardware components are housed in a service enclosure (Sec. 3.3). Software components for controlling equipment and processing data are executed by the computing services.	12
Figure 3.2	Iris Coordinate System. The iris was modelled as a perfectly circular aperture of diameter d . The variable s is the number of <i>steps</i> from a reference position (in this case - the reverse limit switch).	14
Figure 3.3	Iris Position Conversion Fit Residuals. Fitting a one-piece exponential form ($d = ae^{bs} + c$) to the iris measurements yielded large residuals (blue). Fitting a two-piece exponential to the data left smaller residuals (yellow).	15
Figure 3.4	Unsynchronized vs Synchronized Camera Trigger. The amount of light collected by the camera (overlapping purple area) can vary greatly when the frame acquisition state (blue) is not synchronized with the macro-pulse (red).	16
Figure 3.5	Cabinet Lighting Circuit. One of two identical cabinet lighting circuits. The intensity of the lighting stations is adjusted by a PWM dimmer. The dimmer accepts a 0-10 V control voltage which is supplied by a 8-bit DAC. The stations are individually activated with solid-state relays.	20
Figure 3.6	Cabinet Temperature Calibration.	20

Figure 3.7 A Drawing of the View Screen Cabinet. This first-generation enclosure holds every piece of control system hardware required to service sixteen view screen units. The components are coloured by the service supplied: computing (blue), power supply and distribution (red), iris control (green), lighting (purple) and connectivity (orange).	22
Figure 3.8 A Drawing of the Second Generation View Screen Service Enclosure. This second-generation enclosure holds the motor controllers, a signal processing board, environmental sensors and can service up to eight view screen units.	23
Figure 3.9 Beam-space Coordinate System. The beam-space plane (dashed area) is transverse to the nominal beam direction. The origin of the beam-space coordinate system is approximately the centre of the beam tube. Note that the x-axis is reversed due to the orientation of the local beam coordinate system.	25
Figure 3.10 Beam-space Region of Interest. The extents of the beam-space region of interest are set by an operator during device calibration. For a right-handed coordinate system x_l will be greater than x_r . Note that the extents of beam-space are aligned with pixel boundaries, not pixel centres.	26
Figure 3.11 Line segments from the beam-space coordinate system (left) experience a flip, rotation and magnification upon transformation to the raw-image coordinate system (right). These distortions must be removed before properties of the beam profile can be measured.	27
Figure 3.12 A custom port was machined into the front face of all ELBT diagnostic boxes, facilitating calibration target illumination. . .	28
Figure 3.13 Two images of the same calibration target illuminated from a different source. The image of the front-lit calibration target (a) is over-exposed at the top and under-exposed at the bottom. Some control points are nearly indistinguishable from the background. It is easier to distinguish control points from the background when the calibration target is back-lit (b).	29

Figure 3.14	Cutaway Side-view of the Calibration Target. The holes in the calibration target (shown at the left) are chamfered, allowing light from the rear of the target to pass through the hole unimpeded. This is necessary because the targets are rotated 45 degrees about the vertical axis.	29
Figure 3.15	Simulated Images of the Centre Hole in an ELBT Calibration Target. When the fully illuminated hole (a) is partially obscured by the chamfer edge, some light does not reach the camera. Left unaccounted for, this would introduce an offset (b) into all control point locations.	30
Figure 3.16	Three Sub-components of a Calibration Target Hole. The features used by the image recognition algorithm must be described as a combination of the hole, edge and target-face images. . . .	31
Figure 3.17	Operator Description of the Calibration Target Hole Features. An operator describes the feature by selecting values for α, β_1, β_2 and γ	32
Figure 3.18	Bilinear Interpolating Pixel Intensities. One method of calculating pixel intensities is interpolating the value between four neighbouring pixels.	34
Figure 3.19	Geometric Correction Table. The table entry at row v , column u contains the information to construct pixel I'_v . Each two-component column contains the coefficient $c_{v\mu}$ (green) and array index (blue) of an input pixel which contributes to I'_v . A maximum count of eight input pixels may produce a single output pixel.	35
Figure 3.20	ELBT Pixel Area Density. The side of the target which is closer to the view screen camera box (positive x-values, in this case) experiences more magnification than the other side. An object which is on the closer side will fill more pixels on the Charge-Coupled Device (CCD) sensor, appearing dimmer.	37
Figure 3.21	Not all of the light emitted by the target foil (right) reaches the camera's CCD sensor (upper-left).	39

Figure 3.22 Due to the angle of the target foil, geometry of the ELBT camera box and shape of the scintillation light emission distribution only a small fraction of light emitted by the target foil will be collected by the camera. The roughly linear x -dependence and y -independence is caused by one side of the foil being closer to the optical system than the other. This is discussed further in Section 4.1.2. 40

Figure 3.23 The normalized collection efficiency table E for scintillation light passing through the ELBT (low-energy) optical configuration with an iris setting of 16 mm. 44

Figure 3.24 Demonstration of a poorly designed colour map. An operator may perceive a series of flat bands of beam intensity when, in reality, the gradient of the test image changes slowly. 46

Figure 3.25 Four new colour maps. 46

Figure 3.26 Beam Styling Overlays. These features present additional information to the beamline operator. 47

Figure 4.1 Comparing the Measured Collection Efficiency to Simulated Data. The ratio $\frac{I_{i,\nu}}{I_{1,\nu}}$ is a good approximation to the ratio of collection efficiency present during the acquisition of each image. 51

Figure 4.2 The change in collection efficiency in lines across beam space. 52

Figure 4.3 Estimating the Change in Collection Efficiency using Geometry. 54

Figure 4.4 Simulated and Measured Change in Collection Efficiency. 55

Figure 4.5 Systematic bias in the beam centroid measurement of a Gaussian shaped beam incident on a scintillating target imaged by an ELBT camera box. The vector difference between a measured beam centroid and the mean position of the corresponding simulated Gaussian beam was visualized using an arrow originating at the mean position of the simulated beam, pointing in the direction of the vector difference and having a magnitude proportional to the magnitude of the difference. The results of two different iris diameters and beam widths are shown. 59

Figure 4.6	Relative change between the horizontal width ($\sqrt{U_{20}}$) of the reconstructed beam profile and the true width (σ) of a Gaussian beam intercepted by a scintillating target foil and imaged with an ELBT optical configuration.	61
Figure 4.7	Relative change between the horizontal width ($\sqrt{U_{20}}$) of the reconstructed beam profile and the true width (σ) of a Gaussian beam intercepted by a scintillating target foil and imaged with an EHBT optical configuration.	62
Figure 4.8	Systematic bias in the beam centroid and relative change in beam width after the effect of Point Spread Function (PSF) blurring is removed. This was calculated with the simulation of Gaussian shaped beams incident on a scintillating target imaged by an ELBT optical configuration.	64
Figure 4.9	Camera Focusing by Minimizing the Entropy. The location corresponding to the minimum entropy in the image of a calibration target was found by fitting a quadratic curve to the simulated data.	66
Figure 4.10	Uncertainty in the Measurement of the PSF X-Centroid. The x-centroid of the PSF shifts as the camera moves during the focusing procedure (4.10a). Uncertainty in the final camera position leads to uncertainty in the x-centroid bias correction (4.10b).	67
Figure 4.11	Centroid Measurement Accuracy for the ELBT and EHBT Optical Configurations. The required accuracy of the beam centroid measurement is satisfied at the inner portion of the target for the expected beam sizes. The requirement is not satisfied (red area) when the beam profile is clipped at the edge of the target. The results are shown for simulated Gaussian-shaped beams.	70
Figure 4.12	Maximum Precision of the Centroid Measurement for the ELBT and EHBT Optical Configurations. These results include the uncertainty in the camera focusing procedure but not the uncertainty in light emission due to target surface quality. The beam centroid measurement contains two components: M_{10} and M_{01} ; for that reason, the larger of $\sigma_{M_{10}}$ and $\sigma_{M_{01}}$ is shown.	72

Figure 4.13	Maximum Precision of the Centroid Measurement for the ELBT and EHBT Optical Configurations. These results include the uncertainty in the camera focusing procedure but not the uncertainty in light emission due to target surface quality. The beam centroid measurement contains two components: M_{10} and M_{01} ; for that reason, the larger of $\sigma_{M_{10}}$ and $\sigma_{M_{01}}$ is shown.	73
Figure 5.1	Electron Beam Profile Cathode Grid Pattern.	75
Figure A.1	An $N \times M$ grid of equally spaced points in probability space (a) are used to generate a set of photons following the Lambertian distribution (b).	79

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to:

Dr. Dean Karlen for your patience, guidance, and motivation in the process of conducting research and writing my thesis.

D. Storey for your excellent thesis work, which provided the foundation for my own.

the employees of the UVic TRIUMF laboratory and electronics shop for your fundamental and tireless contributions.

my family for your constant source of support and inspiration.

DEDICATION

To my family: Robbin, Hannah, Felicity, Jocelyn and Lorelai.

Chapter 1

Introduction

The topic of this thesis is the development of a control system for the e-Linac view screen system: a diagnostic device which images, in real-time, the transverse profile of an accelerated electron beam. The control system is the implementation of a set of hardware and software interfaces which control equipment, process data and integrate with the external control system.

Design of the system was guided by external and internal requirements. External requirements were provided by the TRIUMF laboratory and were common to all of the e-linac diagnostic systems. Internal requirements pertaining to the distribution of electrical, mechanical and computing support services were set by the prior design of the view screen devices.

A few requirements, such as the software control system framework, were satisfied with tools provided by the greater particle accelerator community. Where solutions did not exist, such as that for the data processing algorithm, it was necessary to develop purpose-built tools.

Support service hardware was specified and integrated within a large control cabinet. It was necessary to distribute most services over long (up to 70 m) cables due to the magnitude of radiation present during normal accelerator operation. A reference implementation of the data processing algorithm existed but was not yet integrated into the software control system. The new implementation allowed images of the electron beam to be processed and displayed in real-time.

The quality of measurements taken by the view screen system was studied in order to evaluate the satisfaction of functional diagnostic requirements. Sources of systematic bias were observed through detailed simulation of the view screen devices; accounting for bias dramatically improved the quality of measurements under certain

conditions.

Chapter 2 provides the motivation for the development of the system and a birds-eye view of TRIUMF (the lab where the diagnostic device is located) and the larger project framework. External and internal system requirements are specified within this chapter.

Implementation of the control system is described in Chapter 3. This includes the support services, service enclosures, cabling, image processing system and additional computer software (such as the operator interface).

Two experiments are described in Chapter 4:

- validating a portion of the data processing algorithm and
- quantifying the quality of beam profile measurements.

Chapter 2

ARIEL, e-Linac and the View Screen System

This chapter describes the Advanced Rare IsotopE Laboratory (ARIEL) project, e-linac beamline and presents the requirements which influenced the development of the view screen control system.

2.1 ARIEL

TRIUMF is Canada's national laboratory for particle and nuclear physics. Current research programs are based around the existing charged-particle accelerators: a 500 MeV hydrogen-ion cyclotron, a suite of linear accelerators for rare isotopes at the Isotope Separator and Accelerator (ISAC) facility, and low-energy medical cyclotrons. A new facility, the ARIEL, will augment the ISAC Rare Isotope Beam (RIB) science program. The flagship of the ARIEL project is the construction of a 50 MeV, high-average current (10 mA) continuous wave (CW) linear accelerator called the e-linac.

2.1.1 Rare Isotope Beam Production

ARIEL will boast two new RIB producing target stations: one for the e-linac and one for the existing 500 MeV proton beamline. The target stations will be complemented by new RIB selection and acceleration modules that will deliver high-energy and high-intensity RIBs to medical, nuclear, and material science experiments.

While a proton beam already exists, there are advantages in driving RIB production with an electron beam. A few of these advantages are briefly described. The

first advantage is lower stopping power. Near 20 MeV, electrons experience a much lower stopping power than protons when passing through air and metal. This allows the electron beam to pass through an air gap or thin metal window, potentially separating the production beam from the RIB target - simplifying cooling and electrical isolation. The second advantage is reduced beam rigidity; an electron beam is easier to bend, allowing it to be “scanned” onto the target, increasing target lifetime by increasing the amount of material available for Bremsstrahlung production and making cooling easier by spreading the thermal payload over a larger area. The third advantage is the method of RIB production. An electron beam stimulates RIB production via Bremsstrahlung and the (γ, n) photofission nuclear reaction. This two-stage process allows the RIB production station to be separated into two physically separate targets: a “converter” target which produces the gamma radiation and a second RIB target which completes the process. Keeping the converter target isolated from the RIB target potentially simplifies the cooling, electrical isolation, shielding and remote handling of the RIB production station.

The RIB beam produced by photofission will have less isobaric contamination (containing fewer atoms with the same atomic number) as compared to one created via spallation or fission.

An electron beam accelerates to $\beta \approx 1$ much sooner than a heavier ion-beam, lowering construction costs because all accelerating cryo-modules are similar in design.

One disadvantage of RIB production via electron beam is the smaller range and depth of products. The production efficiency (in terms of products per electron) is also lower but this is counter-balanced by being able to operate with a higher beam current.

2.1.2 The e-Linac

The e-linac beam-line is loosely separated into four parts: the electron gun where electrons are produced, an ELBT section which consists of the buncher and injector, an EHBT section inside of which the beam is brought to 30 MeV (50 MeV after an energy upgrade to the machine) and the EHBT section which brings the beam to the RIB-producing target station.

The electron gun contains a hot cathode which produces free electrons via thermionic emission. The free electrons are removed from the cathode with a bias voltage modulated at 650 MHz and accelerated to 300 keV. Once free from the gun, the electrons are

passed through a (non-superconducting) buncher before being accelerated to 10 MeV within the injector cryo-module.

The e-linac beam is designed to operate in two modes: a full-current “continuous-waveform” (CW) mode and a reduced-current “pulsed” mode. Continuous-waveform is used under typical operating conditions while pulsed mode is used for commissioning and diagnostics. Pulsing the beam can prevent damage to sensitive diagnostic devices and reduce radiation caused by excessive beam loss. The beam is pulsed by modulating the 650 MHz CW-beam with a 1 Hz to 10 000 Hz “macro-pulse” (of minimum length 1 μ s).


2.2 The View Screen System

The view screen system is a diagnostic device which images (in real time) the transverse profile of a charged particle beam at fixed locations along the beamline. The system is comprised of the individual view screen devices (which are referred to as view screen *units* when it’s necessary to distinguish them from the view screen system) and the the control system.

The view screen unit (Fig. 2.1a) measures beam-profile properties at a single location along the beam-line. It collectively describes the camera box (Fig. 2.1b), target actuator (with targets), shielding and diagnostic/calibration light sources. View screen unit design was the subject of previous work [8]. The method of beam profile imaging is now described.

An electron beam of constant energy travelling through vacuum does not spontaneously emit light; it must be forced to do so. This is accomplished by actuating a thin target foil into the evacuated beam pipeline; as the beam passes through the target, light is emitted either indirectly through the process of scintillation or directly via Optical Transition Radiation (OTR). The particular method of light production is dependent on the beam energy: scintillation is used for the lower-energy sections of the beam-line while OTR is used for the medium and higher energy sections. An optical pathway transports the light around a ninety-degree bend (allowing the camera box internals to be shielded from radiation) and through a system of lenses after which it is captured by a CCD sensor. The image of the light is processed and then properties of the beam profile (such as position and width) are extracted. Data may be viewed by an operator and also may be recorded to permanent storage.

The view screen control system is the subject of this thesis. It encompasses the



(a) A view screen unit is comprised of a target actuator, a set of target screens, a shielded camera box and calibration light sources.

(b) A View Screen Camera Box. Light enters from the right, is reflected upward through the optical components (an iris and two lenses), before being collected by a camera. Layers of lead and polyethylene shielding to protect the optics and camera from radiation damage.

Figure 2.1: Components of a view screen unit (a) and of a camera box (b).

software and hardware which services and controls the individual view screen units and the implementation of the interface with the TRIUMF diagnostic system. Figure 3.1 shows the services provided and the necessary interfaces. The specification and implementation of the control system was based on given requirements.

2.3 Control System Requirements

One of the main purposes of the control system is to “bridge the gap” between the view screen diagnostic devices and the external e-Linac/ARIEL diagnostic system; therefore, the requirements are naturally divided into those which are motivated by external (project-level) needs and those motivated by the internal aspects of the view screen system.

2.3.1 External Requirements

External requirements arise from role of the view screen system (and other diagnostic devices) as a piece of the larger project framework. A few external requirements imposed by the ARIEL project evolved over time, as lessons learned during the VECC collaboration were incorporated into later phases.

A formal description of the necessary external interfaces can be developed from the e-Linac Diagnostic Requirements Document [4]. Among the many requirements are those which are relevant to the view screen control system. They are grouped into the following categories: environmental, control and functional.

Only one environmental requirement is important to the view screen control system. It is (Req. 1) “Radiation resistance of electronics within 1-3 metres of [the e-linac beam-line]. Hard failure (death) or severe data corruption less than 1/year in radiation field up to 10 mSv h^{-1} ” [4, Sec 2.4] . This requirement is one of the reasons that the camera boxes are shielded - to protect the camera from radiation damage.

There are four control requirements:

- Req. 2. “all beam diagnostics devices are to be controlled by the EPICS based control system.” [4, Sec 3.1],
- Req. 3. “all set points of each device shall be settable by EPICS. Set points will be classified as either Operator or Expert.” [4, Sec 3.2],
- Req. 4. “data will be transferred from devices to control displays and archiving by the EPICS system and will be time-stamped by same. The basic interrogate/refresh rate shall exceed 10 Hz.” [4, Sec 3.3] and
- Req. 5. “acquisition and processing diagnostics are correctly synchronized with the beam [pulse]...” [4, Sec. 3.5].

The first control requirement specifies the use of EPICS as the software interface. This ensures that all diagnostic devices have a consistent interface which is compatible with existing control room tools. The second control requirement is relevant to the control system software and design of the OPERator Interface (OPI). It is necessary to specify control system software which allows the classification of device set points and this must be complimented with tools that respect the permissions of different classes of beam-line operators. The third control requirement (through the interrogate/refresh rate) places a restriction on the minimum processing power of the computing devices

in the view screen control system. If the computer is not powerful enough or the bandwidth of the network is too low, the system will be unable to report beam profile measurements at the specified rate of 10 Hz. The fourth, and last, control requirement ensures that beam measurements are not impacted by the beam mode. The e-Linac is a high-power machine; thus, diagnostic devices which are designed to operate in the high-power regime should also be able to operate in the low-power, pulsed diagnostic mode.

The remaining functional requirements relate to the quality of measurements taken from the beam profile.

- Req. 6. “Diagnostic devices shall provide the stated functionalities within the absolute and relative uncertainties as specified (68 % CL).” [4, Sec 6];
- Req. 7. “measure electron beam position along beam-line (at 10 mA): absolute uncertainty with respect to external survey markers is 0.2 mm; relative uncertainty is 25 μm at 10 Hz display rate (i.e. averaged over samples)” [4, Sec 8.1];
- Req. 8. “For the centre: required absolute uncertainty with respect to external survey markers is 0.2 mm; relative uncertainty is 25 μm .” [4, Sec 9.1.1]
- Req. 9. “The appropriate requirement on beam size and structure within the beam profile is resolution = 5 % of the anticipated r.m.s beam size” [4, Sec 9.1.2]
- Req. 10. “Screens shall have resolution adequate to support dithering; requirement 25 μm .” [4, Sec 9.2];
- Req. 11. “Peak-signal to noise-floor ratio within the 2D image of the beam profile from a screen-type monitor shall exceed 100:1.” [4, Sec 9.4].

Where the requirements describe relative quantities, the values in Table 2.1 are used as a reference. The functional requirements are motivated by the needs of the operators as well as accelerator physicists when carrying out commissioning tasks, orbit calculations and accelerator model validation. The aspect of the control system which is most impacted by these functional requirements is the software image processing algorithm; additionally, they motivate the study of overall systematic error.

As mentioned in the introduction, there were a few informal requirements which evolved over time. These requirements relate to the design of the view screen service enclosures and the implementation (and development) of the software.

Section	Beam Size (rms, mm)	Resolution (mm)
ELBT	4.0	0.20
EMBT	1.0	0.05
EHBT	0.40	0.02

Table 2.1: Host Computer Hardware Components

Due to limited space in the VECC collaboration testing area, and to clearly distinguish TRIUMF diagnostic systems from the UVic view screen system, it was necessary (Req. 12) to place the entire control system within a single, wall mounted enclosure. This requirement was dropped when maintenance of the system was eventually transferred to TRIUMF. It is expected that future view screen service enclosures will be placed on standard equipment racks.

Aside from the dependency on EPICS, there were initially very few other requirements placed on the software system. A document specified the “look and feel” of the operator interface; otherwise, the software component of the system was expected to be a “black box”. Migration of the software to the TRIUMF development and deployment systems (Req. 13) was later added to the e-linac phase of the project. This specified the host computer operating system (a TRIUMF flavour of Linux), software version control system (CVS) and EPICS record management software.

2.3.2 Internal Requirements

Other requirements are motivated by the needs of the view screen units. These requirements are related to camera control, iris control. The view screen requirements can be grouped by the applicable subsystem:

- Camera control:

Req. 14. cameras transmit image data and accept configuration commands through an Ethernet interface,

Req. 15. each camera will be individually powered with the ability to hard-power-off in the event of loss of control and

Req. 16. each camera has two user-programmable output TTL signals and two input triggers.

- Lighting:

Req. 17. each view screen unit has up to two lighting stations,

Req. 18. each lighting station has a separate power line and

Req. 19. the intensity of the lights should be adjustable.

- Iris control:

Req. 20. each view screen unit contains a mechanical iris which is actuated via
stepper motor and

Req. 21. each iris assembly has a forward and reverse limit switch.

The problem now becomes one of specifying hardware, software and image processing algorithms that satisfy the diagnostic requirements.

Chapter 3

A Control System for the e-Linac View Screen System

This chapter discusses the high-level system design, the implementation of control system services, and the construction of a service enclosure. The implementation of the image processing algorithm is discussed in Section 3.5.

3.1 High-Level System Design

A fair amount of flexibility in specifying hardware which met the system requirements was available. A natural area to begin the task was high-level system architecture. In designing the logical layout of the system it was convenient to consider the interfaces requested by the view screen units, the external control system and the internal control system itself. This allowed the design to be divided into a set of orthogonal services which implemented the requested interfaces. The services, as shown in Figure 3.1), are camera control, computing, lighting control, electric power distribution, iris control, and environmental monitoring.

Considering the high-level physical design, the control system was physically separated from the beam-line to ensure that radiation would not damage sensitive hardware (Req. 1). Electrical power, logic signals and data are brought to each view screen by long (up to 70 m) cables. Power to each camera and light is individually switched via mechanical relay, satisfying Req. 15 and Req. 18. Digital and analog signals would be brought into the software through input/output cards.

Based on the three control requirements (Req. 2 through Req. 4), the separation

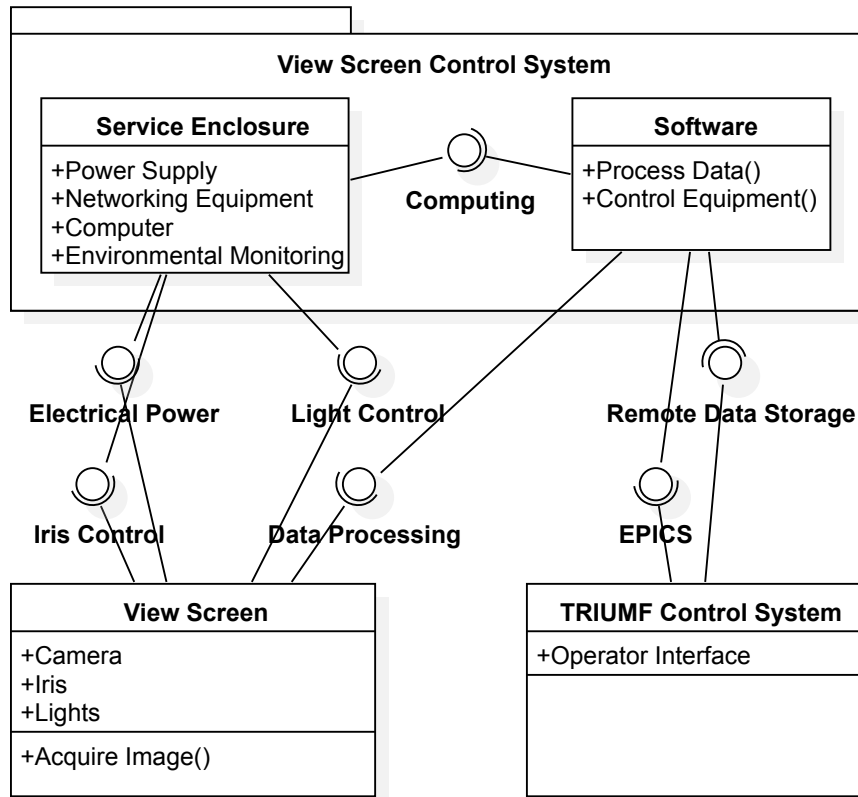


Figure 3.1: System Level Diagram of Control System Interfaces. A majority of the necessary interfaces are implemented as services by the view screen control system. Hardware components are housed in a service enclosure (Sec. 3.3). Software components for controlling equipment and processing data are executed by the computing services.

from (Req. 12) and due to the large amount of hardware and software compatibility, it made sense to use a standard commodity computer as the central access point for the system. The cameras and motor controller cards would be part of a local area network which was isolated from the external TRIUMF network.

3.2 Service Implementation

Necessary control system services were developed based on the requirements in Sec. 2.3.

3.2.1 Iris Control

Iris control is provided by Galil DMC 2183 motor controller cards and computer software.

The Galil DMC-21X3 family of motor controller cards were specified due to the Ethernet interface, their use at TRIUMF for other diagnostic devices and the availability of EPICS compatible software. The cards are populated with Galil SDM-20242 stepper driver modules which provide power to the motor windings. Each SDM-20242 is capable of driving up to four stepper motors. A Galil DB-28040 daughter-board provided an additional forty digital I/O signals and eight analog inputs. These extra signals were used by the camera control and lighting services.

Forward and reverse limit switches prevent the mechanical irises from actuating beyond the intended range of motion. The limit switches are configured as “active-high” and will abort motion when grounded or disconnected. A mechanical homing switch was absent from the first version of the iris actuator. To work around this, a custom software homing routine was written, actuating the iris arm onto the forward limit switch before backing it off until the switch is deactivated. A later version of the iris actuator added a homing switch to the design, allowing the on-board homing routine to be used.

The motor controller software maintains axis position in units of “steps”, but it is more natural for an operator to express the position in terms of the diameter of the iris aperture (in mm). The iris aperture is not perfectly circular, because it is constructed with overlapping, metallic leaves. For this reason, the diameter was measured at multiple locations and the average was used in lieu of a “true” diameter.

The conversion between iris diameter and stepper motor steps is not linear due to the construction of the iris actuator. The relationship was modelled empirically with a function of the form

$$d_1(s) = ae^{bs} + c \quad (3.1)$$

where d is the iris diameter and s is the stepper motor position (see the diagram in Figure 3.2). An advantage of this form is that it is easily inverted, providing the conversion from diameter to steps:

$$s(d) : \mathbb{R} \rightarrow \mathbb{Z} = \text{Round} (d^{-1}(s)) = \text{Round} \left(\ln \left(\frac{d - c}{a} \right) / b \right). \quad (3.2)$$

The relationship was measured by first homing the iris, then measuring the average

diameter every four steps with vernier calipers. Fitting Eq. 3.1 to the measured data yielded residual larger than half of a step (Fig. 3.3). To improve the fit, the transformation function was split into multiple pieces, yielding

$$d(s) = \begin{cases} a_1 e^{b_1 s} + c_1 & : s < s_1 \\ a_2 e^{b_2 s} + c_2 & : s_1 \leq s < s_2 \\ a_3 e^{b_3 s} + c_3 & : s_2 \leq s \end{cases} \quad (3.3)$$

The result of $d^{-1}(s)$ must be rounded to the nearest integer; therefore, residuals with magnitude less than 0.5 steps will agree with measured data. After manually optimizing s_1 and s_2 , a two-piece form ($s_1 == s_2$) was selected and the a_i, b_i and c_i parameters were determined by performing the method of least-squares on the measured data. The continuity of $d(s)$ was maintained at the boundary, rendering one of the parameters redundant. The consistency of $d(s)$ between different irises




Figure 3.2: Iris Coordinate System. The iris was modelled as a perfectly circular aperture of diameter d . The variable s is the number of *steps* from a reference position (in this case - the reverse limit switch).

was not evaluated.

Software support for the motor controller card was provided by a combination of the EPICS motor module [1] and a module produced by Australian Synchrotron (AS) [2]. The AS module was extended by adding support for interacting with the additional I/O on the DB-28040 daughter-board.

3.2.2 Camera Control and Triggering

The bulk of camera control is provided by the software driver; however, a few features are provided by hardware in the control cabinet: camera power, diagnostic readback




Figure 3.3: Iris Position Conversion Fit Residuals. Fitting a one-piece exponential form ($d = ae^{bs} + c$) to the iris measurements yielded large residuals (blue). Fitting a two-piece exponential to the data left smaller residuals (yellow).

and triggering.

Software support for the AVT Manta camera was found within the areaDetector EPICS support module [7]. The driver was extended by adding camera frame timestamp synchronization and proper handling of camera connect / disconnect events. Along with an EPICS interface to the camera driver, areaDetector provided an EPICS-controlled image processing framework. This framework was instrumental in implementing the image processing algorithm (Sec. 3.5).

Individual camera power lines are switched with solid-state relays (satisfying Req. 15). This allows an operator to disable or “hard reset” a camera without entering the (potentially) contaminated beam-line area. The relay logic is managed by digital output lines on the motor controller cards.

Camera acquisition can be configured to start and stop on internal or external events. This is used to satisfy the beam synchronization requirement (Req. 5). It is important that the camera frame is triggered with the beam macro pulse when the frame exposure time is less than or equal to the macro-pulse length and the beam is intercepted by an OTR or fast scintillation target. If the frame exposure is not triggered by the beam, successive camera frames will not collect light from the same portion of the macro-pulse (Fig. 3.4). The beam intensity will appear to strobe and the longitudinal beam profile could introduce additional irregularities into the image.

There are two trigger inputs on the AVT Manta 046B camera. One is wired to a BNC connector on the camera box enclosure and the other is brought back to the cabinet through the umbilical cable; although, currently the cabinet trigger line is not wired to an external connector.




Figure 3.4: Unsynchronized vs Synchronized Camera Trigger. The amount of light collected by the camera (overlapping purple area) can vary greatly when the frame acquisition state (blue) is not synchronized with the macro-pulse (red).

Two camera output signals are brought back to the cabinet through the umbilical cable and read back by the motor controller cards. These operator-configurable signals are used to provide imaging status information, such as: trigger ready, exposing and frame readout.

3.2.3 Computing

Computing services host the system software, manage the local network and interact with the external control system. These services are implemented with software and hardware components.

A computer was needed to execute the system software. Diagnostic requirements and software support limitations place restrictions on the computing hardware. The AVT Manta camera driver required an x86 architecture. The minimum refresh rate (Req. 4) placed a restriction on the minimum processing power. A processor supporting vector operations would allow the software to execute with greater efficiency as most of the image processing algorithms operate on data which is stored in memory sequentially.

The minimum amount of system memory is constrained by the image processing software. Each image processing plugin maintains a pool of N_q multi-dimensional (NDArray) data structures. Let $W \times H$ be the beam-space image dimensions (in px), b be the image bit-depth (in bit px⁻¹), N_p be the number of image processing plugins and N_v be the number of view screen units managed by the host computer. The total

amount of memory consumed by the NDArray pools is roughly

$$M = (W \times H) \times b \times N_q \times N_p \times N_v$$

$$\approx 2^{19} \times 2^4 \times 2^3 \times 2^4 \times 2^4 = 2^{34} \text{ bit} = 2 \text{ GiB}$$

for a system with 16 view screen units. Additional system memory will be consumed by the operating system.

The computer operating system was chosen to be Ubuntu 12.04 based on the long term support, availability of software packages and familiarity.

The computers, cameras and motor controller cards communicate through Ethernet interfaces. The network configuration (addressing and host-name resolution) is managed by a DHCP server (ISC DHCP) running on a host computer while the physical layer is established by a dedicated network switch. All networked devices reside on a local area network solely accessible through the host computer. This segregation protects the devices from accidental (or intentional) tampering and prevents the large amount of camera traffic from saturating the external control network. The host computer must have two network adapters, allowing it to communicate with each of the networks.

Under typical operating conditions, each camera consumes roughly 50 Mbit s^{-1} of bandwidth. Driving the frame rate higher can increase this amount by an order of magnitude. The local network was wired with Gigabit Ethernet cable, ensuring that it wouldn't be saturated with camera frame data.

Data storage requirements are motivated by the need to keep copies of processed data for offline (or postmortem) analysis and the necessity of a location from which the operating system and view screen software can be loaded. It is not necessary that the storage requirements are implemented with a single storage solution. In fact, it's desirable to keep them separate; the view screen data may need to be accessible by multiple operators while the software should only be modified by expert users. Access speed must also be considered because writing sequential data sets generates a burst of file-system activity.

The processed data and system software were initially stored on a single, solid-state hard drive but this implementation evolved over time. The processed data storage location became switchable between the local hard drive, a network file system folder and an operator-defined location. When the view screen software was later deployed on an in-house Linux distribution, it dramatically reduced the operating

system footprint and enabled the host computer to be bootstrapped over the network via Preboot Execution Environment (PXE). The local solid-state drive was eventually replaced with a RAM drive.

The computing resources listed in Tables 3.1 and 3.2 meet the aforementioned requirements and were used in the reference hardware implementation.

Function	Component
Processor	Intel i3 2100 3.1 GHz
Motherboard	Intel DH61DLB3
Power Supply	M4-ATX 250 W
Memory	Mushkin 8 GB DDR-3 1333
Case	Polywell ITX-500B (G4100)
Storage (initial)	Kingston SSD V100 64 GB
Ethernet Adapter 1	Intel 82579V (on-board)
Ethernet Adapter 2	Intel 82574L

Table 3.1: Host Computer Hardware Components

Function	Device
Network Switch	Allied Telesis AT-GS900/24
Network Storage	NFS 3 via RHEL Server
Network Boot	PXE Server

Table 3.2: Additional View Screen Computing Resources

Interaction with the external control system was enabled by the EPICS Channel Access protocol. Channel access allows remote input/output controllers to share contents of EPICS records. Other forms of communication include interfacing with an external file storage server (over SSH) and the transmission of the operator interface to external displays (using the X-11 forwarding protocol).

3.2.4 Electric Power Distribution

Electrical power is required by components of the view screen unit as well as the control system. For view screen units, it's required by the cameras and diagnostic lights; for the control system, it's required by components such as the computing devices and switching relays. Distribution of electricity to the view screen units was somewhat complicated by the distance between the control system hardware and the

view screen units (more about this in Section 3.4). Power distribution was divided into three individually fused circuits:

- a 24 V DC circuit for the computer systems, motor controller cards, cameras and fan,
- a 12 V DC circuit for the lighting sub-system(s) and
- a 5 V DC circuit for the camera digital I/O lines.

Power was supplied to the 24 V and 12 V DC circuits by Cosel PBA600F-24 and PBA150F-12 power supplies, respectively. The 5 V circuit was powered by a standard the 5 V, 2 A power adapter. Each circuit is protected by a fuse with a rating based on the expected power load.

3.2.5 Lighting

The lighting system was designed to satisfy Req. 17 through Req. 19. Two identical lighting circuits were implemented (Fig. 3.5), one for each of the two possible lights on each view screen unit. Lights are individually switched with solid-state relays controlled by the Galil DB-28040. Light intensity is adjusted on a circuit-wide basis by means of a pulse-width modulation (PWM) dimmer. The dimmers are adjusted by purpose-built, 8-bit digital-to-analog converters (DAC) which are also controlled by the Galil DB-28040.

Within a circuit, individual lighting stations are wired in parallel. To reduce the circuit load, software controls prevent more than one station from being activated at a given time.

3.2.6 Environmental Monitoring

The two components of system were monitored: temperature and the current in the lighting circuits.

An Omega EWS-TX temperature transmitter monitored the temperature of the enclosure which housed the control system hardware. The 0 V to 10 V output of the transmitter was read by an analog input channel on the Galil DB-28040. Temperature conversion was calibrated in-situ with a heat gun and infrared thermometer (Fig. 3.6). All specified hardware components have a maximum operating temperature of at least

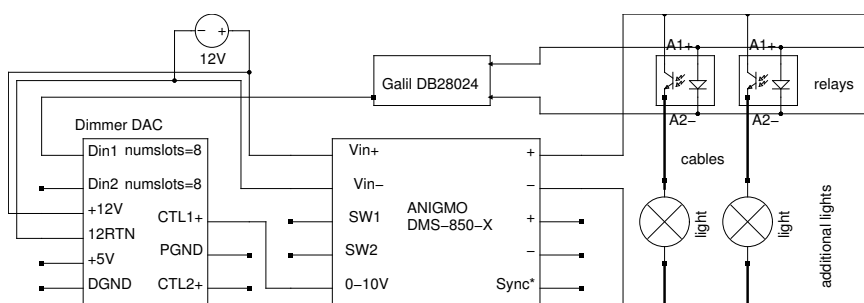


Figure 3.5: Cabinet Lighting Circuit. One of two identical cabinet lighting circuits. The intensity of the lighting stations is adjusted by a PWM dimmer. The dimmer accepts a 0-10 V control voltage which is supplied by a 8-bit DAC. The stations are individually activated with solid-state relays.

50 °C. The EPICS record which monitored the temperature was set to alarm when the temperature was too high.




Figure 3.6: Cabinet Temperature Calibration.

Hall probes measure the current passing through each of the two lighting circuits, allowing an operator to remotely diagnose halogen bulb filament expiration.

3.3 Service Enclosures

Service enclosures house the hardware which provides control system services. Two generations of service enclosures were designed.

The first generation service enclosures were implemented as a single, wall-mounted electrical cabinet (hereafter called, simply, *the cabinet*). The large (36 × 60 inch) Hammond 1418T10 cabinet contains every piece of equipment needed to support the sixteen individual view screens. It was designed to meet the initial requirement of complete separation from the external TRIUMF control system (Req. 12). It was deployed in the VECC test area to service five view screens and later, during e-linac

commissioning, the cabinet was redeployed to the e-Hall rooftop where it serviced sixteen view screen units.

As mentioned previously, every piece of view screen control system hardware was placed within the cabinet. Internal components include those for computing (two computers and a network switch), power distribution, iris control and connectivity (a large amount of break-out wiring). A purpose-built signal processing board steps down the temperature transmitter voltage and averages the feedback from the current sensors.

Ethernet and umbilical cables are brought in through glands on the right wall of the cabinet. The glands provide cable stress relief and act as a dust-guard. The interior of the cabinet is kept at a slight negative pressure by an adjustable-speed Axial fan. Cool air passes through a filter on the bottom right side, flows over the computer power supplies and motor amplifiers before exiting the top of the cabinet. AC power is introduced to the cabinet through a six-port Hammond power strip to which the DC power supplies and Ethernet switch are connected. Transportation of the cabinet is possible by crane via hook on the top or by means of a purpose-built, wheeled sled.

An image of the cabinet is shown in figure 3.7.

The second generation of service enclosures divide the system into smaller pieces. This came as a result of Req. 12 being dropped when more integration with the TRIUMF control system was desired. Power supplies and computing hardware were moved into standard equipment racks; motor controllers were put into individual enclosures (Fig. 3.8). These enclosures, each of which is able to service eight view screen units, acted as the “hub” for control system services. Dividing the system allows components to be serviced individually and placed with similar services (for example, the view screen computing resources can be placed on a rack with the computing resources belonging to other diagnostic devices). It is envisioned that the second generation of service enclosures will be used in the higher-energy section of the beam-line as well as in the future addition of an energy recirculating accelerator.

3.4 Cabling

Control systems for the diagnostic devices were kept separated from the beam line due to the large amount of radiation present during operation. Each view screen was attached to the service enclosure by two long (up to 70 m) cables: a Gigabit Ethernet

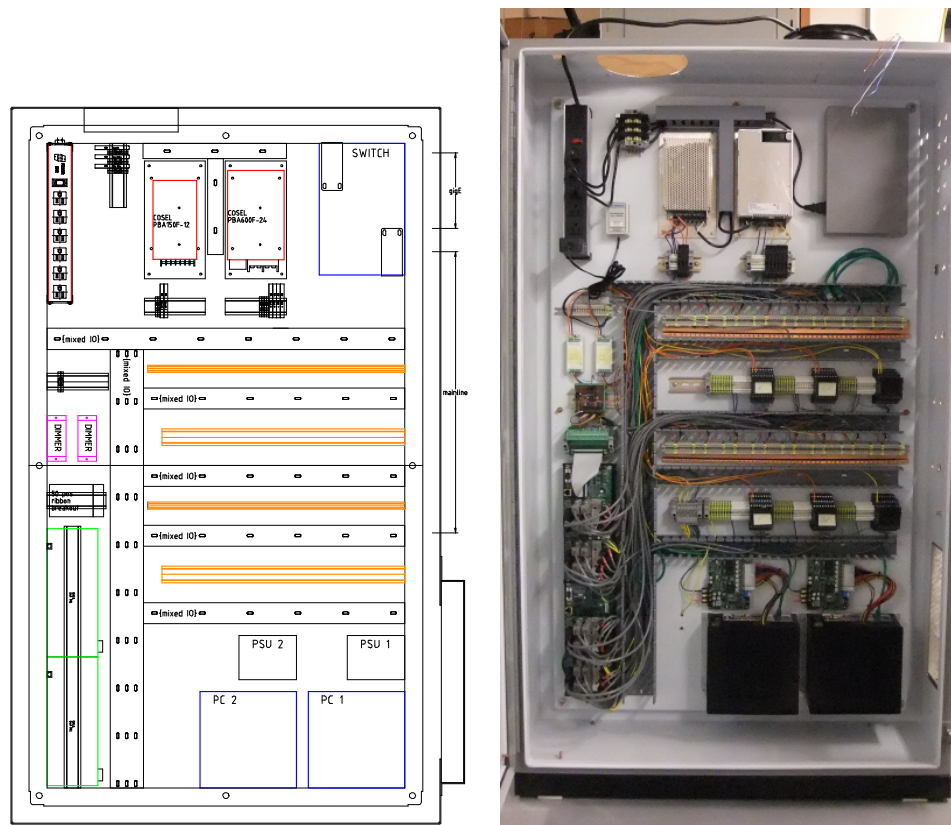


Figure 3.7: A Drawing of the View Screen Cabinet. This first-generation enclosure holds every piece of control system hardware required to service sixteen view screen units. The components are coloured by the service supplied: computing (blue), power supply and distribution (red), iris control (green), lighting (purple) and connectivity (orange).

cable (which carries camera data) and bulky “umbilical” cable.

The umbilical cable carries the iris power, iris limit switch signals, camera power, camera triggers and camera programmable output signals. It contains five 18 gauge (American Wire Gauge (AWG)) and five 20 gauge individually shielded twisted pairs. Additionally, the cable has an overall shield. The camera box end of the cable is terminated by a military-grade Amphenol MIL-DTL-26482 connector. The umbilical cable was wired directly into the first generation service cabinet. The second generation saw the enclosure-end of the cable terminated with a Wieland revos industrial connector.

Cable durability was kept in consideration due to portions of the cable living within the radioactive beam-line area.




Figure 3.8: A Drawing of the Second Generation View Screen Service Enclosure. This second-generation enclosure holds the motor controllers, a signal processing board, environmental sensors and can service up to eight view screen units.

3.5 Image Processing

Similar to many imaging systems, the image collected by the CCD sensor must be processed before meaningful information can be extracted. An image processing routine transforms the raw input image into a calibrated representation of the transverse beam profile (hereafter called the *beam-space* image), calculates beam statistics and styles the image before presentation to the beam operator.

Image reconstruction is complicated because there is typically insufficient information to uniquely determine the beam profile which produced the captured image. Light collected by a particular CCD pixel may have originated from a multitude of starting positions. Additionally, unwanted light may be introduced into the CCD image through sensor noise, internal reflection within the lenses, and photo-emission from background radiation.

Many different algorithms for image reconstruction exist. Iterative techniques, which are popular in computed tomography, can reconstruct the “most likely” object which produced the captured image. Techniques for removing background noise and accounting for a position-dependent point spread function can be integrated into the

process. Unfortunately, iterative algorithms are computationally expensive and may be difficult for a beam-operator to use.

A more traditional approach to image reconstruction was chosen, dividing the reconstruction routine into three stages. In the first stage, a geometry correction algorithm removes distortions and transforms the image into the correct orientation. The second stage, magnification correction, adjusts the intensity of each pixel to ensure that the non-uniform magnification introduced by rotating the beam target does not unintentionally dim or brighten portions of the image. The third stage corrects for the light collection efficiency of the optical system. One drawback of this approach is that no attempt was made to remove “blurring” caused by the point-spread-function causing systematic bias to be introduced into beam profile measurements. This bias was measured in Section 4.2.2.

3.5.1 Coordinate Systems

Coordinate systems are defined in the two spaces of interest, image and beam space.

The CCD camera records images by measuring the intensity of light in each pixel. The AVT Manta G-046B camera used in the view screen system has $8.3\ \mu\text{m}$ wide square pixels with a resolution of 780×580 and uses eight or twelve bits to represent the intensity on each pixel. The 2D space defined by the CCD imager is called image space. Pixels within image space are labelled with row j and column i (indexed from zero). A more convenient way to reference a pixel is to use a single index $\mu = jn_c + i$ where n_c is the number of columns of pixels in the image. For an image I , the notation I_μ refers to the intensity of pixel μ . Written this way, μ is assumed to be a natural number within the interval $[0, 780 \times 580 - 1]$. Sub-pixel positions are represented by the extension of i and j to real numbers. The origin of image space is the centre of pixel 0.

The local e-linac beam coordinate system is right-handed; the positive z (or s) basis vector points along the nominal beam direction, positive x points to the “left” (when looking down the beamline) and positive y points vertically upward [5]. Transverse beam-space, hereafter called beam space, is a two dimensional surface which is perpendicular to the nominal beam direction (Fig. 3.9). The coordinate system for beam-space is the x and y coordinates of the local beam coordinate system with an origin at the approximate centre of the beam pipe.

An expert operator chooses the region of interest (ROI) of beam space that will

be represented by the image produced by the image reconstruction algorithm. The ROI is defined by the beam space extents and number of rows (n'_r) and number of columns (n'_c) of the image. For presentation of the image in beam-space, a pixelated coordinate system is defined by row number u and column number v or by a single index $\nu = un'_c + v$, in a manner similar to the image-space coordinate system.

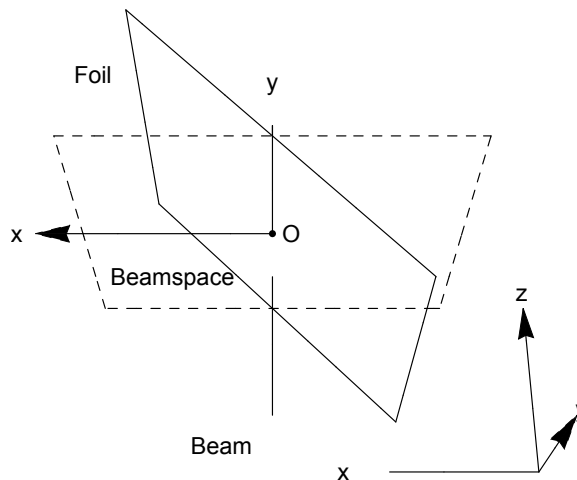


Figure 3.9: Beam-space Coordinate System. The beam-space plane (dashed area) is transverse to the nominal beam direction. The origin of the beam-space coordinate system is approximately the centre of the beam tube. Note that the x-axis is reversed due to the orientation of the local beam coordinate system.

Coordinate Transformations

The transformation between a point (x, y) in the beam-space coordinate system and a point (i, j) in the *image-space* coordinate system takes the form of a second degree multivariate polynomial:

$$\begin{bmatrix} i \\ j \end{bmatrix} = \vec{g}(x, y) = \begin{bmatrix} a_0 + a_1x + a_2x^2 + a_3y + a_4xy + a_5y^2 \\ b_0 + b_1x + b_2x^2 + b_3y + b_4xy + b_5y^2 \end{bmatrix} \quad (3.4)$$

where the a_0, \dots, a_5 and b_0, \dots, b_5 coefficients are determined during device calibration (Sec. 3.5.2).

The transformation between a point within the operator-defined beam-space region-of-interest (Fig. 3.10) and the centre of pixel (u, v) in the *beam-space* image I' is

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x_l + \left(u + \frac{1}{2}\right) \frac{x_r - x_l}{n'_c} \\ y_t - \left(v + \frac{1}{2}\right) \frac{y_t - y_b}{n'_r} \end{bmatrix} \quad (3.5)$$

where y_t , x_l , y_b and x_r are the top, left, bottom and right extents of the beam-space region of interest, respectively.




Figure 3.10: Beam-space Region of Interest. The extents of the beam-space region of interest are set by an operator during device calibration. For a right-handed coordinate system x_l will be greater than x_r . Note that the extents of beam-space are aligned with pixel boundaries, not pixel centres.

3.5.2 Geometry Correction

Any light emitted by the target foil must travel through the optics system before being captured by the camera's CCD sensor. Along the way, light rays may deviate from the ideal path due to misalignment of optical components and non-linear refraction caused by off-axis traversal of the optical system. Additionally, the relative orientation of the target foil, mirror and camera will rotate and flip the collected image (Fig. 3.11).

The resulting perspective distortions are removed by applying a geometry correcting algorithm to the raw camera image. When done correctly, the algorithm produces a two-dimensional representation (beam-space image) of the transverse beam profile.

The geometry correction algorithm uses information about the transformation between image-space and beam-space coordinates to reconstruct the beam profile. For this reason, the device must be calibrated before any corrections can be applied.




Figure 3.11: Line segments from the beam-space coordinate system (left) experience a flip, rotation and magnification upon transformation to the raw-image coordinate system (right). These distortions must be removed before properties of the beam profile can be measured.

Calibration

The device calibration process determines parameters a_i and b_i for the coordinate transformation function \vec{g} . It does so by extracting a set of control points from an image of the calibration target (Fig. 3.13), associating each one with a point in beam-space and applying the method of least squares.

The control points are a set of holes which are machined through the calibration target face. Extracting control points from an image of the calibration target is made difficult by consequences of calibration system design constraints.

One constraint is the limited number of diagnostic ports in the low-energy diagnostic boxes. When an existing port was not available for the diagnostic light, an additional port was machined diagonally through the front-face of the diagnostic box (Fig. 3.12). A light attached to this port is not pointed directly at the calibration foil, preventing uniform calibration target illumination and allowing light to reflect about the interior of the box. This renders control point identification difficult (Fig. 3.13a) and inaccurate. As a workaround, the hot cathode was used as a light source for the view screen units close to the electron gun.




Figure 3.12: A custom port was machined into the front face of all ELBT diagnostic boxes, facilitating calibration target illumination.

Another difficulty is introduced by a calibration target design constraint. The holes in the calibration target are chamfered, allowing light originating from the rear of the target to pass through unimpeded (Fig. 3.14). To ensure that the calibration holes are positioned and sized accurately, the chamfer is not drilled clear through the target. Despite its size being only a fraction of the hole depth, the remaining edge is large enough to introduce a bias into the control point centroid measurement (Fig.




Figure 3.13: Two images of the same calibration target illuminated from a different source. The image of the front-lit calibration target (a) is over-exposed at the top and under-exposed at the bottom. Some control points are nearly indistinguishable from the background. It is easier to distinguish control points from the background when the calibration target is back-lit (b).

3.15).




Figure 3.14: Cutaway Side-view of the Calibration Target. The holes in the calibration target (shown at the left) are chamfered, allowing light from the rear of the target to pass through the hole unimpeded. This is necessary because the targets are rotated 45 degrees about the vertical axis.

Any algorithm, whether executed by an operator or computer, must remain robust, repeatable and easy-to-use within a range of lighting conditions and take known systematic bias into account. During initial view screen system deployment, control point positions were extracted manually by an expert operator. Manual extraction is error-prone, time-consuming and imprecise; a single image processed by two different operators could yield different calibrations. Software methods of control point identification present their own challenges:

- basic image statistics, such as the mean and standard deviation, are influenced




Figure 3.15: Simulated Images of the Centre Hole in an ELBT Calibration Target. When the fully illuminated hole (a) is partially obscured by the chamfer edge, some light does not reach the camera. Left unaccounted for, this would introduce an offset (b) into all control point locations.

by artifacts (such as the bright reflections in the front-lit image);

- deviations in the calibration target position or rotation angle may cause feature extraction to fail catastrophically;
- new calibration target geometries necessitate software changes;
- for cases where only a few holes are unrecognized, procedures for expert-operator intervention need to be developed.

An operator-guided control point extraction algorithm was developed, aiming to harness the strength of both techniques while mitigating the drawbacks. The algorithm uses a feature-recognition algorithm to locate the control points (holes machined through the calibration target) within an image of the target. The features (images of the calibration holes) change under different lighting conditions. To ensure the algorithm's effectiveness under all lighting conditions, a physical model of a calibration hole was developed, dividing it into three discrete components: the face of the hole, the chamfer edge and the surrounding target surface. An image of each component under uniform illumination was produced with the optical simulation (Fig. 3.16). A complete image of the feature, K , is expressed as a combination of the three component images:

$$K = \alpha K_{face} + \gamma K_{hole} + f(\beta_1, \beta_2, K_{edge}) \quad (3.6)$$

where α, β_1, β_2 and γ are real numbers within the interval $[0, 1]$. The function $f(\beta_1, \beta_2, K_{edge})$ applies a gradient to K_{edge} by scaling the pixel intensities linearly

between β_1 (the intensity of the left-most pixel) and β_2 (the intensity of the right-most pixel). This parameterization accounts for shading on the edge component caused by directional lighting.




Figure 3.16: Three Sub-components of a Calibration Target Hole. The features used by the image recognition algorithm must be described as a combination of the hole, edge and target-face images.


An operator describes the feature by estimating values for α, β_1, β_2 and γ based on the calibration target lighting conditions (Fig. 3.17) and a digital image correlation routine measures the “distance” between the feature (represented by the $h_K \times w_K$ matrix K) and sub-regions of the calibration image. The metric (distance measure) between K and η , a submatrix of the calibration image I , is the normalized, squared, Euclidean distance. This is calculated for each pixel in I and the results are stored in a matrix M . The elements of M are

$$M_{ji} = \frac{((\eta - \bar{\eta}) - (K - \bar{K}))^2}{2((\eta - \bar{\eta})^2 + (K - \bar{K})^2)},$$


where η is the $w_K \times h_K$ submatrix of I centred on (i, j) and the following matrix notation is used: \bar{X} is a vector containing the mean of each column of X , $X - \bar{x}$ is a matrix with the vector \bar{x} subtracted from each column of X and X^2 is the square of the Euclidean norm ($X^2 = \|X\|_E^2$).

Elements of M fall within the range $[0, 1]$; values closer to 0 imply that η and K are similar (the distance between η and K is shorter). Control point identification is now accomplished by finding local minima within M .

A set of N initial search points, each corresponding to a calibration target control point, is produced upon selection of the calibration target geometry and camera box orientation. For each search point:



(a) EGUN:VS1 Calibration Target, Front-Lit. Simulated feature (left) with $\alpha = 0.2, \beta_1 = 0, \beta_2 = 1$ and $\gamma = 0.45$. Measured feature (right).



(b) EMBD:VS2 Calibration Target, Back-Lit. Simulated feature (left) with $\alpha = 1, \beta_1 = 0.35, \beta_2 = 0.35$ and $\gamma = 0$. Measured feature (right).

Figure 3.17: Operator Description of the Calibration Target Hole Features. An operator describes the feature by selecting values for α, β_1, β_2 and γ .

1. the smallest element within a 13×17 neighbourhood of the search point is found,
2. a paraboloid is fit to the data surrounding the smallest element and
3. the point for which the fitted paraboloid is at its minimum is taken as the estimated feature location.

The set of estimated feature locations, $\{(i, j)_{1...N}\}$ and set of corresponding beam-space locations $\{(x, y)_{1...N}\}$ are used to find the coefficients of Eq. 3.4) by applying the method of least-squares to the over-determined linear systems

$$\begin{bmatrix} 1 & x_1 & x_1^2 & y_1 & y_1x_1 & y_1^2 \\ 1 & x_2 & x_2^2 & y_2 & y_2x_2 & y_2^2 \\ \vdots & & & & & \vdots \\ 1 & x_N & x_N^2 & y_N & y_Nx_N & y_N^2 \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{bmatrix} = \begin{bmatrix} i_1 \\ i_2 \\ \vdots \\ i_N \end{bmatrix}$$

and

$$\begin{bmatrix} 1 & x_1 & x_1^2 & y_1 & y_1x_1 & y_1^2 \\ 1 & x_2 & x_2^2 & y_2 & y_2x_2 & y_2^2 \\ \vdots & & & & & \vdots \\ 1 & x_N & x_N^2 & y_N & y_Nx_N & y_N^2 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \end{bmatrix} = \begin{bmatrix} j_1 \\ j_2 \\ \vdots \\ j_N \end{bmatrix}.$$

(3.7)

Geometry Correction Algorithm

The geometry correction algorithm uses the transformation function, \vec{g} , and a set of operator-defined parameters to transform the raw CCD image into an image representation of beam-space.

The operator-defined parameters are the extents of the desired beam-space region of interest and the dimensions (size) of the beam-space image. These parameters are defined during device commissioning and are a part of the view screen software device configuration file. When the device configuration is loaded, the algorithm is initialized by iterating through the beam-space image pixel and determining which raw pixels contribute to its construction.

A general expression for the intensity of pixel ν of the transformed image I' is a weighted sum of the raw image pixel intensities

$$I'_\nu = \sum_{\mu} c_{\nu\mu} I_\mu \quad (3.8)$$

where the coefficients $c_{\nu\mu}$ can be determined by a number of means. Two methods of determining the $c_{\nu\mu}$ were investigated: the method of bilinear interpolation and the method of area overlap.

For the method of bilinear interpolation each pixel of I' receives a contribution

from (at most) four pixels in I , determined by the following algorithm. The centre of pixel ν is transformed into beam-space with Eq. 3.5 and then into the image-space containing I with Eq. 3.4. If (i, j) is the transformed centre of ν then

$$I'_\nu = \begin{aligned} & (i_2 - i)(j_2 - j)I_{\mu_{11}} & + (i - i_1)(j_2 - j)I_{\mu_{21}} \\ & + (i_2 - i)(j - j_1)I_{\mu_{12}} & + (i - i_1)(j - j_1)I_{\mu_{22}} \end{aligned} \quad (3.9)$$

where I_{μ_*} are the intensities of the nearest four pixels having centres (i_1, j_2) , etc (Fig. 3.18). The values of the four non-zero coefficients of $c_{\nu\mu}$ for a given ν are the coefficients of Eq. 3.9. The coefficients are normalized in the sense that $\sum_\mu c_{\nu\mu} = 1$.

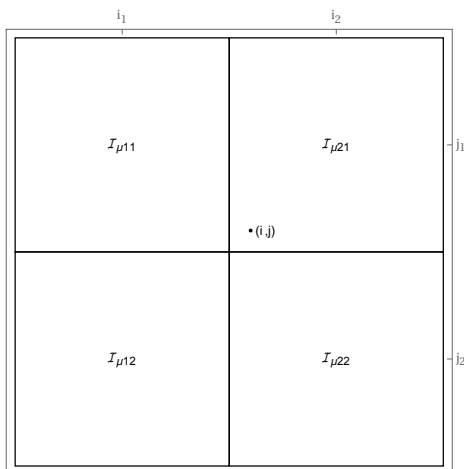


Figure 3.18: Bilinear Interpolating Pixel Intensities. One method of calculating pixel intensities is interpolating the value between four neighbouring pixels.

Efficiency, in terms of execution speed and memory consumption, is an advantage of this method. The expression for each coefficient is calculated with two subtractions and one multiplication operation, avoiding time-consuming division and branching instructions. One difficulty is handling the “edge case” where ν transforms to a location outside of I . A drawback is that straight lines in beam-space are not preserved in I' . This suggests that there are higher-order corrections which are being dropped by the bi-linear interpolation.

A second method of calculating $c_{\nu\mu}$ is the method of area overlap. In this method, each coefficient is proportional to the area of the intersection between I'_ν and I_μ . The corners of the square representing I'_ν are converted to beam-space (Eq. 3.5) and then transformed to raw image-space with the transformation function \vec{g} . In raw image-space, the quadrilateral formed with the transformed corners of I'_ν is intersected with

nearby “pixels” using the General Polygon Clipper library ([6]). From the resulting intersections, the amount of overlapping area, $a_{\nu\mu}$ is calculated.

To ensure that pixels near the edge of the image are not artificially dimmed, the $a_{\nu\mu}$ coefficients are normalized:

$$c_{\nu\mu} = \frac{a_{\nu\mu}}{\sum_{\mu} a_{\nu\mu}}.$$

The coefficients, as determined by the method of area overlap, are identical to the values obtained by the method of bi-linear interpolation when the geometric transformation function is linear and the pixel density of the input and output images (as measured in px mm^{-2}) are identical.

Implementation

The geometric correction algorithm implements the ViewScreenConfiguredNDPlugin interface, which is a specialization of the AreaDetector array processing plugin. When a new image is received, a compatibility check is performed. This ensures that the dimensions of the image match those which were provided in the device configuration file. A table of geometric correction information (Fig. 3.19) is constructed during device configuration. The table entry at row v , column u contains the data for constructing pixel I'_ν ($\nu = v \times \text{Width}(I') + u$) within the beam-space image. The data consists of a fixed-length list of pairs where the first element is the coefficient $c_{\nu\mu}$ and the second element is the index μ .

		u = 0								1				
v = 0		0.34915	0.15779	0.32654	0.16651	0.	0.	0.	0.	0.314	0.10804	0.41473	0.16323	...
		440724	440725	441504	441505	441505	441505	441505	441505	439944	439945	440724	440725	...
1		0.37473	0.13239	0.35121	0.14168	0.	0.	0.	0.	0.33526	0.086947	0.44356	0.13423	...
		440725	440726	441505	441506	441506	441506	441506	441506	439945	439946	440725	440726	...
		⋮								⋮				⋮

Figure 3.19: Geometric Correction Table. The table entry at row v , column u contains the information to construct pixel I'_ν . Each two-component column contains the coefficient $c_{\nu\mu}$ (green) and array index (blue) of an input pixel which contributes to I'_ν . A maximum count of eight input pixels may produce a single output pixel.

The beam-space image is constructed pixel-by-pixel by looping through the correction table, performing the sum in Equation 3.8. By fixing the number of coefficients in each correction table entry, the program compiler is free to unwind the sum and perform the calculation more efficiently.

3.5.3 Magnification Correction

The target foil is rotated 45° toward the optical pipe in order to maximize the amount of cross sectional area presented to the beam and the optical system. A consequence of the rotation is that non-uniform magnification is introduced into the image collected by the optical system due to one side of the target being closer to the camera than the other. Assuming that the target is emitting light uniformly, areas which experience more magnification will appear dimmer than those with less magnification because the same amount of light is distributed over a larger area of the CCD sensor (Fig. 3.20). If the non-uniform magnification is not accounted for, beam position measurements could be biased toward the edge of the foil which is furthest from the target.




Figure 3.20: ELBT Pixel Area Density. The side of the target which is closer to the view screen camera box (positive x-values, in this case) experiences more magnification than the other side. An object which is on the closer side will fill more pixels on the CCD sensor, appearing dimmer.

Differences in magnification are removed by multiplying the intensity of each pixel in the beam-space image by a per-pixel scaling factor, M_v , which is proportional to the magnification of the area represented by the pixel. While it is possible to estimate the magnification from the design parameters of the optical system (such as lens strength and position), any per-device sources of bias (such as component alignment, and camera position) will be left unaccounted for. For this reason, the magnification is calculated from the geometric transformation function (Eq. 3.4) which is measured during device commissioning.

The Jacobian determinant of \vec{g} describes how the size of an area element $dx dy$ in

beam-space is scaled when converted into image-space. Let

$$m(x, y) = \left| \frac{d\vec{g}(x, y)}{d\vec{x}} \right| = \left| \begin{bmatrix} a_1 + 2a_2x + a_4y & a_3 + a_4x + 2a_5y \\ b_1 + 2b_2x + b_4y & b_3 + b_4x + 2b_5y \end{bmatrix} \right| \quad (3.10)$$

be the pixel area density of a point (x, y) on the target foil plane (measured in units of px/mm²). The scaling factor

$$M_\nu = \frac{m_\nu}{\bar{m}} \quad (3.11)$$

will correct for non-uniform magnification on pixel $p_{\nu u}$ while preserving the overall image intensity; the numerator

$$m_\nu = \frac{\int_\nu m(x, y) dx dy}{\int_\nu dx dy} \approx m(x_\nu, y_\nu) \quad (3.12)$$

is equal to the pixel area density averaged over pixel ν ((x_ν, y_ν) is the centroid of ν converted to beam-space with Eq. 3.5) and the denominator

$$\bar{m} = \frac{\int_{ROI} m(x, y) dx dy}{\int_{ROI} dx dy} \approx \frac{\sum_\nu m_\nu}{\sum_\nu 1} \quad (3.13)$$

is equal to the pixel area density averaged over the entire region of interest.

3.5.4 Collection Efficiency Correction

Not all light emitted by a point on the target foil will reach the camera's CCD sensor (Fig. 3.21). The *collection efficiency* $e(\mathcal{O}, \mathcal{E}, x, y, d)$ of the optical system is the probability of collecting a photon for a given optical configuration \mathcal{O} , light emission distribution \mathcal{E} emitted within the interval $[x, x + dx] \times [y, y + dy]$ of the target foil and passing through an iris aperture of diameter d .




Figure 3.21: Not all of the light emitted by the target foil (right) reaches the camera's CCD sensor (upper-left).

Due to the positional dependence of the collection efficiency, the intensity of the observed beam profile will be dimmer where the efficiency is lower and brighter where the efficiency is higher. To quantify this effect other sources of distortion, such as the point-spread-function, are temporarily ignored. It will also be assumed that the geometry and magnification corrections have been perfectly accounted for and that the amount of light emitted by a small region of the target foil increases linearly with the amount of current passing through it.

The observed beam profile can be found by considering the beam profile as a probability density function (PDF) $b(x, y)$ where $\iint_A b(x, y) dA$ is the portion of beam which passes through beam-space region A . The PDF of the measured beam profile is equal to the conditional probability of light from interval $[x, x + dx] \times [y, y + dy]$ reaching the CCD sensor in the presence of collection efficiency e :

$$b'(x, y) = \frac{b(x, y) e(\mathcal{O}, \mathcal{E}, x, y, d)}{C_e} \quad (3.14)$$




Figure 3.22: Due to the angle of the target foil, geometry of the ELBT camera box and shape of the scintillation light emission distribution only a small fraction of light emitted by the target foil will be collected by the camera. The roughly linear x -dependence and y -independence is caused by one side of the foil being closer to the optical system than the other. This is discussed further in Section 4.1.2.

where $C_e = \iint_{-\infty}^{\infty} b(x, y) e(\mathcal{O}, \mathcal{E}, x, y, d) dx dy$ is the total probability of a photon being collected by the optical system.

Equation 3.14 was calculated assuming that x and y are continuous. It is modified to operate on discrete pixels $\{\nu\}$ by integrating b and b' over the beam-space interval covered by each pixel ν . The pixels of I are $I_\nu = g \iint_{\nu} b(x, y) dA$ where g is a conversion factor (with units of pixel intensity) that is proportional to the total beam intensity and imaging parameters such as gain and exposure time into account. The pixels of I' are

$$\begin{aligned}
 I'_\nu &= \iint_{\nu} b'(x, y) dA = \frac{g \iint_{\nu} b(x, y) e(\mathcal{O}, \mathcal{E}, x, y, d) dA}{C_e} \\
 &\approx \frac{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d) g \iint_{\nu} b(x, y) dA}{C_e} = \frac{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d) I_\nu}{C_e}
 \end{aligned}$$

where it has been assumed that e is constant over the beam-space interval covered by ν . A method of reconstructing an image of the beam profile (up to a scaling factor C) is found by solving the definition of I'_ν for I_ν :

$$I_\nu = \frac{C_e}{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d)} I'_\nu \propto \frac{C}{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d)} I'_\nu = E_\nu I'_\nu. \quad (3.15)$$

The per-pixel correction factors E_ν implicitly depend on \mathcal{O} , \mathcal{E} and d and are updated from pre-computed tables as the state of the view screen unit changes. The overall intensity of I_ν is set by C (C_e is not a convenient scaling factor because it depends on b) and is set in a manner which is convenient to beam operators (discussed later). The choice of C is somewhat arbitrary because it does not impact the calculation of beam profile statistics.

If E_ν is known then applying equation 3.15 to an uncorrected image will produce an image in which the collection efficiency within each pixel is the same. An optical simulation was developed and used to measure the collection efficiency (e) for a variety of optical geometries and emission distributions. A validation of this process is carried out in Section 4.1.

Collection Efficiency Correction Implementation

The collection efficiency correction implementation is composed of two parts: pre-computed tables of position-dependent collection efficiency for a particular state (\mathcal{O} , \mathcal{E} and d), and an online algorithm which monitors the machine state and performs the image reconstruction.

The pre-computed tables are calculated by means of the ray-tracing simulation described in Section A.1. The simulation is initialized by selecting an optical geometry and light emission distribution. Next, a grid of points is defined on the front-face of the target foil. Each point becomes the source of a count of n rays which are propagated individually through the optics before being intersected with the plane of the camera CCD sensor. Rays which do not reach the sensor are removed from the tally; those that remain are binned into the CCD pixels, producing a simulated image I_μ . For this operation, the size of the CCD is artificially increased through the introduction of a margin of virtual pixels. This allows the simulation to capture the full point-spread-function of points which lie partially beyond the field-of-view boundary.

The collection efficiency of the optical system for a grid point (x_i, y_i) on the target

foil is estimated as the ratio of the total count of rays which are binned on the CCD sensor to the number of rays emitted:

$$\hat{e}(\mathcal{O}, \mathcal{E}, x_i, y_i, d) = \frac{\sum_{\mu} I_{\mu}(x_i, y_i)}{n}$$

where $I_{\mu}(x_i, y_i)$ is the count of rays emanating from (x_i, y_i) and landing in pixel μ .

The table of collection efficiency estimates, $\{\hat{e}(\mathcal{O}, \mathcal{E}, x_i, y_i, d)\}$, and simulation parameters are stored in a file. Storing the simulation parameters allows run-time validation of the software configuration. Binned PSF information for the target foil points is zero-suppressed, sorted and stored in a separate file. Files belonging to a particular optical configuration and light emission distribution are placed in a common directory.

These tables of collection efficiency estimates are used by the online collection efficiency correction algorithm to create the per-pixel correction values E_{ν} . For each type of target foil currently installed in the target actuator the contents of files corresponding to the particular optical geometry and light emission distribution are validated and loaded into data structures.

The contents of E needs to be updated when a new target is inserted, the iris diameter changes or the beam energy changes (in the case of OTR light); therefore, a few properties of the accelerator diagnostic state must be available in real time. Read-access to these parameters is provided by monitoring EPICS records using the channel access protocol. When one of the records is updated (or when the software is initially loaded), a callback function within the efficiency correction processor is executed, calculating E with the following algorithm:

1. Read the current iris diameter, d , and
2. select the data structures corresponding to the currently inserted target and beam energy.
3. For each pixel, ν , in the input image:
 - (a) transform the pixel coordinates into beam-space using Eq. 3.5;
 - (b) perform a three-dimensional linear interpolation of the data structures to estimate the collection efficiency $\hat{e}(\mathcal{O}, \mathcal{E}, x_{\nu}, y_{\nu}, d)$ at the given beam-space location;

- (c) calculate the normalized efficiency correction table entry $E_\nu = \frac{\hat{e}(\mathcal{O}, \mathcal{E}, 0, 0, d)}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d)}$ where (x_ν, y_ν) is the centroid of the beamspace region represented by pixel ν Equation 3.5.

As mentioned previously, the overall magnitude of E_ν is set by a scaling constant C and does not impact the calculation of beam profile statistics. For that reason, the C term is set to produce an image intensity that reflects the overall collection efficiency. This is accomplished by setting C equal to the estimated efficiency at the centre of the foil. If it was left at $C = 1$, the large values $1/\hat{e}(\mathcal{O}, \mathcal{E}, x, y, d) \approx 10^4$ would produce pixel values that overflow 8 or 16-bit storage types. Based on this choice, the overall image intensity reflects the current state of the mechanical iris. This is shown with the following calculation. If I_1 and I_2 are images of the same beam taken at iris diameters d_1 and d_2 ($d_1 < d_2$), corrected with the *non-normalized* tables

$$E_{1,\nu} = \frac{1}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_1)}$$

and

$$E_{2,\nu} = \frac{1}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_2)}$$

then

$$\frac{I'_{2,\nu}}{I'_{1,\nu}} = \frac{E_{2,\nu} I_{2,\nu}}{E_{1,\nu} I_{2,\nu}} = \left(\frac{I_{2,\nu}}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_2)} \right) \left(\frac{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_1)}{I_{1,\nu}} \right) \approx 1$$

because the fractional increase in pixel intensity, $\frac{I_{2,\nu}}{I_{1,\nu}}$, is roughly equal to the fractional increase in collection efficiency: $\frac{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_2)}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_1)}$.

When $I_{2,\nu}$ and $I_{1,\nu}$ are corrected with normalized correction table entries (Fig. 3.23) the fractional increase in pixel intensity becomes

$$\frac{I'_{2,\nu}}{I'_{1,\nu}} \approx \frac{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_2)}{\hat{e}(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_1)} \quad (3.16)$$

and this provides a visual indication to the operator that the iris is functional because the image intensity will increase as the iris is opened.

3.5.5 Calculation of Beam Profile Statistics

Beam profile statistics are calculated for both the raw camera image and the final reconstructed image of the beam profile.




Figure 3.23: The normalized collection efficiency table E for scintillation light passing through the ELBT (low-energy) optical configuration with an iris setting of 16 mm.

For the case of the raw camera image, a histogram of pixel intensities is calculated and used to check for image saturation. An online algorithm checks for saturation by calculating the fraction of pixels at maximum intensity and comparing the result against a user-defined threshold S_{thresh} (0.05, by default).

If the fraction of pixels at maximum intensity is calculated as the ratio of pixels at maximum intensity to the total number of pixels on the CCD sensor (eg. $\frac{b_{i_{max}}}{n_w \times n_h}$, where $b_{i_{max}}$ is the height of the last histogram bin) then picking a suitable S_{thresh} will be problematic. Lower values of S_{thresh} will increase the likelihood of images of a large beam or images containing many “stuck” pixels as incorrectly being flagged as saturated. On the other end, larger values of S_{thresh} will increase the likelihood of saturated images of a small beam to be labelled as unsaturated. To reduce the rate of these errors, only pixels with an intensity greater than the operator-defined parameter $S_{I_{min}}$ (128, by default) are used to calculate the fraction. Thus, the saturation check becomes

$$\frac{b_{i_{max}}}{\sum_{i=i_{min}}^{i_{max}} b_i} < S_{thresh} \quad (3.17)$$

where i_{min} is the bin number containing $S_{I_{min}}$. When the saturation check fails (when the image is saturated) an alarm is presented to the operator.

Beam profile statistics are extracted from the final reconstructed image by a purpose-built image processing plugin. Both raw and central bivariate moments are calculated up to the second order. Also calculated is the correlation coefficient ρ_{XY} . The plugin is given access to the current device calibration, allowing all calculations to be done in the beam-space coordinate system. These statistics are exported into the EPICS database for use in beam monitoring and tuning algorithms.

3.5.6 Image Styling

Styling, in the form of false colouring and graphical overlays, may be applied to the raw or reconstructed image. This gives the operator additional tools for beam-tuning and diagnostics.

False Colouring

The monochromatic nature of the camera images makes it difficult to visually distinguish pixels with similar intensities. This lack of contrast can be improved by falsely colouring the image. An existing image processing plugin from the `areaDetector` package allows the modification of the image colour format while optionally applying one of two pre-existing false-colour maps. This plugin was extended by adding new false-colour maps and a user-positionable legend. The new color maps were chosen based on the following two design principles:

- low-intensity pixels should appear dimmer than higher-intensity pixels and
- the mapping between monochrome and coloured pixel intensities should be one-to-one (no colours are repeated).

These design principles ensure that no artifacts (such as those in Figure 3.24) are added unnecessarily to the image when it is falsely coloured. Figure 3.25 shows four of the eight new colour maps. While the Cyan, Magenta, Yellow and Key (CMYK) and rainbow colour maps shown in figures 3.25a and 3.25d violate the first design principle they offer the advantage of having high contrast over the entire intensity range. This would be useful for visualizing beam halo, for example.

Graphical Overlays

Graphical overlays display additional information to the beam operator. Two types of graphical overlays are available: coordinate axes and crosshairs.




Figure 3.24: Demonstration of a poorly designed colour map. An operator may perceive a series of flat bands of beam intensity when, in reality, the gradient of the test image changes slowly.




Figure 3.25: Four new colour maps.

Coordinate axes are available for both the raw and beam-space image. Axes and tick mark colour are customizable through the operator interface. An optional 3-D “shadowed” appearance allows the axes to be visible regardless of image intensity. Coordinates axis ticks are presented in units of pixels for the raw image and in units of mm for the beam-space image.

Crosshairs can be placed in a static position or may track the beam centroid in real time. One use of the crosshair overlay feature is “degaussing” a solenoid magnet during device commissioning. A solenoid magnet adjusts the spread of the electron beam by applying a transverse force to the constituent particles. As the beam spread is adjusted the beam centroid may also change as a consequence. The position drift can be cancelled by pairing the solenoid with a pair of dipole magnets. The compensating force is found by repositioning the beam to the original location with the dipole magnets.

This procedure is streamlined by the use of the crosshair graphical overlay. The

beam centroid is visualized by activating the centroid tracker and the initial beam position is marked with a static crosshair. As solenoid current is adjusted and the beam drifts it is returned to the original position by lining up the centroid tracker with the static crosshair.

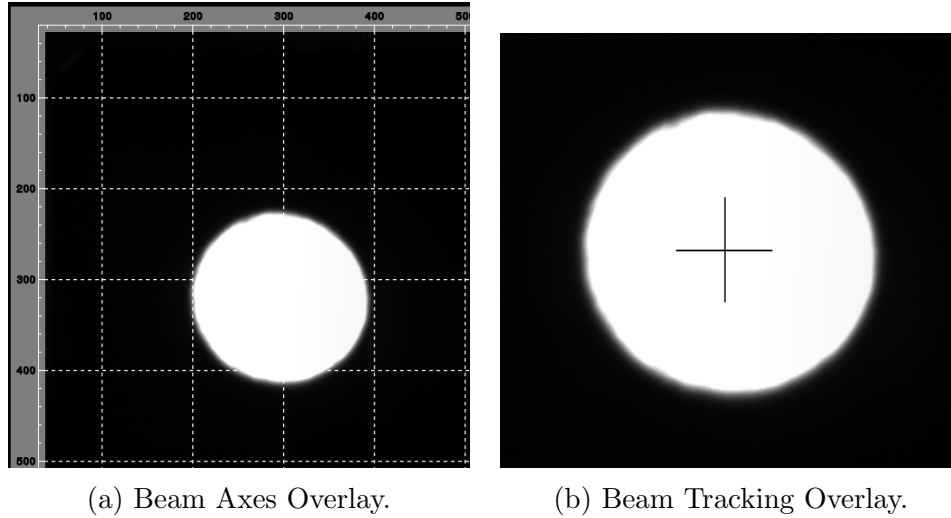


Figure 3.26: Beam Styling Overlays. These features present additional information to the beamline operator.

Chapter 4

Experiments

4.1 Collection Efficiency Validation

The collection efficiency (Sec. 3.5.4) is a function of the optical pathway, emission distribution and primary emission location. If the beam-profile measurements are to be trusted, then the corrections applied during image reconstruction must be validated. An experiment to validate the collection efficiency of the optical system operating with scintillating targets was designed and two analyses were performed on the data collected during the experiment. One analysis measured the relative increase in collection efficiency as the mechanical iris was opened; the other attempted to quantify position dependence.

4.1.1 Data Collection

An experiment to perform an “iris scan” was carried out in the VECC testing area by an expert operator. An iris scan is a procedure where the beam is held constant while the mechanical iris is “scanned” through different aperture sizes. At each aperture size, a set of data is collected. If the beam is stable, this allows the affect of the iris aperture to be measured.

A stable set of beam parameters was found by monitoring the average beam current with a Faraday cup. The beam was defocussed with a solenoid magnet, allowing a larger portion of the target foil to scintillate. The iris was opened to the maximum aperture and the camera exposure was adjusted so that the maximum pixel value within the raw image was just below the saturation level. For each iris diameter d_i within the range 12 mm to 32 mm in steps of 4 mm the following procedure was

carried out:

- temporarily disable the beam with the cathode bias voltage,
- capture an image of the background light (such as that produced by the hot electron gun cathode),
- enable background light subtraction using the recently collected image,
- re-enable the beam and
- record five data sets of the beam profile at the current iris setting.

This process was repeated for three different view screen units in the low energy section of the e-Linac prototype: VS1A, VS1B and VS3.

4.1.2 Analysis

If the true beam profile and light distribution were known then comparing a reconstructed image of the true beam with one produced by an optical simulation would be possible. The difference in intensities between the two images would be caused by the collection efficiency of the true optical system. Unfortunately, producing a particular beam distribution is very difficult due to the large number of parameters involved. Verifying the distribution poses additional difficulties. A slit scanner or wire scanner can measure marginal beam distributions but the only full profile monitor is the view screen system. Using the view screen system to validate the beam distribution prior to measuring the collection efficiency would introduce bias into the analysis.

Both parts of the analysis procedure were performed offline with Mathematica. Images taken under similar conditions (identical iris diameter and view screen device name) were averaged to remove transient beam fluctuations.

Measuring the Collection Efficiency Ratio

Measuring the ratio of the collection efficiency for a particular state to that of another state with a different iris aperture diameter does not require knowledge of the true beam profile but relies on being able to hold the beam profile constant over time and confidence in the mechanical iris control. The iris aperture was “scanned” through different sizes and it was expected that the collection efficiency of the system would increase with the diameter of the mechanical iris.

An image of the beam I_i , was collected at iris diameter d_i and to it the first two stages of the image reconstruction algorithm (geometry and magnification correction) were applied. All other things held constant, the ratio of the intensity of pixel ν from the collected image to that of the same pixel in a reference image (I_1) is mostly due to the change in collection efficiency at the corresponding beam-space location :

$$\frac{I_{i,\nu}}{I_{1,\nu}} \approx \frac{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_i)}{e(\mathcal{O}, \mathcal{E}, x_\nu, y_\nu, d_1)} \quad (4.1)$$

where d_1 is the diameter of the mechanical iris during acquisition of the reference image.

The collection efficiency ratio (Eq. 4.1) was calculated for each image and the images taken with the smallest iris diameter setting $d_1 = 12$ mm were used as the baseline of the comparison.

The resulting plots of the efficiency ratio of each pixel were compared to similar data produced with the collection efficiency maps (Fig. 4.1). It was observed that the simulated efficiency data was very similar to the results of measurement when the iris aperture was smaller than 28 mm in diameter. At larger iris diameters the measured increase in collection efficiency is lower than that which is predicted by simulation.

The divergence at large iris diameters could be due to the point spread function. As the iris aperture opens past 20 mm, most of the additional light is distributed at the edges of the PSF. This isn't a problem when the beam profile is uniform because convolving the beam distribution with a slowly-changing PSF will not introduce large changes in the beam-space image. In the case of a non-uniform beam-profile and non-isotropic PSF, the measured collection efficiency will also have been convolved by the PSF. This has the effect of artificially decreasing the measured collection efficiency near the beam centroid and increasing it near the edges.

Measuring Scintillation Collection Efficiency Position Dependence

The previous portion of the collection efficiency study validated the relative increase in collection efficiency as the iris is opened. This portion of the study attempts to validate the position dependence of $e(d, x, y)$. Validating the relative increase in collection efficiency as the iris is opened demonstrated that the ratio in equation 4.1 agreed with simulation. Unfortunately, it did not predict anything about the particular shape of $e(d, x, y)$. This portion of the analysis attempts to do so. A major hurdle in validating the shape of e is that the true beam profile is unknown. If it were




Figure 4.1: Comparing the Measured Collection Efficiency to Simulated Data. The ratio $\frac{I_{i,\nu}}{I_{1,\nu}}$ is a good approximation to the ratio of collection efficiency present during the acquisition of each image.

known then it would be possible to predict the measured beam distribution $b'(x, y)$ and calculate e (up to a scaling factor) with the equation

$$b'(x, y) = \frac{b(x, y) e(d, x, y)}{\iint_{\infty} b(x, y) e(d, x, y) dx dy}$$

$$\implies e(d, x, y) = \frac{b'(x, y)}{b(x, y)} \iint_{\infty} b(x, y) e(d, x, y) dx dy$$

As mentioned previously, determining the true beam profile is difficult for a number of reasons. This means that an alternative method of measuring e must be developed.

Simulations of scintillation light passing through the low-energy optics suggest that the collection efficiency has minimal dependence on y , and is roughly linear in x (Fig. 4.2) for all values of the iris diameter.

Based on these observations, a model of e was produced:

$$e(d, x, y) = e_{\mu} (m(x - \mu_x) + 1) \quad (4.2)$$

where e_{μ} is the collection efficiency at μ_x , the x centroid of $b(x, y)$, and m characterizes




Figure 4.2: The change in collection efficiency in lines across beam space.

the (potentially) linear dependence of e on x .

This form of e predicts the measured beam distribution

$$b'(x, y) = \frac{b(x, y) e(d, x, y)}{\iint_{\infty} b(x, y) e(d, x, y) dx dy} \quad (4.3)$$

$$= b(x, y) e(d, x, y). \quad (4.4)$$

The marginal distribution of b' in X is used in lieu of the two-dimensional distribution because e is assumed to be independent of y :

$$b'_X(x) = \int_{-\infty}^{\infty} dy b(x, y) e(d, x) = b_X(x) e(d, x) \quad (4.5)$$

Estimating the form of $b_x(x)$ will produce a distribution that can be compared against the measured data. An initial attempt of fitting a normal distribution to a sample of the VLBT:VS1A data indicated a problem with this technique. The correlation between the m and μ parameters is large. To remove these parameters from the fit, a relationship between M'_n (the n th raw moment of b'_X) and M_n was

determined

$$M'_n = \int_{-\infty}^{\infty} dx x^n b_X(x) e(d, x) = mM_{n+1} + (1 - m\mu_x) M_n. \quad (4.6)$$

The raw moments M'_n are measured from the collected data, leaving one free variable due to the M_{n+1} term in equation 4.6. An additional constraint was acquired based on the assumption that $b_X(x)$ is symmetric about μ_x . This causes all odd central moments of $b_X(x)$ to disappear, leading to the following equations:

$$M_n = - \sum_{j=0}^{n-1} \binom{n}{j} (-1)^{n-j} M_j M_1^{n-j} : n \text{ is odd.}$$

These extra conditions were used to solve Eq. 4.6 for M_1 and M_2 as functions of M'_n, a and m .

Based on observations of the data an ansatz was chosen for b_X (Eq. 4.9) and the single parameter model of b'_X was fit to each of the VLBT:VS1A images.

Another check was performed by estimating the change in efficiency based on the geometry of the ELBT optics and the theoretical scintillation emission distribution. The intensity of scintillation light emitted by the target foil is proportional to $\cos \theta d\Omega$ where θ is the polar angle of emission, relative to the surface normal of the target foil (Fig. 4.3) and $d\Omega$ is the solid angle subtended by the aperture of the mechanical iris.

The polar angle of the ray which passes through the centre of the first lens can be calculated via geometry:

$$\cos \theta = \frac{\hat{n} \cdot \vec{P\hat{L}}}{|\vec{P\hat{L}}|} = \frac{(P_x - L_x) \sin \theta_t - P_x \cos \theta_t \tan \theta_t}{|\vec{P\hat{L}}|} = \frac{-L_x \sin \theta_t}{R(x)} \quad (4.7)$$

where $L_x = 534 \text{ mm}$ is the Manhattan (rectilinear) distance between the centre of the target foil and the centre of the first lens, $\theta_t = -\pi/4$ is the rotation angle of the target foil and $R(P_x) = |\vec{P\hat{L}}| = \sqrt{(P_x - L_x)^2 + P_x^2 \tan^2 \theta_t}$ is the distance between a point on the foil and the centre of the first lens.

The solid angle $d\Omega = dA/R(x)^2$ decreases as $R(x)$ increases due to the size of the mechanical iris aperture dA being independent of the target foil position. Combining the expressions for $\cos \theta$ and $d\Omega$, and dividing by the intensity at $x = 0$ yields the

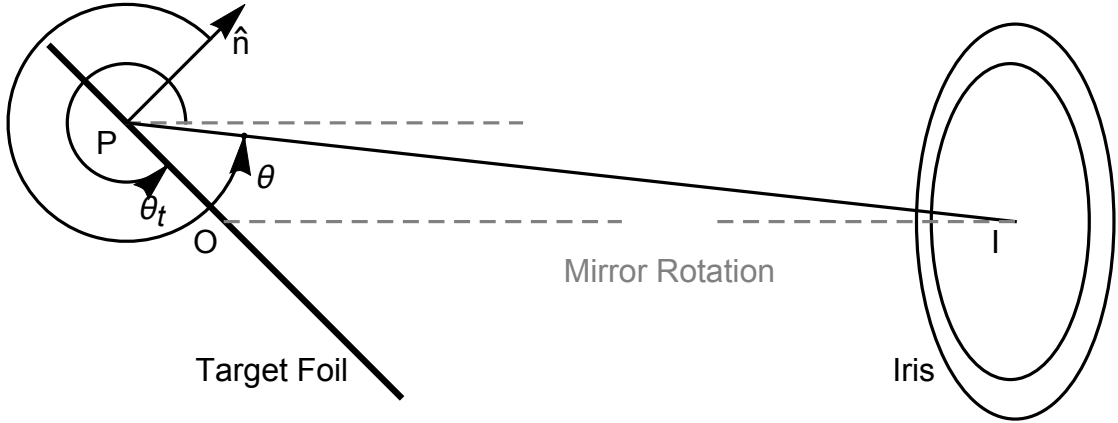


Figure 4.3: Estimating the Change in Collection Efficiency using Geometry.

relative intensity

$$\mathcal{I}(x) = \frac{\cos \theta \, d\Omega}{\cos \theta \, d\Omega|_{x=0}} = \left(\frac{R(0)^2}{R(x)^2} \right)^{3/2}$$

and differentiating the relative intensity yields the relative change in intensity:

$$m_{\mathcal{I}}(x) = \partial_x \mathcal{I}(x) = \frac{-3 (R(0)^2)^{3/2} (x(1 + \tan^2 \theta_t) - L_x)}{(R(x)^2)^{5/2}}. \quad (4.8)$$

The relative change in intensity as measured from the iris scan data, the simulated collection efficiency maps and calculated from the geometry of the optics and theoretical emission distribution are shown together in Figure 4.4.

$$b_X(x) = \begin{cases} \frac{3((\sqrt{5}\sigma)^2 - (x - \mu_x)^2)}{4(\sqrt{5}\sigma)^3} & |x - \mu_x| \leq \sqrt{5}\sigma \\ 0 & \text{otherwise.} \end{cases} \quad (4.9)$$




Figure 4.4: Simulated and Measured Change in Collection Efficiency.

4.2 Evaluation of the Quality of Beam Profile Measurements

A study was carried out to determine the quality of beam profile measurements taken by the view screen system. The results of the study were used to evaluate the satisfaction of the functional diagnostic requirements Req. 6 through Req. 11. While evaluating the system in-situ would have been ideal, producing reference beams with well-defined properties (such as position and shape) would be difficult due to the large number of degrees of freedom in the accelerator. To work around this the optical simulation was used to produce raw images of a variety of beams, as well as images of the calibration target, under different configurations of the optical components.

The first part of the study measured the systematic bias present in ideal (with perfect focus, component alignment and calibration) versions of both the low-energy and high-energy optical component configurations. The second part of study attempted to improve accuracy by accounting for the leading sources of systematic bias. The third part quantified the loss of precision occurring as a result of uncertainty in the camera focusing procedure. The fourth, and final, part of the study evaluated the satisfaction of the functional diagnostic requirements.

When necessary, raw beam images were reconstructed with a Mathematica implementation of the image reconstruction algorithm in lieu of the EPICS-based implementation. This was done due to the ease of validating the analysis and the ability to efficiently manipulate the large amount of data and produce figures of the results. Four of the five independent beam profile properties have accuracy and precision requirements. These properties are the beam centroid (M_{10}, M_{01}) and the beam width in each dimension ($\sqrt{U_{20}}$ and $\sqrt{U_{02}}$). No requirements on the fifth beam profile measurement, ρ , was documented.

4.2.1 A Measurement of the Systematic Bias Present in Ideal Optical Configurations

The initial part of the study measured the systematic bias present in configurations of optical systems where the camera focus and component alignment was perfect, background sources of light were not present, the distribution of light emitted from the surface of the target foils was ideal and the calibration procedure was executed without any error. For each of the two ideal optical configurations (henceforth named

ELBT for the low-energy camera boxes and EHBT for the high-energy camera boxes) a number of optical simulations of various beams imaged under various imaging conditions (Tables 4.1 and 4.2) was carried out.

Beam Shape	Gaussian
Beam Widths (σ)	0.25, 0.5, 1, 1.5, 2, 2.5, 3, 3.5 and 4 mm
Beam X-Positions (μ_x)	-15 to 15 mm in steps of 3 mm
Beam Y-Positions (μ_y)	-24 to 24 mm in steps of 3 mm
Iris Diameters	12, 20 and 28 mm
Light Emission Distribution	Scintillation

Table 4.1: Beam Parameters and Imaging Conditions for the Ideal ELBT Optical Configuration. The beam centroid was constrained to fall within the boundary of the target foil: -17.68 to 17.68 mm in the x-dimension and -25 to 25 mm in the y-dimension.

Beam Shape	Gaussian
Beam Widths (σ)	0.2, 0.4 and 1.0 mm
Beam X-Positions (μ_x)	-8 to 8 mm in steps of 2 mm
Beam Y-Positions (μ_y)	-12 to 12 mm in steps of 2 mm
Iris Diameters	12, 20, 28 and 36 mm
Light Emission Distribution	Scintillation

Table 4.2: Beam Parameters and Imaging Conditions for the Ideal EHBT Optical Configuration. The beam centroid was constrained to fall within the boundary of the target foil: -8.84 to 8.84 mm in the x-dimension and -12.5 to 12.5 mm in the y-dimension. The range of the iris diameter was increased to include 36 mm because the diameter of the iris is expected to be larger to collect more light when imaging targets in the higher energy sections of the beamline.

The raw data was reconstructed and profile measurements were extracted. Two aspects of the analysis are presented: observations of the difference between the measured beam centroid and the mean position of the simulated Gaussian beam and observations of the difference between the measured beam widths (in both dimensions) and the width of the simulated Gaussian beam.

Centroid Bias in the ELBT and EHBT Optical Configurations

The vector difference between a measured beam centroid and the mean position of the corresponding simulated Gaussian beam was visualized using an arrow originating at the mean position of the simulated beam, pointing in the direction of the vector

difference and having a magnitude proportional to the magnitude of the difference. Two such representative graphics (showing the vector difference observed for the ELBT optical configuration collecting scintillation light) are shown in Figure 4.5. It was observed that the vector difference

- at the inner portion of the foil
 - was directed radially outward,
 - increased in magnitude as the beam moved away from the foil centre,
 - was minimally dependent on the beam width and
 - increased in magnitude as the iris was opened;
- and at the edge of the foil
 - was directed radially inward (generally) and
 - had a much larger magnitude than the difference at the inner portion of the foil.

It was hypothesized that there are two leading factors contributing to the difference between the reconstructed and simulated centroid (henceforth called centroid *bias*): beam profile clipping and PSF blurring.

Beam profile clipping occurs when a portion of the beam is not imaged by the camera because it does not intersect with the target foil. If a portion of the beam is not collected then the centroid measurement will be biased toward the collected portion of the beam profile: toward the middle of the foil. Additionally, the magnitude of the centroid bias will increase as a larger portion of the beam is clipped by the beam moving closer to the edge of the target foil. These characteristics agree with the observations of the centroid bias near the edge of the target foil. From Figure 4.5a, it is observed that the magnitude of the centroid bias of a 3 mm Gaussian beam (which is expected for the ELBT optical configuration) is on the order of 100 μm when the beam is roughly 5 mm to 6 mm from the edge of the target foil. What was not shown (because it cluttered the Figure) is that the magnitude increased to 1 mm to 2 mm when the beam centroid is roughly 1 mm from the edge of the target foil. A direct consequence of beam profile clipping is the potential loss of usable target foil area. If the true beam distribution was known then removing the bias could be accomplished by calculating the amount of missing beam and adding it to the centroid

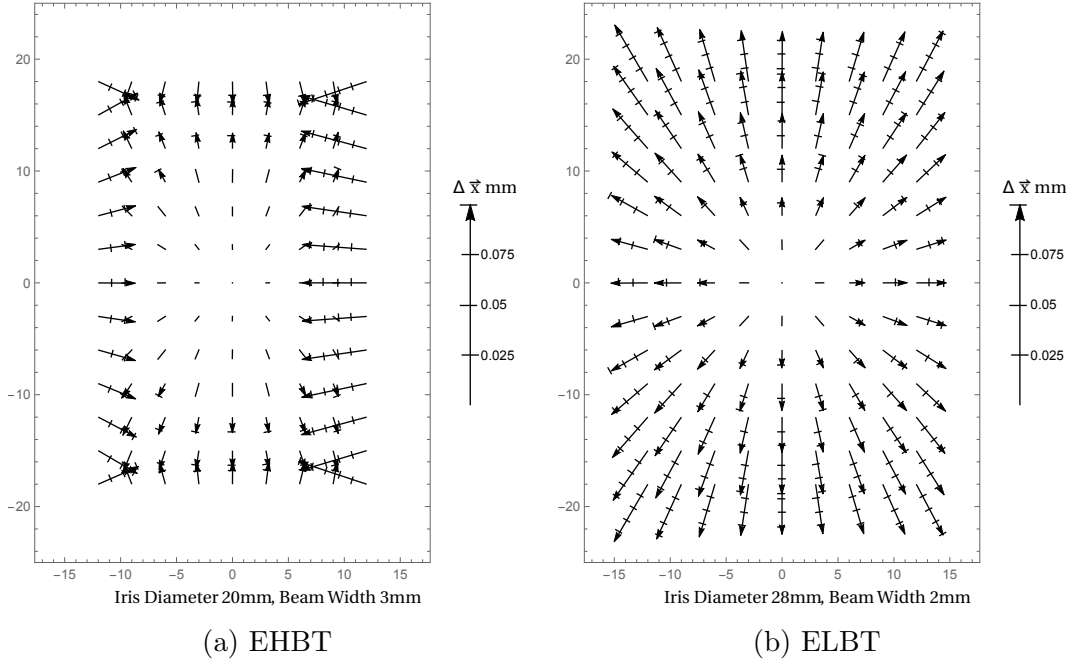


Figure 4.5: Systematic bias in the beam centroid measurement of a Gaussian shaped beam incident on a scintillating target imaged by an ELBT camera box. The vector difference between a measured beam centroid and the mean position of the corresponding simulated Gaussian beam was visualized using an arrow originating at the mean position of the simulated beam, pointing in the direction of the vector difference and having a magnitude proportional to the magnitude of the difference. The results of two different iris diameters and beam widths are shown.

measurement. Unfortunately this is difficult (or impossible) in practise given that the true beam distribution is not known exactly and can vary wildly in shape.

Point-spread-function blurring occurs when light emitted by a single point on the target foil intersects the camera CCD sensor at different places. The image processing algorithm made no attempt at PSF removal; therefore, effects of the PSF are expected to be present in the reconstructed image. Characteristics of the PSF for the e-Linac view screen imaging system were determined in previous work [8]: it was directed radially outward, increased in magnitude as the iris was opened and spread outward as it moved outward from the beamspace origin. These characteristics agree with the observed behaviour of the centroid bias within the inner portion of the target foil. From Figure 4.5 it is observed that the magnitude of the bias caused by PSF blurring is smaller than that of beam profile clipping. When the iris is opened to 28 mm the magnitude of the bias increases to 100 μm near the edge of the foil. The magnitude quickly decreases to less than 50 μm when the iris is reduced to 20 mm.

The centroid bias for the collection of scintillation light by EHBT optical configuration was observed to be similar to the ELBT optical configuration with the exception of having a much smaller bias caused by PSF blurring.

Beam Width Bias in the ELBT and EHBT Optical Configurations

The difference between the measured beam width (square root of the variance) in both dimension and the width of the simulated Gaussian beams may be best represented in a manner analogous to the corresponding functional requirement. Req. 9 states that “The appropriate requirement on beam size and structure within the beam profile is resolution = 5 % of the anticipated r.m.s beam size”. For this reason the figure of merit studied was the relative change in beam width $\frac{\sqrt{U_{20}} - \sigma}{\sigma}$ where U_{20} is the measurement of the second central moment (variance) of the reconstructed beam profile in the x-dimension. A similar quantity exists for the variance in the y-dimension. It may be more natural to consider the bias in the variance ($U_{20} - \sigma^2$ or $U_{02} - \sigma^2$). In fact, a method of improving the beam width measurements by removing the bias from the variance is discussed later.

The relative change in the horizontal beam width was calculated from the simulated images of Gaussian shaped beams being intercepted by a scintillating target in the ELBT (Fig. 4.6) and EHBT (Fig. 4.7) optical configurations. It was observed that the relative change

- at the inner portion of the foil
 - was positive,
 - increased as the beam centroid moved away from the centre of the target foil and
 - decreased as a function of beam width;
- and decreased substantially, becoming negative, at the edge of the foil.

It was hypothesized that the “bias” in the relative change in beam width was caused by the same leading factors which contributed to the centroid bias: beam profile clipping and PSF blurring. Beam profile clipping at the horizontal (vertical) edges of the target foil would cause the beam width to be underestimated in the horizontal (vertical) direction. PSF blurring would cause the reconstructed beam profile to be wider than the true profile, leading to an overestimation of the beam

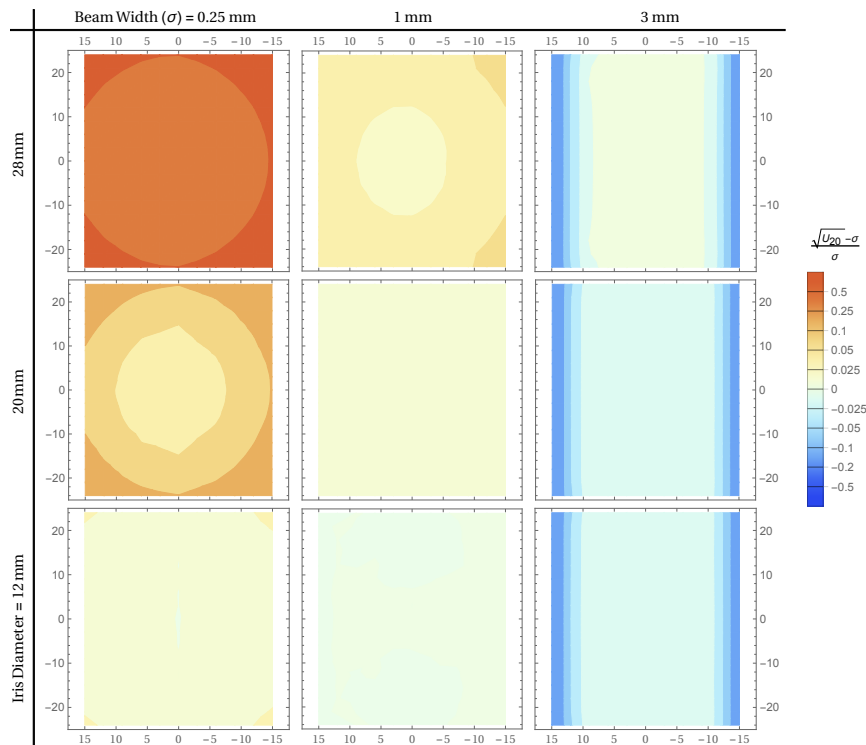


Figure 4.6: Relative change between the horizontal width ($\sqrt{U_{20}}$) of the reconstructed beam profile and the true width (σ) of a Gaussian beam intercepted by a scintillating target foil and imaged with an ELBT optical configuration.

width. The manner in which the reconstructed profile is widened depends on the shape of the PSF. It is expected that the spread of the PSF extends radially outward and is proportional to the distance to the centre of the target foil; therefore, a similar rotational symmetry is expected in the relative change in beam width away from the foil edges.

For the ELBT optical configuration the magnitude of the PSF blurring and profile clipping effects are dictated by the width of the imaged beam and the diameter of the mechanical iris. The relative change caused by PSF blurring is much more dramatic when the iris is opened and the beams are small; ranging from 0.25 to 0.5 for a $\sigma = 0.25$ mm beam and 28 mm iris setting. Fortunately, 0.25 mm wide beams are not expected in the lower energy section of the accelerator and the relative change in beam width of more typical beam sizes ($\sigma \approx 1$ mm to 3 mm) is within the range of -0.025 to 0.05 for most locations within the inner portion of the foil. The relative change due to beam profile clipping is not visible on the plots of the $\sigma = 0.25$ and 1 mm data because the extents of the simulated beam centroids (± 15 mm) did not allow




Figure 4.7: Relative change between the horizontal width ($\sqrt{U_{20}}$) of the reconstructed beam profile and the true width (σ) of a Gaussian beam intercepted by a scintillating target foil and imaged with an EHBT optical configuration.

substantial clipping to occur for small beam sizes. The effect of profile clipping on the $\sigma = 3$ mm beams is apparent: the relative change in width begins to drop when $|M_{10}| > 10$ mm and falls below -0.10 at the edge of the foil.

The behaviour of the relative width is similar when measured from data collected by the EHBT optical configuration. There is one difference: the magnitude of the PSF effect is below 0.025 when beam of width greater than 0.2 mm are imaged when the diameter of the iris is less than 36 mm.

4.2.2 Improving the Quality of Beam Profile Measurements by Accounting for the Systematic Bias

The evaluation of the systematic bias present in the ideal ELBT and EHBT optical configurations (when imaging light produced via scintillation) demonstrated that systematic bias was present in measurements taken from the reconstructed beam profile. It was hypothesized that the two leading contributions to the bias were PSF blurring and beam profile clipping. The second part of this analysis determined if the accuracy of beam profile measurements could be improved by accounting for the systematic bias.

PSF Blurring

Accounting for the bias introduced through the PSF would improve the accuracy of the view screen system measurements when the beam size is small and the iris diameter is wide. The impact of blurring can be estimated by evaluating the approximate effect of the point spread function $p(x, y)$ on the beam profile probability

distribution $b(x, y)$. The probability distribution of the reconstructed beam profile (with magnification and efficiency corrections removed) is equal to the convolution of p and b :

$$b'(x, y) = p(x, y) * b(x, y) = \iint dx' dy' p(x' - x, y' - y)b(x', y'). \quad (4.10)$$

Calculating the mean, variance and correlation coefficient of Eq. 4.10, and making a zeroth order “expansion” of p about the centroid of b' yields the following relationships between the $(n, m)^{\text{th}}$ raw (M_{nm}) and central (U_{nm}) moments of b, b' and p :

$$M_{b,10} = M_{b',10} - M_{p,10}, \quad (4.11)$$

$$M_{b,01} = M_{b',01} - M_{p,01}, \quad (4.12)$$

$$U_{b,20} = U_{b',20} - U_{p,20}, \quad (4.13)$$

$$U_{b,02} = U_{b',02} - U_{p,02} \text{ and} \quad (4.14)$$

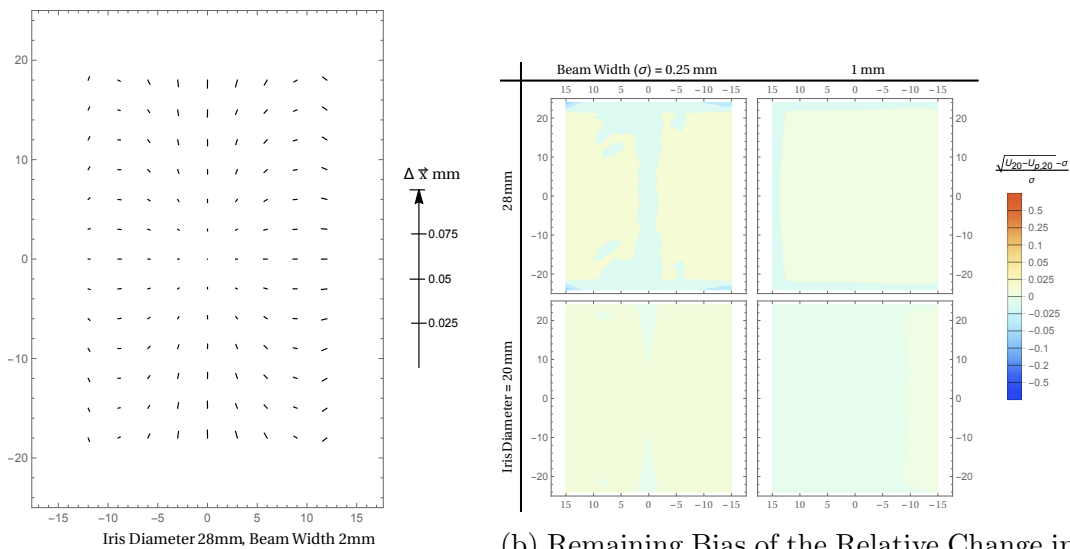
$$\rho_b = \frac{U_{b',11} - U_{p,11}}{\sqrt{(U_{b',20} - U_{p,20})(U_{b',02} - U_{p,02})}}. \quad (4.15)$$

These equations provide a means of removing any systematic bias caused by PSF blurring, provided that the properties of p are known.

Images of the ELBT and EHBT optical configuration PSFs were calculated for a range of iris diameters (8 mm to 40 mm in steps of 4 mm) and beam-space positions (an 11×11 grid of points) using the optical simulation. The images can be interpreted as a distribution where the probability of a photon being collected by pixel (u, v) is equal to the fractional amount of the total image intensity contained within. This probability mass function was converted to beamspace by applying the inverse of an ideal geometry transformation function g (Eq. 3.4) to the imagespace coordinates and then shifting the origin to be the point source of the rays used to simulate the PSF. From this distribution the relevant statistical properties ($M_{p,01}$, etc) were measured. The resulting data set was interpolated so that the properties of the PSF could be evaluated throughout beamspace and the full range of iris aperture sizes.

The beam profile measurements taken from the simulated data were corrected using the interpolated PSF properties (Fig. 4.8). The reduction of the bias which was present at the inner portion of the foil when imaging with the ELBT optical configuration and an iris diameter greater than 24 mm is substantial (Fig. 4.8a). Where the magnitude of the bias was once in the range of 0 μm to 100 μm , it was no

greater than $25\ \mu\text{m}$ throughout the inner portion of the target foil. An improvement in the measurement of the beam width was also seen for similar imaging conditions when the beam width was less than 1 mm (Fig. 4.8a). The magnitude of bias at the inner portion of the foil was reduced from 0.5 to less than 0.025 for all beam sizes greater than $\sigma = 0.25\ \text{mm}$.



(a) Remaining Centroid Bias (to be compared with Fig. 4.5b).

(b) Remaining Bias of the Relative Change in Beam Width (to be compared with the upper-left quadrant of Fig. 4.6).

Figure 4.8: Systematic bias in the beam centroid and relative change in beam width after the effect of PSF blurring is removed. This was calculated with the simulation of Gaussian shaped beams incident on a scintillating target imaged by an ELBT optical configuration.

It was observed that the reduction of the bias which is introduced when imaging with the EHBT optical configuration was not as dramatic due to the magnitude of the PSF blurring effect being much smaller than when imaging with the ELBT optical configuration. The effect of the bias when imaging OTR light was not investigated.

Beam Profile Clipping

The bias introduced by beam profile clipping is equal to the negative of the expected value of the measurement of interest in the region of beamspace *not* covered by the foil:

$$\text{bias} = -E[f(X, Y)] = \int_{|x| > 25/\sqrt{2}} dx \int_{|y| > 25} dy (-f(x, y)) b(x, y) \quad (4.16)$$

where $f(X, Y)$ is the measurement of interest (eg. X for the mean), $b(x, y)$ is the beam profile distribution and the limits of integration are based on the cross-section of the target foil. Anything other than an estimate of the bias cannot be calculated due to $b(x, y)$ being unknown beyond the edges of the target foil; however, this effort was not wasted. The effect of beam profile clipping can be accounted for by estimating the degree of belief in the satisfaction of the functional diagnostic requirements. If a model beam distribution is chosen, such as a bivariate normal distribution, an estimate of the bias can be calculated based on the measurements of the unclipped portion of the beam profile. An alarm condition can be added to relevant EPICS records, activating when it is suspected that the functional diagnostic requirements are no longer met (such as when the beam strays too close to the edge of the foil).

4.2.3 Precision of the Systematic Bias Corrections

This part of the analysis studies how uncertainty in the camera focusing procedure leads to uncertainty in the measurement of statistical properties of the point spread function. These uncertainties reduce the precision of the systematic bias corrections. The camera focusing procedure is the process by which an operator determines the vertical camera offset that maximizes the sharpness of raw images collected by the view screen. The natural limitation of fine-tuning the final camera position, along with uncertainty in the true location of the ideal position, combine to reduce the precision of calculations of the PSF properties.

The camera position is adjusted by turning a 0.7 mm pitch M4 screw and a separate set screw locks the final location in place. It was estimated that static friction limits the movement of the screw to 5° increments, leading to a positional uncertainty of $\sigma_y = \frac{0.7 \text{ mm}}{2\sqrt{3}} \frac{5^\circ}{360^\circ} = 2.8 \text{ }\mu\text{m}$.

The sharpest image is taken to be the one with the least Shannon entropy:

$$S = - \sum_i^N p_i \log_2 p_i \quad (4.17)$$

where p_i is fractional number of pixels containing intensity i and N is the count of discrete intensities (equal to $2^8 = 256$ for an 8 bit imager). The optical simulation was used to generate images of the calibration foil for both the ELBT and EHBT optical configurations (and their corresponding calibration target geometries) for a variety of camera positions (17 values within $\pm 200 \text{ }\mu\text{m}$) of the ideal camera position).

The entropy of each image was calculated and was plotted as a function of camera position (Fig. 4.9). The location of the minimum, as well as an uncertainty estimate (σ_S), were determined by fitting a quadratic form to the data.

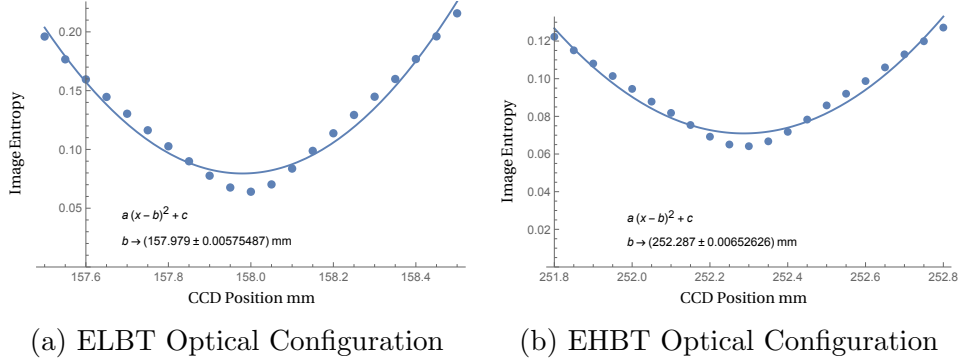


Figure 4.9: Camera Focusing by Minimizing the Entropy. The location corresponding to the minimum entropy in the image of a calibration target was found by fitting a quadratic curve to the simulated data.

Using the optical simulation, the properties of the point spread function were measured as a function of camera position (Fig. 4.10a), iris diameter, light emission distribution and optical configuration. The variance of each PSF property was then calculated based on the assumption that the final camera position was distributed normally about the minimum entropy position with a standard deviation equal to the sum of σ_y and σ_S in quadrature.

For the case of scintillation light travelling through a 24 mm diameter iris being collected by the ELBT optical configuration, the following observations were made. The standard deviation in the PSF x-centroid ($M_{p,10}$) was $0.25 \mu\text{m}$ at the centre of the foil and increased to $1.5 \mu\text{m}$ at the horizontal extents (Fig. 4.10b). The standard deviation of ($M_{p,01}$) was similar in magnitude but increased, instead, toward the vertical extents. Changing the diameter of the mechanical iris had little effect on the standard deviation of the measurement of either quantity. The standard deviation of both $U_{p,20}$ and $U_{p,02}$ was found to lie between $1 \mu\text{m}^2$ to $3 \mu\text{m}^2$: at a minimum in the middle of the target foil and increasing radially to the maximum of the range. Unlike $\sigma_{M_{p,10}}$ and $\sigma_{M_{p,01}}$, it was observed that the range of $\sigma_{U_{p,20}}$ and $\sigma_{U_{p,02}}$ depended on the diameter of the mechanical iris. The range of both quantities decreased to $0.1 \mu\text{m}^2$ to $0.4 \mu\text{m}^2$ for an iris diameter of 12 mm and increased to $5 \mu\text{m}^2$ to $9 \mu\text{m}^2$ for an iris diameter of 36 mm.

For the case of scintillation light being collected by the EGBT optical configuration




Figure 4.10: Uncertainty in the Measurement of the PSF X-Centroid. The x-centroid of the PSF shifts as the camera moves during the focusing procedure (4.10a). Uncertainty in the final camera position leads to uncertainty in the x-centroid bias correction (4.10b).

the following observations were made. The range of both $\sigma_{M_{f,10}}$ and $\sigma_{M_{f,01}}$ was $0.5 \mu\text{m}$ to $1.5 \mu\text{m}$ for all settings of the mechanical iris and no positional dependence was observed. Both $\sigma_{U_{f,20}}$ and $\sigma_{U_{f,02}}$ was observed to lie within the range of $0.01 \mu\text{m}^2$ to $0.04 \mu\text{m}^2$ when the iris was 12 mm in diameter. The range increased (roughly linearly) to $0.025 \mu\text{m}^2$ to $0.175 \mu\text{m}^2$ when the iris was 40 mm in diameter. That the magnitude of each quantity was much smaller than its ELBT counterpart was not unexpected; the PSF of the EHBT optical configuration was much smaller than that of the ELBT configuration.

The accuracy of these results is limited by how well the relationship between the Shannon entropy of the calibration target images and the camera offset is understood. The curves in Figure 4.9 are subject to the limitations of the optical simulation (an absence of dispersion and tracking of secondary rays produced via internal reflection) as well as focusing problems caused by component misalignment and the stability of the calibration light source. It is likely that these factors contribute entropy to the images of the calibration target and that (σ_S) has been underestimated. If σ_S was increased it is expected that the uncertainty in the measurement of each PSF property would increase as a result.

The precision of the results is constrained by the resolution of the CCD sensor, the resolution of the optical system and the precision of the optical simulation.

4.2.4 Satisfaction of Beam Profile Measurement Requirements

The final part of this analysis comments on the satisfaction of the functional diagnostic requirements. These requirements were introduced in Section 2.3.1 and are tabulated in Figure 4.3 for convenience. The results for both the absolute uncertainty (accuracy) and relative uncertainty (precision) requirements are presented.

Figure of Merit	Accuracy (μm)	Resolution (μm)
Beam Position	200	25
Beam Position (1 mA C.W. beam)	200	50
Beam Position (10 μA pulsed beam)	200	100
Beam Width (ELBT)	-	200
Beam Width (EMBT)	-	50
Beam Width (EHBT)	-	20

Table 4.3: Function Diagnostic Requirements of the View Screen System.

Accuracy

The accuracy of the beam profile measurements is limited by the ability to identify, measure and account for systematic bias. There are many sources of potential bias in the view screen system:

- image blurring caused by the point spread function,
- beam profile clipping at the edge of the target foil,
- misalignment of optical components,
- displacement of the beam targets and
- under or over-rotation of the calibration target (or target actuator).

Of these sources only the image blurring caused by the PSF and beam profile clipping are accounted for. Misalignment of optical components would be difficult to measure in-situ and is not expected to be a large source of bias.

A displaced calibration target will introduce bias into the measurement of the beam centroid. A horizontal displacement (movement through the $x - z$ plane) will bias the centroid measurement by the same amount; it will also introduce errors into the collection efficiency correction process. A vertical displacement will introduce

a similar bias into the centroid measurement but will not cause large errors in the efficiency correction process because the corrections for scintillation light are only weakly dependent on the y coordinate (Fig. 3.22). Misalignment of the calibration target is potentially identified by observing a correlated shift in the measured location of each calibration hole during the calibration process (although this is also consistent with an unaligned mirror).


Under or over-rotation of the calibration target (or target actuator) will cause both the measured x-centroid and the horizontal width of the beam to be under or over-estimated. The relative bias will be $\cos(\delta\theta) + \sin\delta\theta - 1$ where $\delta\theta$ is the difference between the ideal target rotation angle and the true angle. Accuracy of the collection efficiency procedure will also be affected because the correct target positioning was used assumed when producing the corrections. Incorrect angular position of the calibration target is identified by observing correlated scaling in the horizontal component of the measured location of each calibration hole during the calibration procedure.

The satisfaction of the beam centroid accuracy requirement was determined based on the simulated data described in Section 4.2.1. No attempt to quantify or correct for bias introduced by potential misalignment of the optical components and target actuator was made, although the potential for a large bias exists. The requirement was considered to be satisfied when the distance between the centroid of the reconstructed beam profile and the centroid of the simulated Gaussian beam was less than 200 μm .


Accuracy of the centroid measurement for the ELBT optical configuration (imaging scintillation light) is shown in Figure 4.11a. With the bias from the point spread function removed, and not including the bias from the misalignment of optical components and the target actuator, the accuracy requirement is satisfied for any beam size less than 3 mm provided that the beam is located at the inner portion of the target foil. The accuracy requirement is not satisfied when the beam profile is clipped at the edge of the target or the image of the beam is clipped at the extents of the camera CCD sensor. Similarly, the accuracy requirement is satisfied throughout the inner portion of the target when imaged by the EHBT optical configuration except that the region of dissatisfaction is smaller because the beams are expected to be smaller.

Although a formal requirement on the accuracy of beam width measurements was not present in the requirements document it is worth mentioning that once the bias from the point spread function was removed the remaining bias was low (-0.25% to

0.25% of the true beam width) at the inner portion of the target foil and suffered from profile clipping at the edges.



(a) ELBT Optical Configuration



(b) EHBT Optical Configuration

Figure 4.11: Centroid Measurement Accuracy for the ELBT and EHBT Optical Configurations. The required accuracy of the beam centroid measurement is satisfied at the inner portion of the target for the expected beam sizes. The requirement is not satisfied (red area) when the beam profile is clipped at the edge of the target. The results are shown for simulated Gaussian-shaped beams.

Precision

A few sources of uncertainty contribute to reduce the precision of beam profile measurements:

- limitation of sensor resolution in the camera CCD,
- surface roughness of target foils and
- precision of the systematic bias corrections.

Loss of precision of beam profile measurements due to sensor resolution (or pixel size) is not expected to be dramatic because the amount of area covered by a single

pixel ($3.8 \times 10^3 \mu\text{m}^2$ to $5.0 \times 10^3 \mu\text{m}^2$ for the ELBT and $0.94 \times 10^3 \mu\text{m}^2$ to $1.04 \times 10^3 \mu\text{m}^2$ for the EHBT optical configuration) is much smaller than the expected beam sizes; however, the ability to resolve beam features (Req. 9) which are $20 \mu\text{m}$ to $50 \mu\text{m}$ in size will be limited.

Surface flatness in the target foils is expected to be at a maximum in solid crystal scintillating target foils, decrease for powder-deposited scintillating foils and vary for thin-metal OTR foils. If the surface is not flat then the light emission distribution will depend on the surface properties. When the surface is uniformly rough then the emission distribution could be modelled by convolving the initial distribution with a function which is isometric about the mean surface normal. This would tend to widen the point spread function, leading to a larger systematic bias and lower precision in beam measurements. The emission distribution of a foil containing large-scale ripples (such as the thin metal OTR foils) may be much more difficult to model. The ripples will cause localized warping in the beam image which will be unaccounted for by the geometry correction algorithm. In either case, the surface quality will add uncertainty to the collection efficiency corrections. The affect of target surface conditions on the precision of profile measurements was not quantified.

It was demonstrated in Section 4.2.1 that beam profile measurements could be made more accurate by removing systematic bias caused by the point spread function; therefore, the precision of the beam profile measurements are limited by the precision of the PSF properties. The relative uncertainty in the measurement of the PSF properties caused by uncertainty in the camera focusing procedure was quantified in Section 4.2.3.

The following results only include uncertainty introduced during the focusing procedure; therefore, they should be interpreted as the *maximum* precision of each measurement.

The maximum precision of the centroid measurement when imaged by the ELBT and EHBT optical configurations collecting scintillation light is shown in Figure 4.12. It is clear that the maximum precision ($0 \mu\text{m}$ to $2 \mu\text{m}$) is much better than the worst-case request of $25 \mu\text{m}$ but the requirement cannot be claimed to be satisfied without quantifying the other sources of uncertainty (which may be significant).

The maximum precision of the beam width measurement when imaged by the ELBT and EHBT optical configurations (again, collecting scintillation light) is shown in Figure 4.13. The maximum precision ranges from $0 \mu\text{m}$ to $5 \mu\text{m}$ for all conditions and is well above the minimum precision of $20 \mu\text{m}$ required in the higher energy section

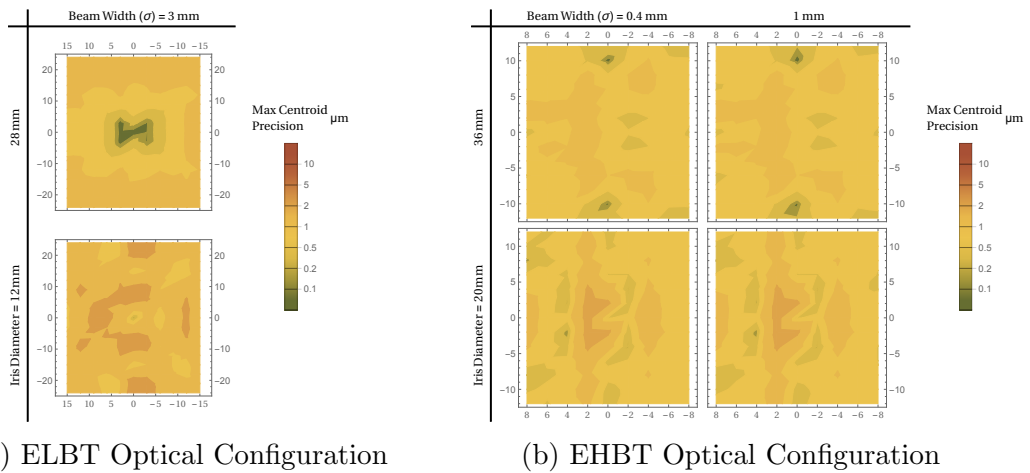
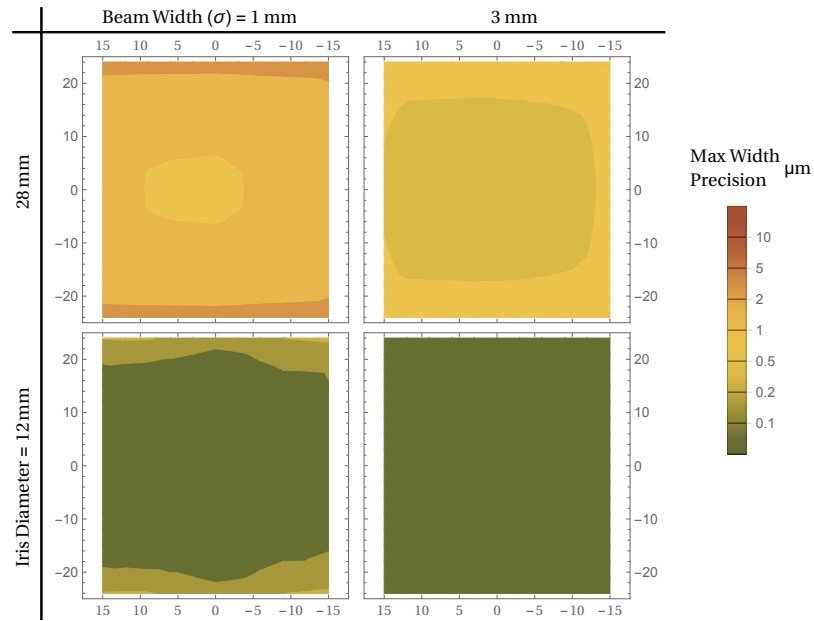


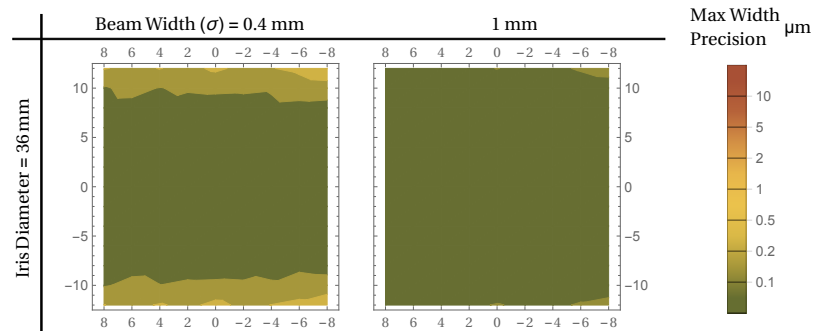
Figure 4.12: Maximum Precision of the Centroid Measurement for the ELBT and EHB Optical Configurations. These results include the uncertainty in the camera focusing procedure but not the uncertainty in light emission due to target surface quality. The beam centroid measurement contains two components: M_{10} and M_{01} ; for that reason, the larger of $\sigma_{M_{10}}$ and $\sigma_{M_{01}}$ is shown.

of the accelerator. Again, this requirement cannot be satisfied without quantifying all sources of uncertainty in the system.

To summarize, systematic bias in beam profile measurements was observed by the simulation of Gaussian shaped beams striking scintillating targets. Two sources of bias were identified. The bias caused by the point spread function could be removed; however, uncertainty in the exact properties of the PSF reduced measurement precision. The corrected bias and maximum precision were compared against the functional diagnostic requirements. The required centroid accuracy is possible provided that the target actuator is properly aligned. It is not known whether the required precision of the beam centroid and width measurements is met without quantifying other sources of uncertainty.



(a) ELBT Optical Configuration



(b) EHBT Optical Configuration

Figure 4.13: Maximum Precision of the Centroid Measurement for the ELBT and EHBT Optical Configurations. These results include the uncertainty in the camera focusing procedure but not the uncertainty in light emission due to target surface quality. The beam centroid measurement contains two components: M_{10} and M_{01} ; for that reason, the larger of $\sigma_{M_{10}}$ and $\sigma_{M_{01}}$ is shown.

Chapter 5

Conclusions

This thesis described the process of specifying and implementing a control system for the e-linac view screen diagnostic system.

It was demonstrated that the diagnostic system satisfied most of the external requirements provided by the TRIUMF laboratory. Where complete satisfaction of the requirements could not be demonstrated, such as those relating to the relative uncertainty in beam measurements, a qualitative analysis showing the path forward was shown.

Internal system requirements were identified and implemented as a set of control system support services. A service enclosure was designed and subsequently commissioned as component of the diagnostic system for the e-linac prototype beamline. A data analysis algorithm was implemented within the Experimental Physics Industrial Control System. Validation of the efficiency correction stage of the image processing sub-algorithm was attempted by comparing real data with simulation results. The quality of beam profile measurements was evaluated by simulation studies and a method of accounting for two sources of systematic bias was designed.

The utility of the view screen system is continuously demonstrated by the frequency at which it's used to commission and tune the e-linac and observe beam profile features which would be difficult to see otherwise (such as the profile of the beam as it leaves the electron gun in Figure 5.1). The quality of beam profile measurements is appreciated by the expert operators who integrate view screen data into high-level applications.

It was mentioned frequently in Chapter 4 that the results of the efficiency correction and systematic bias studies were done assuming that the beam target was emitting scintillation light. A natural extension of this work is the study of the




Figure 5.1: Electron Beam Profile Cathode Grid Pattern.

complicated emission distribution of light produced via optical transition radiation.

Bibliography

- [1] Epics: Motor record. <http://www.aps.anl.gov/bcda/synApps/motor/>. Accessed: 2015-07-23.
- [2] M. Clift, R. Farnsworth, A. Starritt, and L. Corvetti. Motion control using epics and galil controllers. In *ICALEPCS*, page 1, 2009.
- [3] Gerhard M. Haak, Nelson L. Christensen, and Bruce E. Hammer. Experimental studies on the angular distribution of scintillation light from small {BGO} crystals. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 390(12):191 – 197, 1997.
- [4] Shane Koscielniak. E-linac beam diagnostics requirements. <http://elinac.triumf.ca/e-linac/config/operations-modes/requirements/e-linac-diagnostic-requirements-2010-sep-19>, Sept 2010.
- [5] Shane Koscielniak. <http://elinac.triumf.ca/e-linac/wbs-areas/beamlines/coordinate-systems>, July 2013. Last Modified: July 18, 2013 01:42 PM.
- [6] Alan Murta. Gpc general polygon clipper. <http://www.cs.man.ac.uk/~toby/gpc/>. Accessed: 2015-12-14.
- [7] Mark L. Rivers. areadetector: Software for 2-d detectors in epics. *AIP Conference Proceedings*, 1234(1):51 – 54, 2010.
- [8] Douglas W. Storey, University of Victoria (B.C.). Dept. of Physics, Astronomy, and University of Victoria (B.C.). A view screen beam profile monitor for the ariel e-linac at triumph. Master’s thesis, University of Victoria, 2011.

Appendix A

Supplementary Information

A.1 Optical Simulation

An optical simulation was developed to produce sample image data, create tables of collection efficiency data, evaluate optical component configurations and measure properties of the point spread function. This section describes the implementation of the simulation.

The program simulates, in parallel, the propagation of rays produced by a given beam distribution, light distribution and optical configuration. Snell's Law is used to calculate the direction of travel at the interface of two different material. The simulation was initially implemented with the commercial *LensLab* optical simulation framework which plugs into the *Mathematica* symbolic computing framework. When limitations of LensLab could not be overcome the simulation was re-implemented purely in Mathematica. When a need to simulate a large number of beam images arose the performance of Mathematica became insufficient. A 10000 fold reduction in simulation execution time was realized by switching the programming language from Mathematica to C++, writing the ray tracing procedure in OpenCL and executing portions on a dedicated graphics card. Validation was performed regularly by comparing results with those of previous iterations and other work [8].

The simulation accepted as input:

- a file specifying the layout and parameterization of optical components (low-energy and high-energy camera box configurations are included by default),
- the diameter of the mechanical iris diameter and

- a description of the desired beam type.

and produced a raw image of the beam as output. Three beam types were supported: isometric Gaussian, uniform rectangular and a “point spread function” psuedo-beam. The only light emission distribution available was that of scintillation light.

The light emitted by a scintillating surface follows a Lambertian distribution [3]. The intensity is peaked along the direction of the surface normal (\hat{n}) and decreases as the azimuthal angle θ increases (Fig. 123). The probability of a photon being emitted in a random direction D within a solid angle $d\Omega$ centred on the polar angle θ and the azimuthal angle ϕ is proportional to $\cos\theta$:

$$\Lambda(D \in d\Omega) \propto \cos\theta d\Omega = \frac{\cos\theta d\Omega}{\int_0^{\pi/2} d\theta \int_0^{2\pi} d\phi \cos\theta \sin\theta} = \frac{1}{\pi} \cos\theta \sin\theta d\theta d\phi.$$

The angle Φ is distributed uniformly over the interval $[0, 2\pi]$ because Λ is independent of ϕ ; the marginal distribution in Θ is

$$\Lambda_{\Theta}(\theta) = \int_0^{2\pi} d\phi \frac{1}{\pi} \cos\theta \sin\theta d\theta = 2 \cos\theta \sin\theta d\theta.$$

Λ_{Θ} is sampled through an inverse transformation:

$$\begin{aligned} \int_0^{\theta} \Lambda_{\Theta}(\theta') d\theta' &= \int_0^{\theta} d\theta' \frac{1}{\pi} \cos\theta' \sin\theta' = \sin^2\theta = x, \\ \implies \theta &= \arcsin \sqrt{x}. \end{aligned} \tag{A.1}$$

where X is distributed uniformly over $[0, 1]$.

Rather than sample Λ at random, an $(N \times M)$ -length set of (θ, ϕ) pairs was generated by subdividing $[0, 1]$ and $[0, 2\pi]$ into N and M length intervals, respectively. This ensures a uniform sampling of Λ while minimizing the number of samples.

A large portion of the full solid angle is directed away from the view screen optics. To further improve simulation efficiency, photons were only generated within the particular solid angle Ω which was directed toward the optics. The limits on Ω vary as a function of the iris diameter and optical configuration; they were determined empirically. To compare simulations where Ω differed, the results were normalized.

If C is an event in which a photon is collected by the optics and R_{Ω} is an event in which a photon emitted from the screen enters the solid angle Ω which is directed

toward the optics, then

$$\begin{aligned}
 P(C) &= P(C \cap \Omega) + P(C \cap \bar{\Omega}) \\
 &= P(C | R_\Omega) P(R_\Omega) + \cancel{P(C \cap \bar{\Omega})} \\
 &= P(C | R_\Omega) \int_{\Omega} \Lambda(\theta, \phi). \tag{A.2}
 \end{aligned}$$

The term $P(C \cap \bar{\Omega})$ is equal to zero because it is impossible to collect a photon travelling away from the view screen optics.

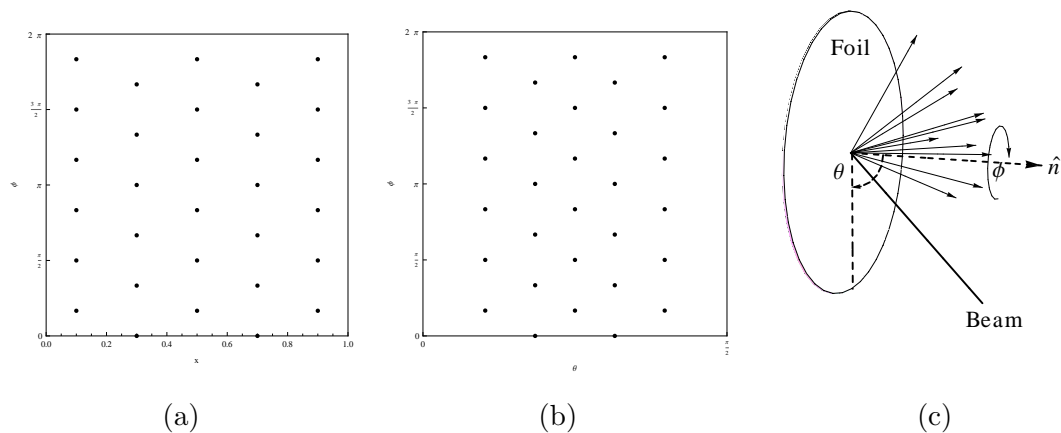


Figure A.1: An $N \times M$ grid of equally spaced points in probability space (a) are used to generate a set of photons following the Lambertian distribution (b).

There were numerous limitations of the optical simulation are now discussed. Dispersion is ignored because each component is described with a constant index of refraction. It is assumed that the wavelength of light is that which the camera CCD is most sensitive to. Partial reflection was enabled in the Mathematica version of the simulation but was dropped in the C++ version to simplify the ray tracing code and improve performance. To study cases where internal reflection produced “ghost” images on the CCD sensor the Mathematica implementation was used.

A.2 External Software Dependencies

External software dependencies used by the software portion of the view screen control system shown in Table A.1 for EPICS software modules and Table A.2 for non-EPICS software.

Name	Brief Description
areaDetector	Contains camera drivers with EPICS integration and a general purpose image processing framework.
asyn	An EPICS dependency for areaDetector. Provides a C++ callback framework for responding to EPICS record state changes.
autosave	Writes a user-defined subset of the EPICS database state to disk for recovery at a later time.
devIocStats	Tracks various characteristics of the host operating system and exports the information to the EPICS database.
EDM	EPICS Display Manager. The main operator interface implementation currently used at TRIUMF.
EPICS	The Experimental Physics Industrial Control System. It provides a mechanism for diagnostic devices to communicate with software through “records”.
galil	EPICS integrated motor controller driver. Exports the digital and analog I/O into EPICS records.
motor	Provides the “motor” EPICS database record type. The motor record contains fields which describe the current motor state. Used by the Galil EPICS driver to manage the stepper motor which actuates the mechanical iris.

Table A.1: EPICS Software Dependencies

A.3 Data

This portion of the appendix contains the measurements (Table A.3) used to calibrate the temperature probe used within the first generation service enclosure cabinet.

Name	Brief Description
GraphicsMagick	A general-purpose image manipulation framework. Used by the areaDetector EPICS module to implement some image storage routines.
isc-dhcp-server	Internet Systems Consortium (ISC) implementation of the Dynamic Host Configuration Protocol (DCHP). Assigns IP addresses to the Ethernet cameras.
X11	The X Windowing System, version 11.
Libxml2	A C-language XML parser and writer that is used to manage the view screen configuration files.

Table A.2: Non-EPICS Software Dependencies

AI (± 0.1 V)	Temp. (± 1.7 °C)
4.4	25.8
5.0	27.2
5.2	27.4
5.4	28.8
5.5	29.0
5.8	30.0
6.1	32.4
7.2	34.0
9.99	42.2

Table A.3: Cabinet Temperature Transmittor Calibration Data

Acronyms

ARIEL Advanced Rare IsotopE Laboratory. 3, 6, 7

AWG American Wire Gauge. 22

CCD Charge-Coupled Device. ix, 5, 23, 24, 26, 33, 37, 39, 41, 44, 59, 67, 69, 70, 78

EHBT Electron High energy Beam Transport. vi, x, xi, 4, 57, 59, 60, 62–73

ELBT Electron Low energy Beam Transport. vi, viii–xi, 4, 28, 30, 37, 40, 44, 53, 57–73

EPICS Experimental Physics Industrial Control System. vi, 7, 9, 13–15, 18, 20, 42, 45, 56, 65, 78–80

ISAC Isotope Separator and Accelerator. 3

OPI OPerator Interface. 7

OTR Optical Transition Radiation. 5, 15, 42, 64, 71

PSF Point Spread Function. x, xi, 42, 58–68, 71, 72

RIB Rare Isotope Beam. 3, 4