

On the Creation and Measurement of Rare Isotopes

or,

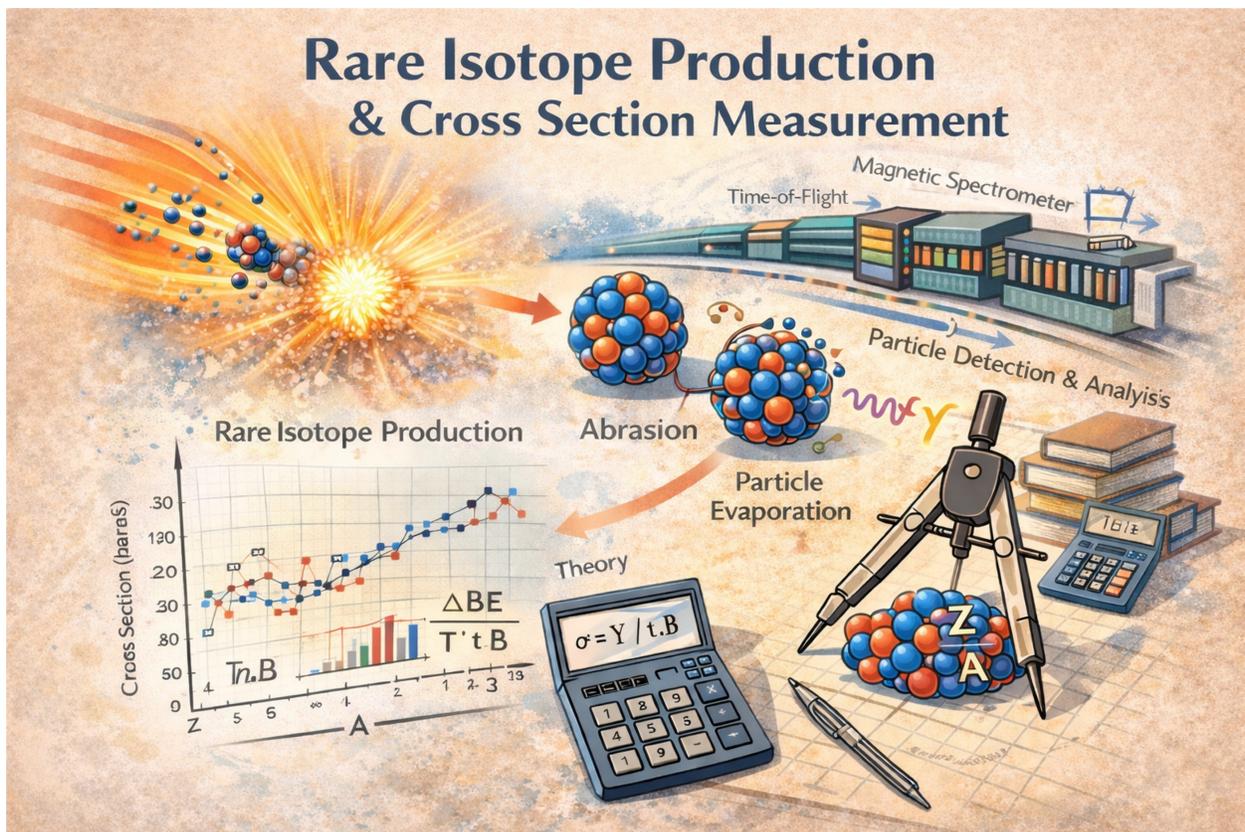
How Rare Is Rare?

A Practical Guide to Cross Sections at the Limits of Stability

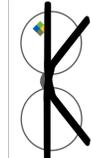
Shane Watters*,

on behalf of the Rare Isotope Research Group

Version 1.0.4 from February 11th, 2026



ChatGPT's Graphical Summary of this Manual



*watterss@frib.msu.edu

Contents

1	Introduction	4
1.1	Data Analysis Workflow	5
2	Experimental Data Analysis	6
2.1	Overview of the Workflow	6
2.2	Data Calibration and Cleaning	7
2.2.1	Magnetic Rigidity Calibration	8
2.2.2	Time-of-Flight Calibration	10
2.2.3	Data Cleaning	11
2.2.4	Total Kinetic Energy Calibration	15
2.2.5	Z Calibration	18
2.2.6	Final PID Resolution	19
3	Data Extraction and Isotopic Yields	20
3.1	Exporting Data from <code>SpecTcl</code>	20
3.2	Obtaining Yield Statistics	20
3.3	PID Deduced Integer Values	22
3.4	Counting Statistics	24
4	Experimental Inputs for Cross Section Determination	25
4.1	Target Thickness	25
4.2	Beam Integral	25
5	LISE⁺⁺ Transmission Analysis	27
5.1	Procedure	28
5.2	Distinguishing Primary and Secondary Fragments	28
5.3	Estimating Transmission Uncertainty	29
6	Cross Section Calculation	32
6.1	From First Principles to Measured Yields	32
6.2	Uncertainty Redux	34
7	Interpretive Analysis	35
7.1	Partial vs Total Cross Sections	36
7.2	Abrasion-Ablation	37
7.2.1	Dissipation and Friction	38
7.2.2	Abrasion-Fission	40
7.2.3	Abrasion–Ablation Calculations in <code>LISE⁺⁺</code>	41
7.2.4	Performing Abrasion-Ablation Fits	43
7.3	Δ BE Systematics	50
7.3.1	Performing Δ BE Fits	51
7.3.2	Mass Table Adjustments	52
7.4	Complementarity of Abrasion–Ablation and Δ BE Analyses	53

8 Concluding Remarks	54
Appendix A Installing and Booting SpecTcl and SpecTk	55
Appendix B Analysis Program Basics	58
B.1 SpecTcl	58
B.1.1 SpecTcl Filtering Output	61
B.2 SpecTk	62
Appendix C Ion Optics and Trajectory Reconstruction Primer	63
C.1 The Transfer Matrix and Other Optical Definitions	64
C.2 The Optics of ARIS	68
C.2.1 Momentum Compression	68
C.2.2 Optical Modes of Operation	68
C.2.3 Secondary Beam Purification	70
C.3 Trajectory Reconstruction	72
C.3.1 Detector Calibration and Connection to Transmission	74
C.3.2 Empirical Transfer Matrix Determination	76
Appendix D Quantitative Gate Analysis	77
D.1 Conceptual Framework	78
D.2 Procedure	78
D.3 Diagnosing Problematic Gates	80
D.4 Best Practices	80
Appendix E Particle Identification	80
E.1 Overview of Measured and Derived Quantities	81
E.2 Mass-to-Charge Ratio and Mass Reconstruction	81
E.3 Atomic Number Reconstruction	83
E.4 Assumptions and Validity	84
E.5 Common PID Failure Modes	84
E.6 Identification of Blobs	85
E.6.1 Primary Identification Methods	85
E.6.2 Secondary (Corroborative) Identification Methods	89
E.6.3 Failure Modes and Best Practices	90
E.7 Summary and Recommended Workflow	90
Appendix F Momentum Integrated Total Cross Section	91
FAQ	93
References	95
Suggested Further Reading	97
Acknowledgments	99

1 Introduction

The purpose of this manual is to guide researchers through the complete analysis of rare isotope production cross sections. The workflow described here is intended to be sufficiently detailed that a new investigator, with a background in nuclear physics, can carry out a full analysis independently, while also serving as a practical reference for experienced users. The procedures described here are tailored to in-flight rare isotope production experiments using magnetic separators and event-by-event particle identification. In particular, this manual has been created for use with the Advanced Rare Isotope Separator (ARIS) [1, 2], and serves as a concrete realization of the analysis concepts discussed throughout this manual. Examples are drawn from its implementation wherever specificity is useful. A schematic of the system is shown in Figure 1. A discussion of ARIS optics is presented in Appendix C.

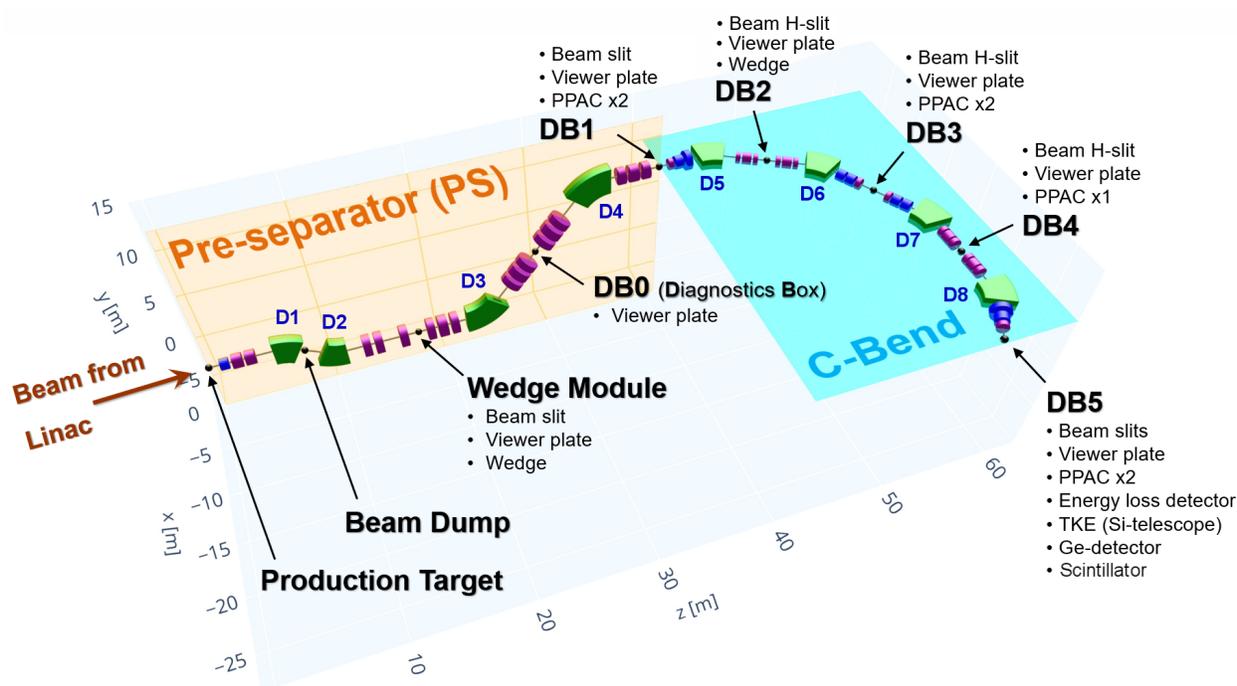


Figure 1: A schematic overview of the Advanced Rare Isotope Separator. Detector and material locations are shown.

The analysis of rare isotope production cross sections is inherently multi-stage and interdisciplinary, combining detector calibration, data cleaning, ion optics, nuclear reaction modeling, and statistical uncertainty analysis. The manual assumes familiarity with basic concepts such as energy-loss, magnetic rigidity, and time-of-flight techniques, but does not assume prior experience with this specific analysis framework. For the first-time analyst, reading the entire manual before starting is recommended to understand the flow and scope of the analysis.

Because of this complexity, errors introduced at any stage can propagate forward and significantly affect the final results. For this reason, the manual emphasizes not only the procedural steps required to complete the analysis, but also the physical reasoning that

motivates each step. The ordering of the analysis steps is intentional; later stages rely critically on the stability and physical validity of earlier calibrations and cleaning.

The manual is organized into five overarching components. These components mirror the natural progression of the analysis from raw data to physics interpretation:

1. Experimental data calibration and cleaning,
2. Extraction of experimental parameters,
3. Transmission analysis,
4. Cross section calculation and error propagation, and
5. Interpretive analysis of extracted cross sections.

Wherever possible, the discussion is framed to be self-contained. Implementation details that are not essential to understanding the analysis logic are deferred to the appendices, while references to external documentation and published resources are provided for readers who wish to explore underlying methods in greater depth.

1.1 Data Analysis Workflow

At a high level, the rare isotope production cross section analysis proceeds through the following sequence:

1. Raw experimental data are calibrated to convert detector signals into physically meaningful quantities, including magnetic rigidity, time-of-flight, energy-loss, and total kinetic energy.
2. A series of physics-motivated gates is applied to remove nonphysical events and background contributions while preserving true reaction products.
3. Cleaned and calibrated particle identification spectra are used to extract elemental and isotopic yields.
4. Transmission corrections are applied to account for ion-optical acceptance, reactions in materials, charge-state effects, and higher-order transport effects that are not directly accessible in the experimental data.
5. Production cross sections are calculated from the corrected yields, and uncertainties are propagated to obtain final error estimates.

The remainder of this manual follows this sequence explicitly. Each section is structured to explain both how the analysis step is performed and why it is necessary, with particular attention paid to identifying common failure modes and sources of systematic uncertainty. Appendix A and Appendix B provide software startup instructions and user-interface references which are used throughout the text; these appendices are recommended reading for

first-time analysts. However, this manual does not attempt to document the full functionality of `SpecTcl`¹, `SpecTk`², or `LISE++`³; instead, it focuses on a validated analysis path that produces reproducible and physically meaningful cross sections.

2 Experimental Data Analysis

The goal of the experimental data analysis is to convert raw detector signals into clean, calibrated fragment yields suitable for cross section extraction. This stage establishes the foundation of the entire analysis. Errors introduced here propagate directly into transmission estimates and, ultimately, into the reported cross sections.

The analysis is performed using `SpecTcl` and `SpecTk`, which together provide tools for histogramming, gating, calibration, and yield extraction. While these programs are highly flexible, the procedures outlined below define a standard, reproducible workflow that should be followed unless a compelling experimental reason dictates otherwise.

2.1 Overview of the Workflow

At a high level, the experimental data analysis proceeds through the following steps:

1. Launch `SpecTcl` and `SpecTk` and load the appropriate experiment configuration files.
2. Verify detector readout integrity and establish basic diagnostic spectra.
3. Perform detector calibrations (timing, position, energy, and charge).
4. Apply physics-motivated data-cleaning cuts to remove unphysical or corrupted events.
5. Construct particle identification (PID) spectra.
6. Extract fragment yields and associated statistical uncertainties.

Before beginning the analysis, ensure that both `SpecTcl` and `SpecTk` are installed and configured correctly. Start up instructions are provided in Appendix A, and user interface basics are presented in Appendix B.

Start `SpecTcl` from the command line. Once `SpecTcl` is running, launch `SpecTk` and connect it to the active `SpecTcl` session.

Upon successful connection, verify that:

- Spectra can be created and updated,
- Event rates appear reasonable,

¹FRIBDAQ and `SpecTcl` documentation are available at <https://docs.frib.msu.edu/daq/newsite/index.php>. All `SpecTcl`-related questions are directed here.

²`SpecTk` can be obtained from its development GitHub repository or from https://lise.frib.msu.edu/spectk/spectk_home.html, which also hosts the official documentation.

³<https://lise.frib.msu.edu/lise.html>

- No error messages related to missing parameters or channels are present.

If unexpected errors occur at this stage, they should be resolved before proceeding. Many downstream analysis issues can be traced back to configuration or mapping problems introduced here.

Other modes for analysis failures originate from subtle issues introduced early in the workflow, such as incomplete data cleaning, miscalibrated timing offsets, or residual kinematic correlations. These issues often manifest only at later stages as poor PID separation, unstable Gaussian fits, or anomalous transmission behavior. For this reason, analysts are strongly encouraged to validate each step before proceeding.

2.2 Data Calibration and Cleaning

Once `SpecTcl` and `SpecTk` are running and experimental data have been loaded, the next step is data calibration. This includes, at minimum, the calibration of magnetic rigidity ($B\rho$), time-of-flight (ToF), and total kinetic energy (TKE). It is important to note that calibrating these measured observables is directly linked to calibrating the particle identification parameters A/q and Z . The relationships between these variables are presented in Appendix E.

Calibration steps should be performed in a consistent order, as later calibrations often depend implicitly on earlier ones. In particular, magnetic rigidity and time-of-flight calibration should be stable before attempting total kinetic energy calibration. These calibrations directly determine the quality of the particle identification and therefore the reliability of extracted fragment yields.

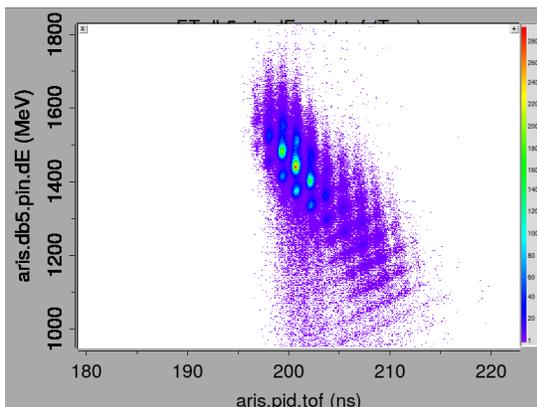


Figure 2: Example of an uncleaned and uncalibrated particle identification plot.

The process of calibration converts raw electronic signals into physically meaningful quantities such as time-of-flight, position, and energy-loss. Calibration procedures vary by detector type but generally follow the same principle: known reference points are used to define a transformation from raw channels to physical units. An example of converted (but uncalibrated) detector output (time-of-flight in nanoseconds and energy-loss in MeV) is shown in Figure 2. Note that the blobs are ambiguous in their identity. The goal of PID is to increase the resolution and to determine who they are.

All calibration files are kept in the `calibration` folder, locally inside the `SpecTcl` build. The exception is the `brho_raw.csv` file, which is shared across builds. This file should not be modified unless the user is confident in the validity of the changes. This file sets the baseline $B\rho$ used in all calculations. Exploratory modifications should be performed on a local copy, with the file path updated in `/calibrPID/brho.tcl`.

A valuable diagnostic tool during detector calibration is the ARIS scaler report. During

data acquisition, scaler readouts record the integrated rates of selected detector and logic channels over the duration of a run. Each scaler counts the number of times a specified logic signal – typically the output of a discriminator, trigger, or coincidence condition – fires within the scaler gate. Because scalers are generally free-running and insensitive to DAQ dead time, they provide a robust means of monitoring detector performance, trigger stability, and overall run conditions. It is important to note that the scalers represent time-integrated rates rather than event-by-event measurements and do not contain energy, timing, or correlation information. Accordingly, scaler values should be routinely monitored both online and during post-experiment analysis. These reports may be accessed from `/user/arisdag/vme/ARISscaler`.

2.2.1 Magnetic Rigidity Calibration

The first experimental observable to be calibrated and enhanced is the magnetic rigidity $B\rho$. The rigidity is one component of the derived mass to charge ratio A/q , which is the ultimate calibration we are concerned with. It is this derived quantity which, in conjunction with atomic number Z and charge-state q , allows us to uniquely identify every fragment.

After ensuring the raw magnetic rigidities have been set accordingly (these values correspond to the measured fields of the dipole magnets), trajectory reconstruction corrections to $B\rho$ may be implemented. For a detailed discussion on ion optics and trajectory reconstruction, see Appendix C. The trajectory reconstruction is obtained via a LISE⁺⁺ transfer matrix calculation in the extended configuration.

2.2.1.1 Trajectory Reconstruction Transfer Matrices

To correct first-order optical correlations in $B\rho$, a trajectory reconstruction based on calculated transfer matrices is required. This is done through the LISE⁺⁺ extended configuration of ARIS.

This file is distributed in the LISE⁺⁺ package (`/examples/FRIB/eARIS/eL_ARIS_k3cb2`). Configure the extended file to match the experimental setup⁴.

⁴The information can be found in an experimental planning LISE⁺⁺ file. Parameters to set include primary beam, energy, fragment of interest, target thickness and composition, wedge thickness(es) and composition,

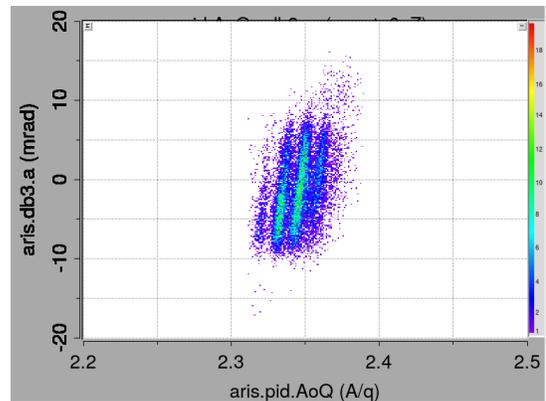


Figure 3: A/q vs DB3 x-angle before kinematic corrections are applied.

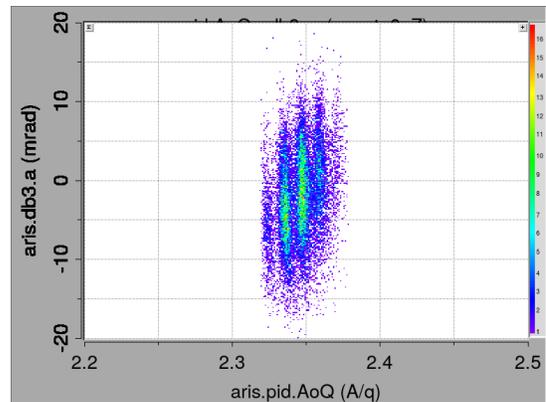


Figure 4: A/q vs DB3 x-angle after kinematic corrections are applied.

Once the general experimental parameters are set, load an experimental saveset into the extended configuration. Savesets contain snapshots of the magnet settings at a given time (e.g., during tuning or for a specific run). Savesets are currently saved in `/departments/ARIS /info/FTC_ARIS`.

The extended configuration is now ready to calculate the transfer matrices between Diagnostic Boxes (DBs) 3-4 and DB4-5. In the main LISE⁺⁺ window, click **Experimental Settings** → **Optics** → **Calculate matrix between two selected blocks**. In the dialog box that appears choose the starting and ending points for the transfer matrix calculation. For DB3-4, use DB3 PPAC 1 to DB4 PPAC0; for DB4-5, use DB4 PPAC0 to DB5 PPAC 1⁵. Convert these matrix elements into a form that SpecTcl can read; previous examples are located inside the SpecTcl calibration folder: `/calibrPID/matrix`.

2.2.1.2 Implementing Kinematic Corrections

To enable the kinematic corrections to $B\rho$, change the `Brho_method` in the `/calibrPID/CSV/pid_values` spreadsheet. Multiple $B\rho$ reconstruction methods have been implemented in SpecTcl to mitigate sensitivity to detector resolution and optical aberrations at different focal planes. Method 34 utilizes the matrix between DB3-4, while method 45 utilizes the matrix between DB4-5. Start with method 34. Reload the experimental data in SpecTcl and gate on an individual Z line in the PID plot (Z vs A/q)⁶. Apply this gate to the position and angle vs A/q plots and look for any correlation in the plots at the focal planes (DB3, 4, and 5). Proper kinematic correction removes first-order correlations, producing vertical distributions with minimal residual aberrations. If the LISE⁺⁺ matrix elements do not perfectly negate the first-order correlations, tweak the values until the correlations disappear. An example of before and after kinematic corrections are applied is shown in Figures 3 and 4, respectively.

Sometimes, there exist second-order aberrations in the data. These usually manifest as tails on the ends of the data, forming an ‘S’-shape. Second-order corrections can be applied to the data. In order to enable them, change the value of ‘fine’ column in the `pid_values` spreadsheet from 0 to 1. Next, make adjustments in the spreadsheet `pid_fine_values` to correct for the aberrations. It is important that these corrections are generally small – on the order of 10^{-4} or smaller.

As the kinematic corrections are applied and refined, the A/q resolution should increase (that is, the sigma of the blobs should decrease). Once `brho_method` 34 is refined, move to method 45 and repeat the process. After refining both methods, take their average with method 345. This average will further enhance the resolution of A/q (if both 34 and 45 are well-behaved). Figure 7 shows the increased resolution between method 1 and method 345.

If either method 34 or 45 is unsuitable for the analysis at hand, use only the well-behaved method. In instances where neither is appropriate (e.g., when the DB4 PPAC is not used), default to method 1, which sets all particle rigidities to the measured value at DB5.

and detectors.

⁵See https://lise.frib.msu.edu/work/SpecTcl/ARISmatrices_withLISE.pdf for an illustrated guide.

⁶The element with the brightest blobs is usually recommended, regardless of whether or not it contains the fragment of interest.

2.2.2 Time-of-Flight Calibration

Time-of-flight (ToF) is the second measured component of the mass-to-charge ratio A/q . It is calibrated between DB3 and DB5 using detector pairs selected for optimal timing resolution. The locations of detectors along the beam path are stored in the `/calibrPID/locations.tcl` file. SpecTcl automatically calculates the path length between pairs of given detectors; this quantity is held fixed for the purposes of the time-of-flight analysis.

Where possible, multiple independent time-of-flight measurements should be averaged to improve timing resolution. This is because each individual timing signal provides an estimate of the same physical quantity – the particle’s flight time – but is subject to uncorrelated statistical fluctuations arising from detector response, electronics jitter, and signal processing. When several such measurements are combined, these random fluctuations partially cancel. As a result, the uncertainty on the averaged ToF decreases approximately as $1/\sqrt{N}$, where N is the number of **independent** timing measurements. The increased resolution of the particle identification when an averaged ToF method is used is shown in Figure 6. Note that the PID shifts due to calibration differences between the time-of-flights. Physically, the increased resolution reflects the fact that while any single detector may register an early or late time by chance, it is unlikely that all detectors fluctuate in the same direction by the same amount. Averaging therefore suppresses random timing noise while preserving the true flight time, leading to narrower ToF distributions and improved downstream quantities such as A/q , charge-state separation, and PID resolution.

After new detectors have been installed, it may be of interest to perform a principal component analysis. This analysis estimates the intrinsic timing resolution of the individual detectors after removing the global jitter from the signals. This analysis allows one to see block jitter – the

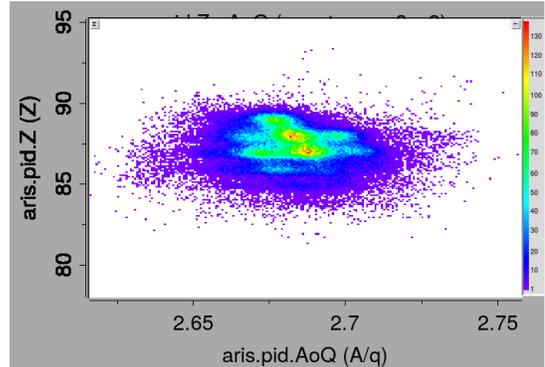


Figure 5: Before kinematic corrections (Brho_method 1) with single time-of-flight method. $\sigma(A/q) = 0.0038$.

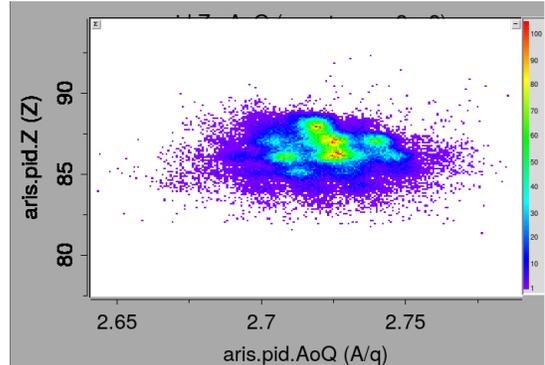


Figure 6: Averaged time-of-flight method used but no kinematic corrections. $\sigma(A/q) = 0.0031$.

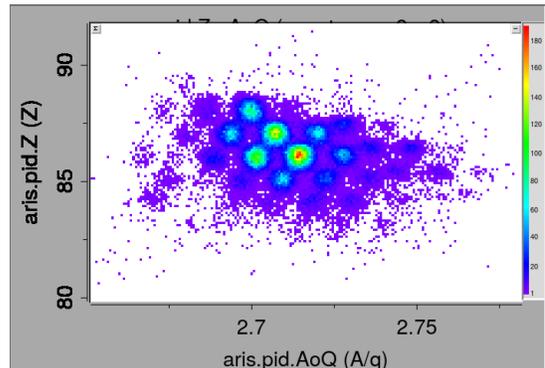


Figure 7: Averaged time-of-flight method and kinematic corrections used. $\sigma(A/q) = 0.0021$.

way all detectors on the same module fluctuate relative to others – as well as reveal any other statistical correlation between the detectors. This has the power to reveal improvements that can be made through electronics adjustments. Additionally, this analysis will inform the best signals to combine to reduce the overall ToF variance.

Alternatively, a brute-force approach may be used in which the standard deviation of every unique combination of detectors is calculated to find the best overall resolution⁷.

The timing method to use is specified in the `pid_values` spreadsheet. The various methods of the ToF are defined in the file `/source/Caris_pid/Caris_pid_LenToF.cpp`. New averaged methods may be added with a new index number. Note that `SpectCl` must be remade after changes are made to a `.cpp` file.

Timing offsets are used to move the measured ToF to agree with physically-reasonable values. These values are determined by a LISE⁺⁺ Monte Carlo simulation. Using the pre-configured experimental LISE⁺⁺ file, calculate the time-of-flight between DB3 and DB5⁸. Use the mean ToF value as the baseline calibration. Final adjustments will be made as dictated by the particle identification.

To calibrate individual pairs of detectors, use the files `/calibrPID/CSV/tof_offsets`. Once this has been achieved (accuracy need not be perfect), the averaged time-of-flight method may be set in the `pid_values` spreadsheet. Note that the global ToF offset is also located in this sheet and is how the averaged method ToF will be adjusted, if necessary⁹.

The final ToF calibration is to ensure that the PID lines up where expected. During beam tuning online, particle identification (PID) should already have been determined¹⁰. Adjust the global ToF offset to align the PID pattern (A/q vs Z). If PID needs to be performed from scratch, see Appendix E.

2.2.3 Data Cleaning

With PID finalized, the next step is data cleaning. Cleaning is performed by applying physics-motivated gates to remove nonphysical or background events (colloquially referred to as “bad counts”) while preserving true reaction products. There are many ways to remove these bad counts; the ones presented here are the most general. Representative examples of common cleaning gates are shown in Figures 8–10.

It must be noted that these gates constitute a first-pass cleaning. They may be refined later as necessary as the analysis progresses. It is important to emphasize that over-cleaning can bias yields by preferentially removing valid low-statistics events. Cleaning gates should therefore be as tight as necessary, but no tighter, and should always be validated by inspecting their effect on PID and yield stability. When in doubt, apply gates conservatively and tighten or loosen them incrementally as warranted by the data. Quantitative gate analysis is discussed in Appendix D.

⁷It should be noted that this becomes infeasible as the number of detectors at both DB3 and DB5 become large. An Excel spreadsheet has been made to calculate standard deviations of up to 5 timing signals.

⁸For quick calibrations, a typical flight time between DB3 and DB5 is approximately 130 ns.

⁹It is good practice to set this value to 0 at the beginning of a new analysis, as this global offset is applied in addition to any local offsets (e.g., detector pair offsets).

¹⁰This information is located in the Olog: <https://controls.frib.msu.edu/logbook/>

For high- Z beams, an additional charge-state gate should also be employed after the total kinetic energy calibration; see subsection 2.2.4.

PPAC Sum Gates

Parallel Plate Avalanche Counters (PPACs) play a critical role in event validation due to their detector design and fast signal response. These detectors consist of thin electrodes separated by a narrow gas gap and operated under a strong electric field. When a charged particle traverses the detector, it produces an electron avalanche that results in a prompt electrical signal [3]. The minimal material budget ensures negligible energy-loss and preserves the fragment trajectory¹¹.

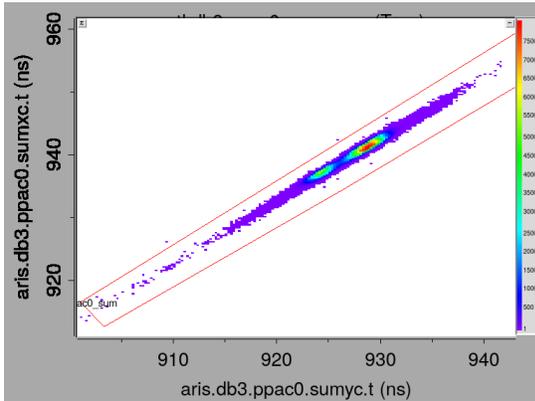


Figure 8: Example of a PPAC sum gate. Valid events lie along the main diagonal.

Events that do not correspond to physically meaningful trajectories are readily identified through deviations in the timing sum. Examples include particles that clip the detector edges, traverse the detector at large angles, scatter upstream, or originate from secondary reactions. Such events often produce delayed, asymmetric, or incomplete signals, leading to broadened or shifted timing-sum distributions. By applying gates on the timing sum, these unphysical events can be efficiently rejected. An example of this is shown in the sum spectra of Figure 8.

When a pair of PPACs is installed at the same focal plane, the particle trajectory angle can be reconstructed from the two position measurements. A particularly powerful diagnostic is the tracking spectrum formed by connecting these two positions, which directly visualizes the reconstructed particle paths between the PPACs. Large-angle particles are readily identifiable in this representation. During beam tuning, this capability provides a fast and intuitive check of the angular distribution of the beam at the focal plane.

When multiple PPACs are installed along the beamline, additional consistency checks become possible. Valid fragments exhibit correlated timing sums and positions across successive PPACs, consistent with the expected transport through the magnetic optics. Events that fail to register in one PPAC, or that show inconsistent timing or position correlations

¹¹It is important to note, however, that for high- Z beams, charge-exchange may occur within the gas volume. This can be observed as a difference in the reconstructed momentum between segments 34 and 45, as seen in Figure 10

between detectors, are indicative of scattering, charge-changing reactions in degraders, or partial loss of acceptance.

PPAC-based conditions therefore act as geometric and kinematic filters that are largely independent of energy-loss observables. This independence is particularly valuable for high- Z beams and high-rate conditions, where pileup and reaction-induced background are more prevalent. Applying PPAC timing-sum and correlation gates early in the analysis chain suppresses artificial broadening in PID and isotopic spectra and improves the stability of downstream calibrations, including total kinetic energy calibration and yield extraction.

Energy-Loss Gates

Energy-loss gates are applied between successive detectors in the silicon stack to suppress light reaction products and other sources of background that do not correspond to the ions of interest. These gates exploit the fact that, for a charged particle traversing thin layers of material at fixed velocity, the specific energy-loss is primarily determined by its charge and velocity, and is only weakly sensitive to small variations in detector thickness or entrance angle. As a result, identical particles passing through adjacent thin silicon detectors are expected to deposit nearly the same amount of energy in each layer, producing a narrow, correlated locus in ΔE - ΔE space (Fig. 9).

In contrast, light reaction products, scattered beam particles, and partially stopped ions typically exhibit different charge-to-velocity combinations or undergo significant energy degradation between layers, leading to broader or displaced energy-loss correlations. By selecting events consistent with the expected ΔE - ΔE behavior of the isotope of interest, energy-loss gates provide an effective and largely model-independent means of rejecting these backgrounds while preserving high efficiency for fully transmitted heavy ions.

Dispersion Gates

Finally, a dispersion-ratio gate between Dip34 and Dip45 is applied when using the $B\rho$ method 345 to ensure optical consistency of the reconstructed trajectories. In a first-order achromatic beam line, the horizontal position at a dispersive focal plane is proportional to the particle's magnetic rigidity, such that ions with unchanged charge-state and momentum experience the same relative dispersion at successive dipoles. Consequently, the reconstructed dispersions at Dip34 and Dip45 are expected to scale linearly, producing a narrow locus along a 45° line when plotted against one another. This is shown in Figure 10, for many charge-states. The ions whose charge-state changed by the same amount form 45° lines in the spectrum. Both the charge-state and dq are labeled.

Events that deviate from this correlation typically correspond to particles that have

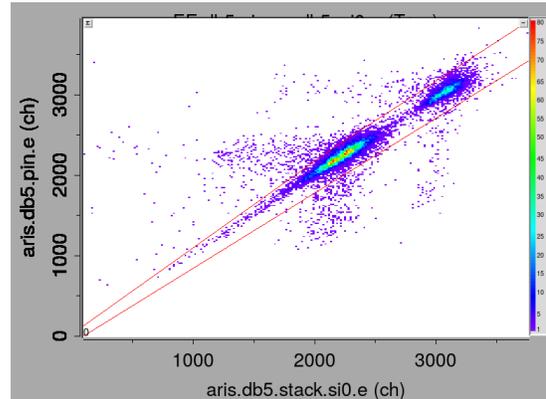


Figure 9: Example of an energy-loss gate. Valid events lie along the main diagonal.

undergone nuclear reactions, charge-changing processes, or large-angle scattering between the two dipoles, all of which alter the particle’s effective magnetic rigidity and disrupt the expected dispersion matching. By gating on the 45° correlation, the dispersion-ratio cut provides a robust means of rejecting such non-ideal trajectories while retaining ions that are cleanly transported through the spectrometer and are therefore suitable for reliable $B\rho$ reconstruction.

2.2.3.1 Sources of Background in Cleaned Spectra

Even after careful calibration and the application of well-designed cleaning gates, residual background is unavoidable in rare isotope data. Understanding the physical and instrumental origins of this background is essential for interpreting low-statistics yields and for avoiding overly aggressive cleaning.

One common source of background arises from the high- σ tails of neighboring PID peaks. Finite detector resolution in time-of-flight, energy-loss, and magnetic rigidity produces non-Gaussian wings that extend beyond the nominal centroid of each isotope. For fragments with large production cross sections, these tails can populate the PID space of rarer neighboring isotopes, particularly along directions of strong parameter correlation (e.g., A/q -ToF or $Z-dE$). This effect becomes more pronounced when the yield ratio between adjacent isotopes is large.

Background may also originate from unresolved charge-states. Partially stripped ions can form secondary bands that overlap or intersect the fully stripped PID loci. If not explicitly separated or accounted for, these charge-state contaminants may leak into otherwise well-defined isotope gates, especially at lower velocities or for high- Z fragments.

Secondary reactions in detectors or degraders constitute another irreducible background source. Nuclear interactions within the target, wedge, or detector material can produce fragments that mimic valid PID signatures but do not correspond to direct projectile fragmentation products. These events often populate diffuse regions of PID space and are difficult to remove without simultaneously sacrificing good events.

Finally, random coincidences and electronic effects – such as pile-up, multi-hit ambiguities, or imperfect signal reconstruction – can generate isolated or sparsely populated background events. While these contributions are typically small, they become significant when extracting yields for the most exotic, lowest-cross-section fragments, and are principally why new isotope discoveries require two counts for positive identification [5].

For these reasons, cleaning gates should be designed to suppress dominant background contributions while preserving physically plausible tails of genuine isotope distributions. Residual background is best treated as a systematic contribution to the yield uncertainty

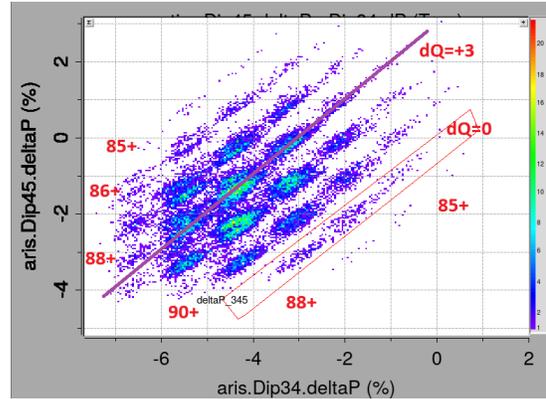


Figure 10: Example of a dispersion gate. The different horizontal and vertical lines represent charge-states. Diagonal lines represent constant charge change.

corrected by adjusting the detector bias voltage. Any change in detector bias invalidates previous energy-loss calibrations, and therefore must be redone after voltage adjustments.

Once proper detector biasing and stopping power is confirmed, TKE calibration may proceed. Conceptually, the TKE calibration enforces consistency between measured detector responses and the expected energy-loss and stopping behavior of identified ions. In practice, this means calibrating detector response functions so that reconstructed energies agree with well-established heavy-ion stopping calculations.

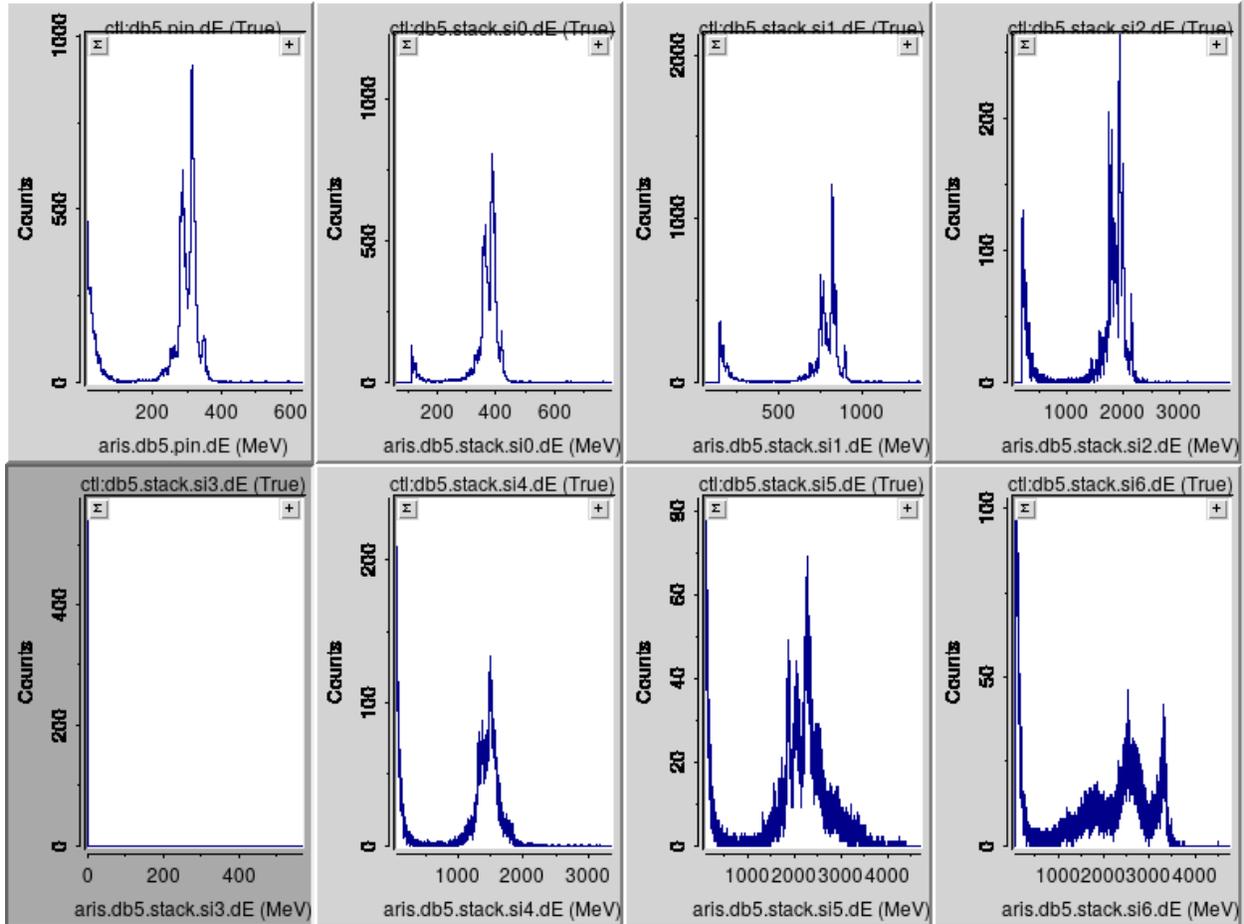


Figure 12: Example of run with invalid energy-loss signals: stack 3 is absent, and it is unclear if the particles stop in the last detector (stack 6)

On the PID spectrum, create 4-6 isotope identification gates spanning the full fragment distribution. These isotope gates, together with the previously defined cleaning gates, are applied to one-dimensional magnetic rigidity and energy-loss spectra to extract average detector channel values for each fragment. Ensure that all detectors registering a valid signal are included, up to and including where the fragment stops¹³.

The TKE calibration is performed using a dedicated Excel workbook that compares measured energy-loss and magnetic rigidity values with expectations calculated by LISE⁺⁺.

¹³Detectors downstream of the stopping location do not register physical energy-loss and should not be included.

The workbook determines detector slopes and offsets that best reproduce the simulated stopping behavior. For correct operation, the `lise64` Excel file must be the first Excel file to be opened; ensure that all macros and external content are enabled. Opening `lise64` while another Excel file is open may Excel may crash. The workbook can be found in `"/Program Files/LISEcute/lise64.xlsm"`.

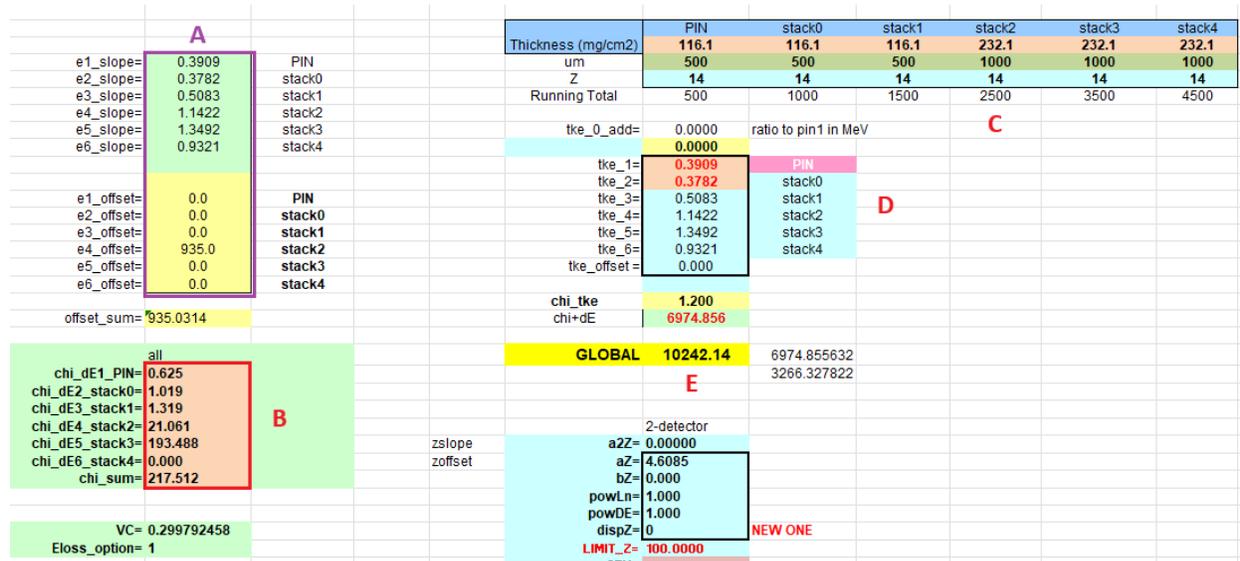


Figure 13: Fitting page (**parameters**) of the total kinetic energy calibration workbook. Box **A** contains the fitted slope and offset parameters, which are ultimately used in `SpecTcl`.

Figure 11 shows an example of the `LISE_eloss` input sheet. The required inputs are highlighted in boxes **A** and **B**. Box **A** contains particle identification parameters and the experimental $B\rho$. Box **B** contains the average channel values from the energy-loss detectors. All remaining cells on this sheet are automatically calculated from these inputs.

After entering the inputs, proceed to the **parameters** sheet (Fig. 13), where the minimization that determines the calibration parameters is performed. The objective of the fit is to transform raw detector channels into energy-loss values in a manner that is internally consistent and physically meaningful.

Detailed algorithmic aspects of the fitting procedure are not discussed here. Users are encouraged to examine the worksheet structure by tracing cell precedents and dependencies (`Formulas` → `Trace Precedents`) to better understand how the calibration is constructed.

In Figure 13, Box **A** contains the adjustable slope and offset parameters to be fitted. Box **B** displays the χ^2 values calculated on the `LISE_eloss` sheet. Box **C** must be verified before minimization begins; it specifies the physical thicknesses of the detectors and directly affects

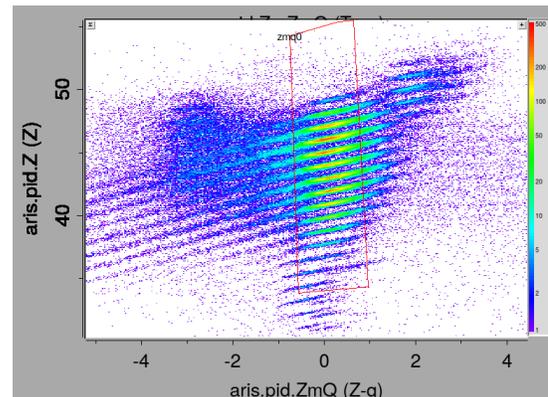


Figure 14: Example of a charge-state gate on fully-stripped ions.

the calculated energy-loss. Box **D** is algebraically linked to Box **A** and must not be modified, as downstream calculations reference these cells. Box **E** contains the summed χ^2 , which serves as the minimization target.

Minimization should be performed iteratively, reducing one χ^2 component at a time while monitoring the behavior of the **GLOBAL** cell. Using the Excel Solver add-in, minimize the selected χ^2 component by varying the parameters in Box **A**. Repeat this process until the total χ^2 is minimized and stable.

Only detectors that register physical energy-loss for the fragment of interest should be included in the minimization. For example, if fragments stop in silicon detector 3, then silicon detector 4 and beyond must be excluded from the fit.

Once the minimization is complete, copy the resulting calibration parameters into the `SpecTcl` calibration file located at `/calibrPID/CSV/db5_si_values.csv`. Reload the data and verify the calibration by inspecting the $Z - q$ vs Z spectrum. For light- Z beams, a well-defined cluster near zero charge-state difference should appear. An example of a well-calibrated spectrum and associated gate are shown in Figure 14. This gate is applied in subsequent data cleaning to isolate the desired charge-state prior to yield extraction.

The contamination of yields by different charge-states is shown in Figure 15. The analysis of charge-state distributions and locations was performed in Ref. [6]. Without proper removal of charge-state contamination, extracted isotopic yields will be skewed.

Once charge-state separation is complete and final cleaning gates are defined, the analysis proceeds to yield extraction, where elemental and isotopic distributions are quantified.

2.2.5 Z Calibration

The final essential PID variable is the atomic number, Z . This quantity is reconstructed from the particle's energy loss in the DB5 detectors. The formal relationship used for the calculation is given in Appendix E.3. In practice, the Z calibration is performed concurrently with the total kinetic energy (TKE) and charge-state calibration described in the previous section, since these quantities are coupled through the velocity dependence of the stopping power.

Once the TKE calibration has been established, the elemental blobs in PID space can be translated into their correct physical positions. The detailed identification procedure is described in Appendix E.6. At this stage, adjustments to the Z calibration parameters serve

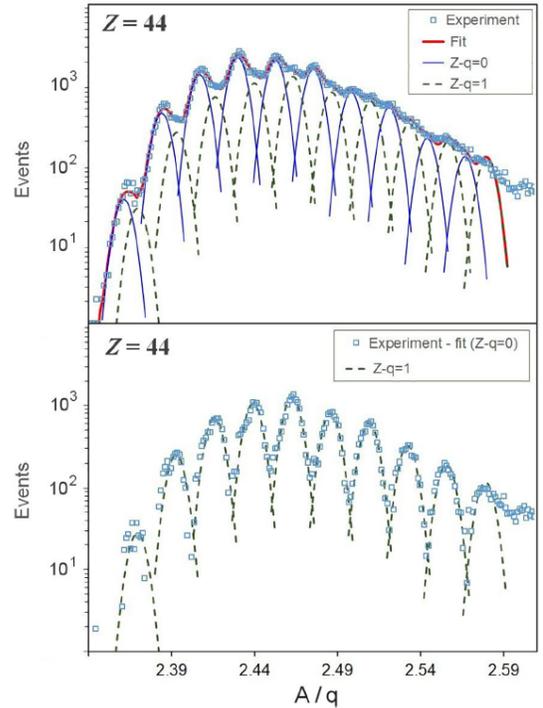


Figure 15: Example of potential charge-state contamination to the yield charge-state statistics. From Ref. [6].

primarily to align the reconstructed PID pattern with known isotopic centroids.

Within the `pid_values` spreadsheet, the calibration parameters are defined explicitly. Equations (19) and (20) show how these parameters enter the reconstruction of the atomic number. In typical operation, $Z_{\log fi}$ is set to unity and $Z_{\text{slope}2}$ is set to zero. The dominant tuning parameters are therefore Z_{slope} and Z_{offset} , which control the overall scaling and translation of the PID pattern in Z . Unless particle identification has been confirmed, large adjustments to Z_{offset} (greater than approximately 0.5 units in Z) are strongly discouraged, as they can mask underlying calibration inconsistencies.

The elemental resolution can also be improved by optimizing the energy-loss input used in the reconstruction. The `pid_values` spreadsheet includes a column specifying the `Z_method`, which determines which energy-loss signal enters Equation (19). In `SpecTcl`, method 0 uses only the PIN detector measurement, while method 1 uses the average of the PIN and silicon stack 0 signals. When both detectors are properly calibrated and functioning reliably, combining their measurements reduces statistical fluctuations and can noticeably improve the elemental peak resolution, as shown in Figure 16. Note that the Z shifts due to calibration differences.

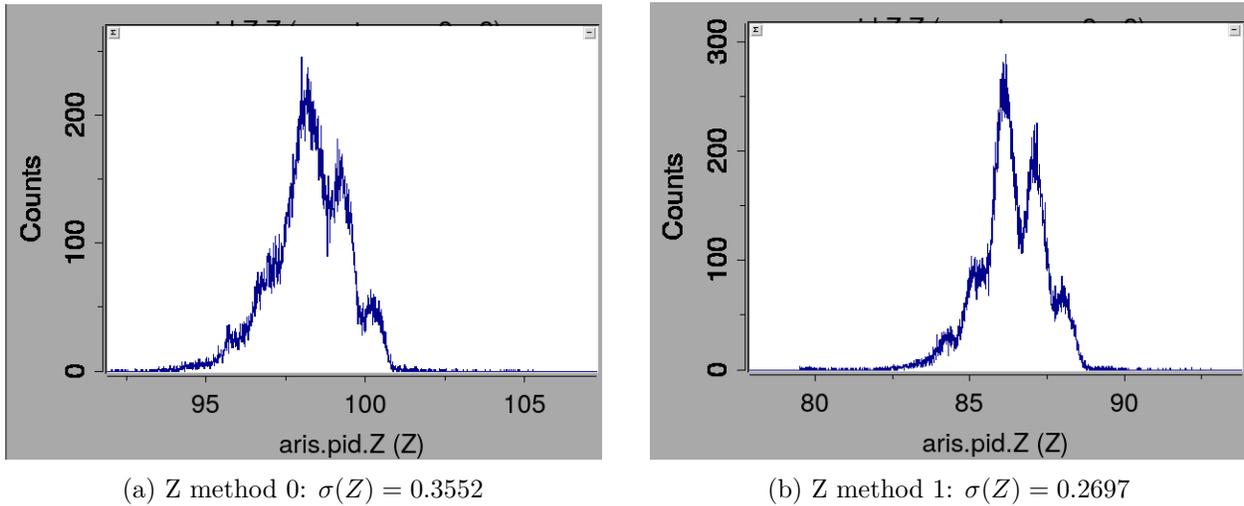


Figure 16: Increased resolution achieved when Z method 1 is used.

2.2.6 Final PID Resolution

As discussed previously, the intrinsic resolution of particle identification is commonly quantified by the widths (σ) of the Gaussian-like distributions observed in PID space. These widths reflect the combined experimental uncertainties in the measured observables used for identification, such as time-of-flight, energy-loss, and total kinetic energy. In addition to the absolute widths of individual PID peaks, the effective resolving power also depends on the separation between neighboring species. This separation is governed by the underlying nuclear properties (A , Z , and charge-state q) as well as the detector and beam-line resolutions that determine how distinctly different species map into PID space. A high-confidence PID is therefore characterized by both small peak widths and sufficient separation between adjacent isotope blobs.

This interplay between peak width and peak separation is illustrated in Figure 7, where gaps between PID clusters significantly improve the reliability of isotope assignment and yield extraction. When such gaps are present, systematic uncertainties associated with peak overlap and background subtraction are substantially reduced, leading to more robust cross section determinations.

A quantitative estimate of the overall PID resolution can be obtained using the `lise64` Excel spreadsheet PID resolution calculator. This tool takes as input the experimentally measured resolutions of the individual observables (ToF, ΔE , and TKE), the assumed identification parameters (A , Z , and q), the magnetic rigidity setting, relevant optical properties of the beam line (including dispersion, magnification, and final image size), and the intrinsic detector resolutions. These contributions are then combined to produce a single, effective PID resolution that reflects the convolution of all experimental and optical effects. This calculated resolution provides a useful benchmark for assessing PID performance and for guiding the optimization of analysis cuts and spectrometer settings.

3 Data Extraction and Isotopic Yields

Once all calibrations and cleaning have been completed, the isotopic yields can be extracted. This extraction procedure assumes that elemental and isotopic peaks are well separated after cleaning and that the residual resolution is dominated by detector response rather than unresolved physical structure. In regions where peaks are severely overlapping or statistics are extremely limited, extracted yields should be treated with additional caution.

Extraction is performed using a dedicated Excel workbook designed specifically for yield determination. Elemental and isotopic peaks are modeled as Gaussian distributions, with all peaks in a given spectrum assumed to share a common width. This reflects the approximately uniform experimental resolution across the relevant charge or mass range after calibration.

3.1 Exporting Data from `SpecTcl`

Begin by creating the one-dimensional atomic number spectra, Z with all previously defined cleaning gates applied. Ensure that the full data range is plotted and that the number of bins is sufficiently large to resolve individual elemental peaks without introducing artificial binning structure. Once the spectrum has been verified, export it from `SpecTcl` using the following command in the command line:

```
swrite -format ascii <fileName><spectrumName>.
```

This command writes the specified spectrum as an ASCII file to the `SpecTcl` root directory.

3.2 Obtaining Yield Statistics

Open the purpose-built Z -extraction Excel workbook, an example of which is shown in Figure 17. This spreadsheet requires the `lise64` Excel file to be open concurrently, as it retrieves Gaussian function and reference information generated by `LISE++`. In the appropriate input fields (blue boxes in Figure 17), specify the axis limits, number of bins, and the relevant

charge-state. These parameters come from the `SpecTcl`-defined spectrum. Copy the raw histogram channel numbers and corresponding yields into columns A and B, respectively.

	Z_fixed	Am3Z	Z_Bins	Area_Fit	Mean_Fit	Delta_Bin	XY_Bin	Yield	Amp	<X>	Ch*Y	<Ch>	Calib_Ch	DeltaZ	DeltaChi	Counts
AxisMin = 0	5	5	5		0.00	0.00	0.00E+00	0.00E+00			0					
AxisMax = 100	6	6	6	10.17	6.34	89.18	8.75E+02	1.38E+02	10.17	6.34	42706	309.464	6.338	-0.338	9.73E-03	138
Bins = 5000	7	7	7	704.08	7.09	2113.55	6.77E+04	9.55E+03	704.08	7.09	3315847	347.136	7.090	-0.090	8.33E-05	9552
k_Am3Z = 0.02	8	8	8	3061.86	8.07	5292.86	3.35E+05	4.15E+04	3061.86	8.07	16468032	396.447	8.075	-0.075	2.75E-05	41539
TotYield = 2.77E+07	9	9	9	22925.55	9.05	14406.51	2.81E+06	3.11E+05	22925.55	9.05	1.38E+08	445.040	9.045	-0.045	3.64E-06	311022
RefChan = 287	10	10	10	76920.10	10.05	36637.19	1.05E+07	1.04E+06	76920.10	10.05	5.17E+08	495.550	10.054	-0.054	2.81E-06	1043545
RefAm3Z = 5.74	11	11	11	284589.97	11.04	92503.35	4.26E+07	3.86E+06	284589.97	11.04	2.1E+09	545.161	11.044	-0.044	9.92E-07	3860921
RefA = 5.889	12	12	12	345881.19	12.04	121902.62	5.65E+07	4.69E+06	345881.19	12.04	2.79E+09	594.909	12.037	-0.037	6.49E-07	4692435
ChanGap = 0.020	13	13	13	446443.68	13.03	130495.64	7.89E+07	6.06E+06	446443.68	13.03	3.9E+09	644.534	13.028	-0.028	3.26E-07	6056727
Zloc = 0	14	14	14	265954.00	14.02	98401.72	5.06E+07	3.61E+06	265954.00	14.02	2.5E+09	694.209	14.020	-0.020	2.14E-07	3608094
ZmQ = 0	15	15	15	279010.61	15.01	72424.92	5.68E+07	3.79E+06	279010.61	15.01	2.82E+09	743.975	15.014	-0.014	9.87E-08	3785228
	16	16	16	106705.22	16.01	38079.54	2.32E+07	1.45E+06	106705.22	16.01	1.15E+09	794.010	16.013	-0.013	1.38E-07	1447628
SigmaA = 0.108	17	17	17	80030.76	17.01	24981.17	1.85E+07	1.09E+06	80030.76	17.01	9.16E+08	843.811	17.007	-0.007	5.08E-08	1085746
Result = 741.611	18	18	18	22236.29	18.00	12364.91	5.43E+06	3.02E+05	22236.29	18.00	2.7E+08	893.434	17.998	0.002	6.67E-09	301671
kn = 0.998	19	19	19	21621.25	18.98	10456.17	5.57E+06	2.93E+05	21621.25	18.98	2.77E+08	942.847	18.985	0.015	4.32E-07	293327
bn = 0.159	20	20	20	8153.40	19.98	9508.24	2.21E+06	1.11E+05	8153.40	19.98	1.1E+08	992.574	19.978	0.022	1.51E-06	110614
Result2 = 29.286	21	21	21	34428.57	20.95	18068.52	9.79E+06	4.67E+05	34428.57	20.95	4.86E+08	1041.523	20.955	0.045	2.97E-06	467079
Fin_Result = 770.897	22	22	22	28969.21	21.94	18952.34	8.62E+06	3.93E+05	28969.21	21.94	4.29E+08	1090.655	21.936	0.064	6.54E-06	393014
	23	23	23	10623.36	22.92	9864.75	3.30E+06	1.44E+05	10623.36	22.92	1.64E+08	1139.931	22.920	0.080	1.69E-05	144123
	24	24	24	738.73	23.90	1867.94	2.39E+05	1.00E+04	738.73	23.90	11913911	1188.776	23.895	0.105	1.10E-04	10022
	25	25	25	41.94	24.80	209.53	1.41E+04	5.69E+02	41.94	24.80	702109	1233.935	24.797	0.203	1.73E-03	569

Figure 17: Z -fitting Excel spreadsheet. Using functions from the `lise64` workbook, Gaussian functions are fit to the data. Set-up information in the colored boxes are taken from `SpecTcl`.

The purpose of this spreadsheet is to model the elemental yield distribution as a sum of Gaussian peaks. All Gaussians are constrained to have a common standard deviation (σ), reflecting the assumption of uniform charge resolution. The only free parameters in the fit are σ , along with a linear slope and offset that control the spacing and absolute positioning of the Gaussian centroids (green boxes in Figure 17). For well-calibrated data, the slope should remain close to unity and the offset close to zero; significant deviations may indicate residual calibration issues. When Excel is performing its fitting, it is suggested to add a constraint that ensures the largest peak in the elemental distribution remains fixed at the centroid calculated in `SpecTcl`. This stabilizes the fit and prevents unphysical drift of the reference peak during minimization. All remaining quantities in the spreadsheet are calculated automatically.

The Excel Solver is used to minimize the difference between the experimental spectrum and the calculated Gaussian sum. Two error functions are provided, and adjustable weights may be applied to each. Varying the relative weights can improve convergence or stability of the fit, particularly in spectra with low statistics or partially overlapping peaks. As with the total kinetic energy calibration workbook, using Excel's precedent and dependent tracing tools can be helpful for understanding how changes propagate through the calculation.

A successful fit is characterized by a calculated distribution that closely reproduces the experimental spectrum. An example is shown in Figure 18. The solid gray curve represents the experimental data, while the dotted red curve shows the fitted Gaussian sum. The dashed green line illustrates an example $\pm 3\sigma$ gate applied to an elemental peak.

In some cases, the Gaussian fitting procedure may fail to converge or may produce unphysical results. The most common indicators of a problematic fit include:

- unusually large or small values of σ ,
- significant deviations of the slope from unity,

- offsets far from zero, or
- visible discrepancies between the calculated and experimental distributions.

If the fitted slope or offset deviates substantially from expected values, this often indicates residual calibration issues in the upstream analysis. In such cases, the relevant calibration spectra should be revisited before proceeding.

Poor convergence or unstable fits can also result from low statistics or overlapping peaks, particularly near the edges of the acceptance. Adjusting the relative weights of the error functions in the Solver may improve stability. In some situations, temporarily fixing the slope and offset to their nominal values and fitting only for σ can help diagnose whether the instability arises from resolution effects or from misaligned peak centroids.

Finally, visual inspection remains an essential diagnostic tool. Even when the numerical minimization converges, the fitted distribution should always be compared directly to the experimental spectrum to ensure that the Gaussian model provides an adequate description of the data.

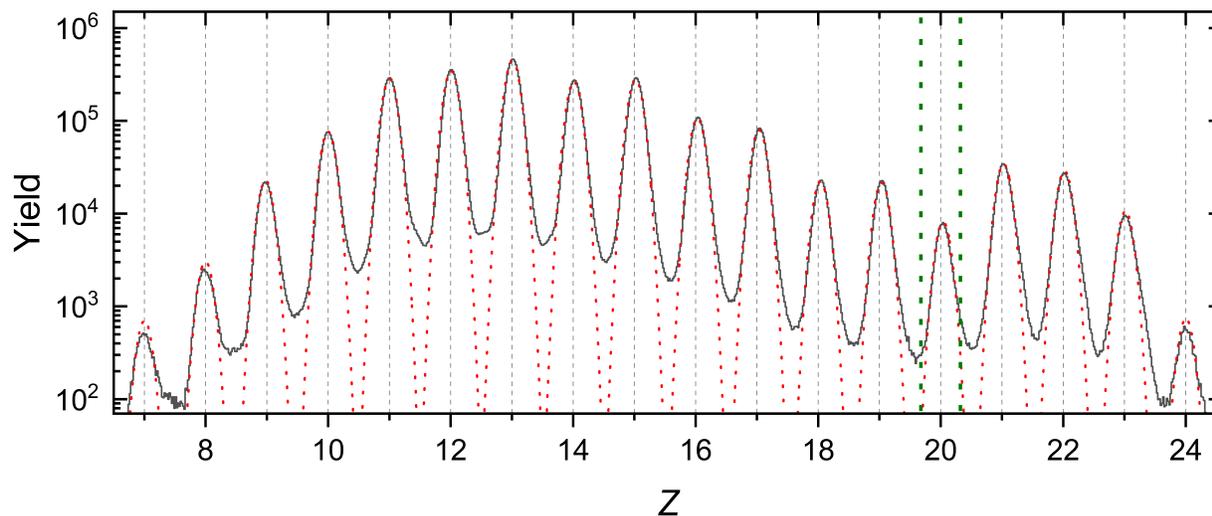


Figure 18: Z distribution (solid gray line) with calculated Gaussian (dotted red line) overlaid. Dashed green line represents a $\pm 3\sigma$ gate.

Once the optimal value of σ has been determined, define elemental gates in `SpecTcl` corresponding to $\pm 3\sigma$ around the mean of each fitted peak. These gates define the experimental elemental widths used in all subsequent yield extractions. A recommended approach is to first place approximate gates visually using `SpecTk`, then manually edit the gate limits in the `SpecTcl Gates` tab using the calculated ($\pm 3\sigma$) values.

3.3 PID Deduced Integer Values

Isotopic distributions are then generated for each element by applying an `And` gate combining the cleaning gates with the corresponding Z gate. These distributions are projected onto a

one-dimensional $A - 3Z$ (for neutron-rich fragments) or $A - 2Z$ (for proton-rich fragments) axis.

In `SpecTc1`, the quantity labeled $A - 3Z$ is constructed from the measured A/q and Z assuming a known charge-state q , according to

$$A - 3Z = ((A/q) - 3) Z. \quad (1)$$

This expression is exact for fully stripped ions ($q = Z$). For partially stripped ions, the resulting value no longer corresponds strictly to $A - 3Z$, but remains a useful PID coordinate, as different charge-states map to distinct loci in the projection.

This projection choice of $A - 3Z$ is motivated by both physical and practical considerations. For neutron-rich fragments, mass and charge are strongly correlated along trajectories of approximately constant $A - 3Z$. Projecting onto this variable effectively aligns isotopes of a given element along a single axis, reducing peak overlap and improving separation in regions where statistics may be limited. Similarly, proton-rich fragments are more cleanly separated using the $A - 2Z$ projection, which better follows the dominant correlation direction in that region of the fragment distribution. These projections also mitigate distortions arising from residual correlations between mass, charge, and time-of-flight, which can broaden peaks when projecting directly onto A . By choosing a projection variable that is approximately orthogonal to the dominant correlation direction, the resulting one-dimensional spectra exhibit narrower, more symmetric peaks that are better described by Gaussian distributions. This improves both fit stability and the reliability of extracted yields.

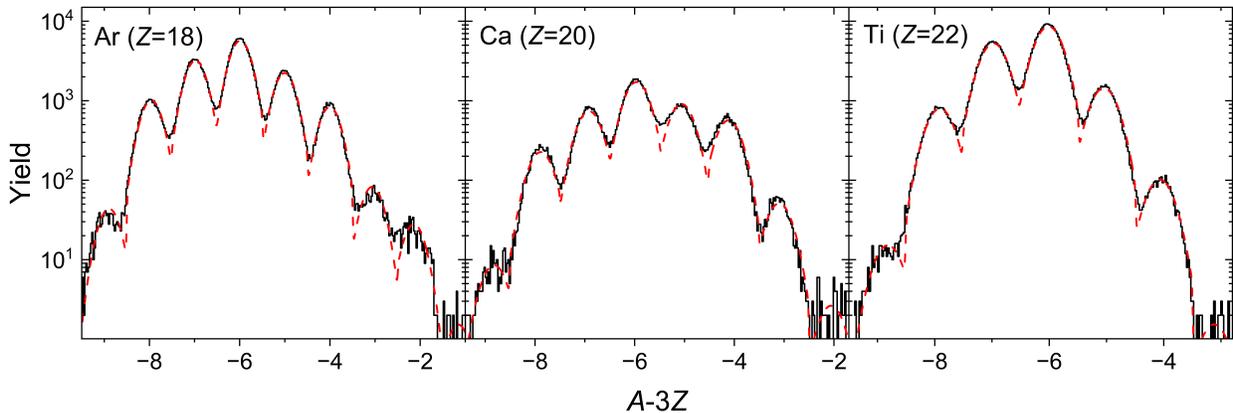


Figure 19: $A - 3Z$ distribution (solid gray line) with calculated Gaussian (dotted red line) overlaid for three elements.

In addition to improving peak separation, projections onto $A - 3Z$ (neutron-rich) and $A - 2Z$ (proton-rich) partially cancel correlated uncertainties in the reconstructed values of A and Z , further reducing the effective variance of the projected variable, as both quantities are derived from overlapping experimental observables.

As with the elemental spectra, the isotopic distributions are exported from `SpecTc1` and analyzed using the same Excel fitting framework as the elemental yields to extract Gaussian parameters. An example of an $A - 3Z$ spectrum with fitted Gaussian peaks is shown in

Figure 19. This representation illustrates the improved separation and symmetry of isotopic peaks obtained using the $A - 3Z$ projection.

The yield for each isotope is defined as the area under the corresponding fitted Gaussian peak, as reported in the green-highlighted column in Figure 17. Conversion to atomic mass A is performed using the gated charge-state. In most cases, the measured charge-state satisfies $Z = q$, allowing A to be obtained directly from the fitted value of $A - 3Z$. For fragments where multiple charge-states are present, only the charge-state selected by the Z gate should be used in the conversion.

Y	absolute		relative		\sqrt{Y}
	$-dY$	$+dY$	$-dY$	$+dY$	
1	0.63	1.75	-63%	175%	1.00
2	1.26	2.25	-63%	113%	1.41
3	1.90	2.30	-63%	77%	1.73
4	1.66	2.78	-42%	70%	2.00
5	2.25	2.81	-45%	56%	2.24
6	2.18	3.28	-36%	55%	2.45
7	2.75	3.30	-39%	47%	2.65
8	2.70	3.32	-34%	42%	2.83
9	2.67	3.79	-30%	42%	3.00
10	3.22	3.81	-32%	38%	3.16
11	3.19	3.82	-29%	35%	3.32
12	3.17	4.29	-26%	36%	3.46
13	3.72	4.30	-29%	33%	3.61
14	3.70	4.32	-26%	31%	3.74
15	3.68	4.32	-25%	29%	3.87
16	3.67	4.80	-23%	30%	4.00
17	4.21	4.81	-25%	28%	4.12
18	4.19	4.82	-23%	27%	4.24
19	4.18	4.82	-22%	25%	4.36
20	4.17	5.30	-21%	27%	4.47
21	4.58	4.58	-22%	22%	4.58

Table 1: Absolute and relative errors for small statistics as a function of Y . Adapted from [7]

Depending on the length of a given isotopic chain, the result block in the spreadsheet may need to be extended. Care should be taken to expand all relevant ranges uniformly to avoid truncation or misalignment of calculated values.

3.4 Counting Statistics

After the yields have been extracted, statistical uncertainties should be assigned. For yields greater than 20 counts, the standard approximation of \sqrt{Y} , where Y is the number of counts, provides a reasonable estimate of the uncertainty. For yields below 20 counts, asymmetric

confidence intervals derived from Poisson statistics should be used instead, as summarized in Table 1. These asymmetric error bars are a result of the residual background that exists from high- σ tails of neighboring peaks, finite detector resolution, unresolved charge-states, and reactions in beamline materials. Details of these sources of background are discussed further in paragraph 2.2.3.1.

These yield uncertainties form the statistical component of the cross section uncertainty, which is propagated together with transmission, target thickness, and beam integral uncertainties as described in section 6.

4 Experimental Inputs for Cross Section Determination

Calculating production cross sections requires two key experimental inputs: the effective target thickness and the integrated beam power. The determination of these quantities is discussed in the following subsections.

4.1 Target Thickness

During the experiment, the target and wedge thicknesses are determined using energy-loss measurements with the primary beam. The beam energy is known from the accelerator settings, and its passage through material results in a predictable energy-loss that depends on the areal density of the material. A more detailed discussion of this method is provided in subsection E.3.

Briefly, the primary beam is first transported unobstructed from the accelerator to the first dispersive focal plane of the separator. The beam is tuned to be centered at the viewer position. When the target is inserted, the beam loses energy in the material. This energy-loss changes the magnetic rigidity of the beam, causing it to shift along the dispersive axis (the vertical y direction in the pre-separator).

To re-center the beam at the focal plane, the magnetic rigidity $B\rho$ is retuned. The required change in $B\rho$ is directly related to the energy-loss in the target and therefore to the effective target thickness. The same procedure is applied to determine the thicknesses of any wedges used in the separator.

The results of these measurements are recorded in the corresponding LISE++ file for the run and are typically also documented in the laboratory logbook (Olog). These values should be used consistently in subsequent cross section calculations.

4.2 Beam Integral

The beam integral quantifies the total number of primary beam particles incident on the target during the experiment. It is calculated by integrating the measured beam power over the time intervals corresponding to active data collection¹⁴.

¹⁴Instructions for extracting the raw data has been compiled in the help sheet. See [/projects/lisedev/Literature/Cross Section Manual](#)

An Excel workbook has been developed to streamline this calculation. The workflow consists of the following steps:

1. The workbook searches for beam power measurements whose timestamps fall between the start and stop times of each data run, as recorded in the experimental log (Figure 20).
2. These measurements are collated into two columns: elapsed time and corresponding beam power.
3. A trapezoidal integration is performed to sum the total beam power over the run. Negative power readings, which are artifacts of the measurement system, must be set to zero before integration.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
1		Index	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
2	# Time	ACS_DIAG:BPM_ENGY7E:POWER_RD Value	3189	3190	3191	3192	3193	3194	3195	3196	3197	3198	3204	3205	3206	3207	3208
18	33:24.5	18640.251	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
19	33:24.7	18640.252	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
20	33:34.5	18689.31	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
31	4/8/2025 23:33	18689.318	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
32	4/8/2025 23:33	18744.618	18744.62	0	0	0	0	0	0	0	0	0	0	0	0	0	0
33	33:44.5	18744.625	18744.63	0	0	0	0	0	0	0	0	0	0	0	0	0	0
34	33:54.3	18815.95	18815.95	0	0	0	0	0	0	0	0	0	0	0	0	0	0
35	33:54.5	18815.943	18815.94	0	0	0	0	0	0	0	0	0	0	0	0	0	0
36	34:04.3	18927.056	18927.06	0	0	0	0	0	0	0	0	0	0	0	0	0	0
37	34:04.5	18927.05	18927.05	0	0	0	0	0	0	0	0	0	0	0	0	0	0
38	34:14.3	19077.57	19077.57	0	0	0	0	0	0	0	0	0	0	0	0	0	0
39	34:14.5	19077.576	19077.58	0	0	0	0	0	0	0	0	0	0	0	0	0	0
50	34:24.3	19110.081	19110.08	0	0	0	0	0	0	0	0	0	0	0	0	0	0
51	34:24.5	19110.074	19110.07	0	0	0	0	0	0	0	0	0	0	0	0	0	0
52	34:34.3	19096.91	19096.91	0	0	0	0	0	0	0	0	0	0	0	0	0	0
53	34:34.5	19096.903	19096.9	0	0	0	0	0	0	0	0	0	0	0	0	0	0
54	34:44.3	19180.202	19180.2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
55	34:44.5	19180.261	19180.26	0	0	0	0	0	0	0	0	0	0	0	0	0	0
56	34:54.3	19428.265	19428.27	0	0	0	0	0	0	0	0	0	0	0	0	0	0
57	34:54.5	19428.271	19428.27	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 20: Excel workbook that collates the beam power measurements falling between the start and stop timestamps of data collection runs.

After compiling the beam power time series, it is good practice to verify the consistency of the elapsed run time. Convert the start and stop timestamps of the data run to elapsed seconds and compare this duration with the sum of the intervals between consecutive power measurements.

A key consistency check is that these two quantities agree within expected rounding or sampling limits. Significant discrepancies may indicate missing power readings, gaps in the data acquisition, or misalignment between the recorded run times and the power measurement timestamps. Performing this check ensures that the integrated beam accurately reflects the total number of primary beam particles incident on the target. This is critical because any mismatch directly affects the precision of the derived cross sections.

The output of this procedure is a time series of beam power and the cumulative integrated energy delivered to the target, shown graphically in Figure 21. Beam power is recorded in watts (W). Integrating the power over time yields the total energy delivered during the experiment. This energy can then be converted to the number of incident beam particles

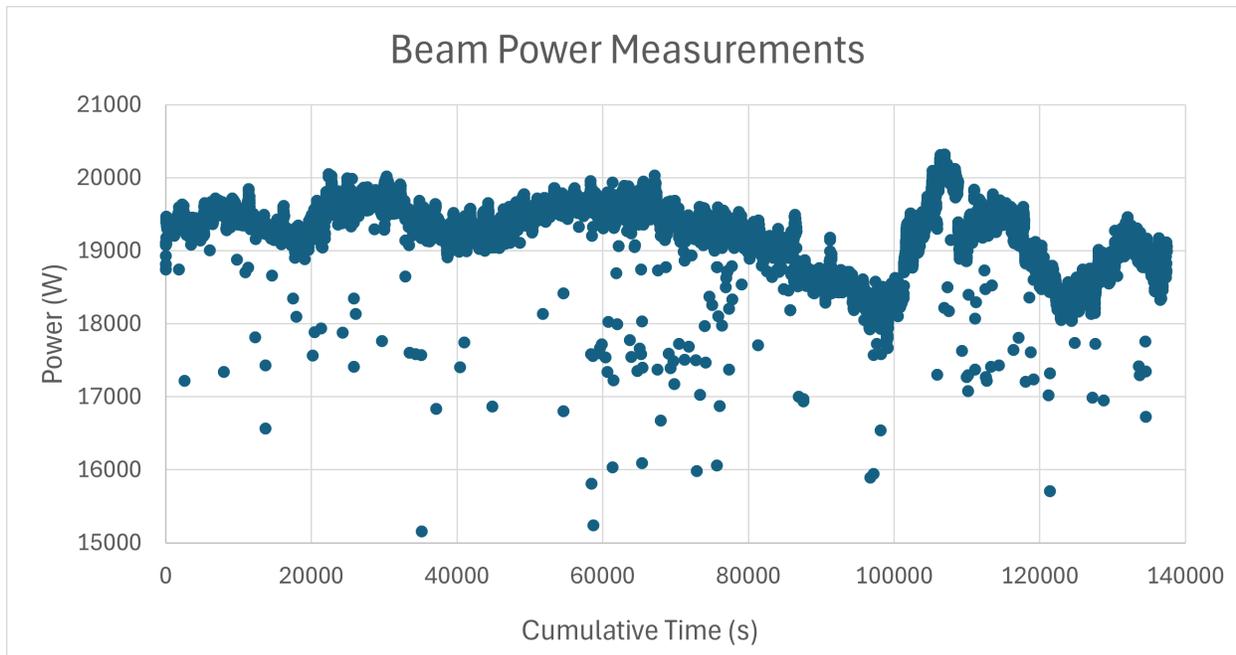


Figure 21: Resulting beam power graph from the beam integral workbook (Figure 20). A trapezoidal integration of this curve yields the total energy delivered to the target.

using the known beam energy and LISE^{++} . It is important to ensure unit consistency when using these values in cross section calculations.

To estimate the uncertainty in the beam integral, a conservative approach is to treat the largest negative power reading (assumed to be an artifact of the measurement system) as the uncertainty for each measurement. In most cases, this yields a negligible contribution to the integrated beam. Nevertheless, performing this calculation is a valuable exercise for first-time researchers, helping them understand the propagation of errors in cross section measurements.

5 LISE^{++} Transmission Analysis

A key component of cross section analysis is determining the transmission from the target to the final focal plane. While the number of counts observed at the final focal plane may be small, the number of particles produced at the target is generally much larger. Correcting for losses between production and detection ensures a more accurate determination of the cross section.

LISE^{++} is a versatile tool for predicting fragment yields and beam characteristics [8–11]. The goal of transmission analysis is to reproduce experimentally observed fragment rates. For isotopes with previously measured cross sections, ensure the corresponding experimental data are loaded into LISE^{++} . Published datasets are included with the package, while theoretical predictions are used for isotopes lacking measurements.

5.1 Procedure

Start by setting the primary beam rate in the segmented experiment file to the calculated beam integral. Then calculate the rate and transmission of all fragments and compare these results to the experimental yields. Initial discrepancies are expected; iterative adjustments are typically required to achieve reasonable agreement.

Parameters commonly adjusted include:

- Target and wedge thicknesses
- Material defects and lateral wedge offsets

Advanced parameters, such as optics settings at DB1, should be modified only after all reasonable material-based adjustments have been explored. For the expert user and research into momentum distributions, the parameters that govern the LISE⁺⁺ momentum distributions may constitute another knob to adjust.

Physical plausibility check: Always ask, “Are my changes physically reasonable?” Small adjustments within experimental uncertainties are valid. For example, changes in target or wedge thicknesses exceeding 5%, or slit positions exceeding 1 mm, are likely unphysical. After each adjustment, record the transmission and rates for all fragments to track trends and sensitivities.

A special word of warning for the momentum slits in the pre-separator: when they are set asymmetrically, the reference coordinates become important. In the saveset, they are set according to physics coordinates; in LISE⁺⁺, they are defined according to the beam coordinates. This amounts to a reversal of the slit positions (e.g., -75 +60 → -60 +75).

Make special note of the target thickness defect that is set for the baseline transmission; this will serve as the uncertainty in the thickness measurement for the cross section error propagation.

5.2 Distinguishing Primary and Secondary Fragments

Not all fragments reaching the final focal plane are primary products from the target. Contamination of the beam can arise from secondary reactions in the target or reactions in wedge materials. These unwanted contaminants either mimic the trajectory of the desired fragments, or are transported because of higher-order optical effects. LISE⁺⁺ provides tools to identify these contributions:

- Secondary reactions in the target (`Physical Models` → `Secondary Reactions in the Target`)
- Reactions in materials (`Calculate reactions in this material` in the material dialog)

Comparing simulations with these options enabled or disabled highlights fragments influenced by secondary processes. Fragments dominated by secondary or material-reaction contributions are shown in light gray in the isotope box (Figure 22, e.g., ⁶⁸Cr) and may be excluded from cross section calculations.



Figure 22: LISE⁺⁺ calculations for rate and transmission with secondary and wedge reactions enabled.

Higher-order optical effects can be assessed using Monte Carlo simulations. In the main LISE⁺⁺ window, calculate the transmission of fragments in the region of interest. Then, go to the Monte Carlo dialog, and set the following options: in the upper left corner, select **Group of isotopes already calculated**; make sure the **After BLOCK** is set to the DB5 slits; set the X-coordinate to be $A - 3q$ and the y-coordinate to be Z . Choose the optics order as appropriate.

Calculate the Monte Carlo distribution to first order. Once sufficient statistics have been collected, stop the calculation and click **MC transmission statistics** on the left side of the dialog. This will bring up the number of counts predicted to reach DB5 as well as the transmission. Save the results and repeat the calculation for fifth-order optics. When comparing the first- and fifth-order transmission, any large increases (decreases) hint at aberrations in the optics boosting (suppressing) the transmission of the respective fragments and mark the isotopes as candidates for exclusion.

5.3 Estimating Transmission Uncertainty

Once primary fragments are identified and LISE⁺⁺ settings reproduce experimental yields (Figure 23), transmission uncertainties can be estimated. These arise from deviations between the simulated parameters and experimentally measured values.

For each adjusted parameter (e.g., target or wedge thickness, slit position), vary the value slightly above and below the nominal setting. Recalculate fragment transmission for these variations. The difference between these values and the nominal transmission defines the contribution of that parameter to the overall uncertainty (Figure 24). For example,

^{67}Cr 's transmission is greatly suppressed when the pre-separator wedge's defect is smaller than what was set (2%). Similarly, its transmission is enhanced when the pre-separator wedge offset is 1mm greater than was set (2mm).

Once all relevant input parameters have been varied and the corresponding transmission values recorded, the total transmission uncertainty is evaluated. This is accomplished using quadrature addition in logarithmic space, which is appropriate given the multiplicative nature of transmission effects. In practice, this calculation is performed using the `lise64` function `S_ArrayAsymErrorLog_Transpose_Quadrature`. This routine takes as input the baseline (mean) transmission value, the associated positive and negative deviations for each parameter variation, and an optional weighting factor, and returns the combined asymmetric uncertainty on the transmission.

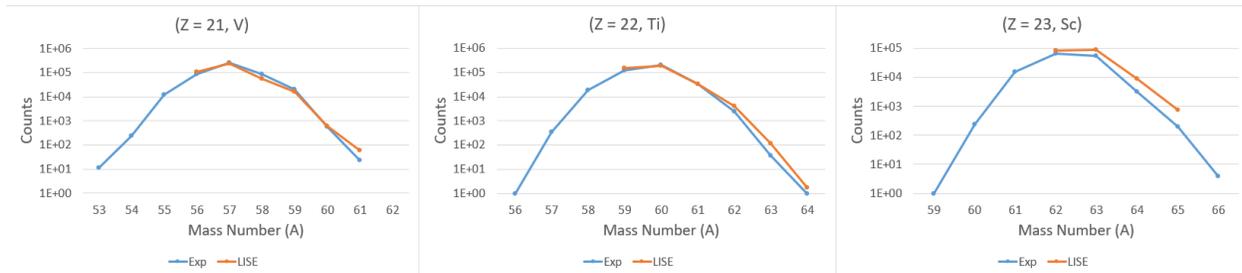


Figure 23: Comparison of experimental yields to final LISE^{++} -predicted yields for primary fragments.

Z*1000+N	Index	A El	PSWedge Defect=1% Trans	PSWedge Defect=3% Trans	PSWedge Offset=1 Trans	PSWedge Offset=3 Trans	Target Defect=1% Trans	Target Defect=3% Trans	Target Thickness +1% Trans	Target Thickness -1% Trans
240044	68Cr									
240043	67Cr		3.37E-05	1.97E-01	6.79E-03	3.20E-01	6.87E-02	6.88E-02	6.06E-02	7.73E-02
240042	66Cr		2.08E-01	3.69E-01	1.12E-01	7.52E-01	3.07E-01	3.39E-01	2.59E-01	3.80E-01
240041	65Cr		6.97E-02	6.87E-02	2.61E-02	1.58E-01	6.34E-02	8.69E-02	5.66E-02	7.37E-02
240040	64Cr		1.80E-03	1.55E-03	7.78E-04	3.28E-03	1.27E-03	2.91E-03	1.20E-03	2.87E-03
240039	63Cr									
240038	62Cr									
230044	67V									
230043	66V			1.08E-01	3.99E-06	4.26E-02	1.85E-03	1.84E-03	1.84E-03	1.86E-03
230042	65V		3.32E-01	2.08E+00	5.20E-01	2.71E+00	1.36E+00	1.35E+00	1.33E+00	1.37E+00
230041	64V		2.29E+00	3.20E+00	1.73E+00	4.70E+00	3.19E+00	3.18E+00	3.08E+00	3.37E+00
230040	63V		2.40E+00	1.91E+00	1.29E+00	3.21E+00	2.17E+00	2.25E+00	1.92E+00	2.24E+00
230039	62V		3.59E-01	2.61E-01	2.14E-01	4.46E-01	2.78E-01	3.76E-01	2.63E-01	4.03E-01
230038	61V		8.53E-03	5.26E-03	4.79E-03	7.94E-03	5.33E-03	9.32E-03	6.57E-03	9.17E-03

Figure 24: Calculation of transmission for various changes in the final LISE^{++} transmission file. These changes correspond to the uncertainty of set parameters. The largest transmission change in a given fragment will constitute the uncertainty for that fragment.

As a concrete example, the transmission uncertainty for ^{67}Cr is shown in Figure 25, which utilizes the Excel workbook `Error_analysis_v7_x64`. The graphs labeled **A** and **B** correspond to the constructed distribution widths, assuming that the error bars constitute σ . The same distributions are shown on linear and log scales, respectively. Graph **C** shows the

input parameters (1st-4th init; these correspond to the transmission deviations illustrated in Figure 24: PS wedge defect, PS wedge offset, target defect, and target thickness, respectively) and the various weighting schemes the workbook provides. In the workbook set up, all weights should be set to unity, with the left σ representing the negative error and the right σ representing the positive error.

For ^{67}Cr , the resulting transmission uncertainty spans a couple orders of magnitude. Such a large uncertainty might render the extracted cross section meaningless, as the final error bars dominate the result. The final determination is made only after the cross section and its corresponding errors have been calculated.

As this example illustrates, transmission effects frequently constitute the dominant source of uncertainty in rare isotope production cross section measurements. To obtain physically meaningful results, the parameter variations used in the transmission uncertainty evaluation must therefore remain within reasonable and defensible physical bounds:

- Target thickness: variations of $\sim 1\%$ are typical
- Optics settings: variations beyond 1–2% are usually unrealistic

Using excessively large changes in parameters will artificially inflate the estimated uncertainties. This, in turn, will artificially inflate the propagated uncertainties on the cross sections.

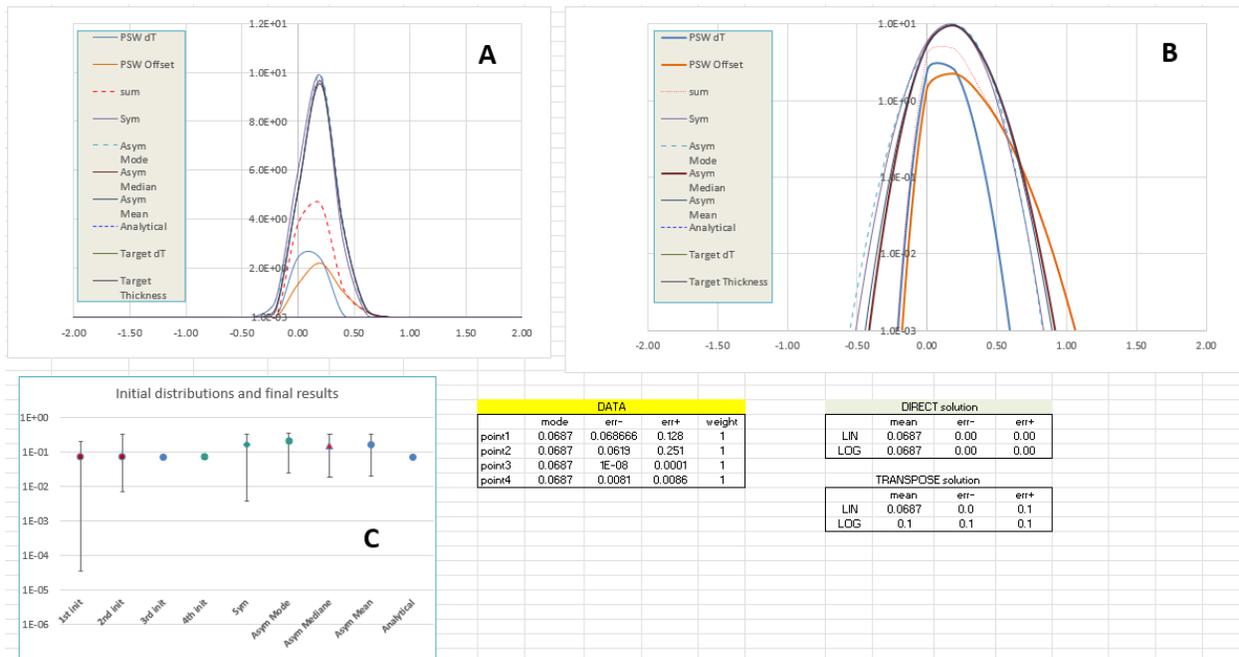


Figure 25: Transmission uncertainty components for ^{67}Cr (parameter indices 1–4) and the final combined transmission uncertainty for various methods.

A final winnowing of fragments is done using the transmission. Certain isotopes that are determined to be predominantly from the target and whose counts are not significantly increased by higher-order optics aberrations, may have low predicted transmission. Fragments

falling into this category – whose transmission is below 0.1% – are generally excluded from cross section analysis on the grounds that they exist on the edge of the fragment separator acceptance. As such, these particular fragments are highly sensitive to small changes in the system, such as material thicknesses or defects (^{67}Cr , in the above example). This translates directly into error bars that can span several orders of magnitude.

It should be mentioned that fragments with low transmissions do not necessarily have low count rates: a large cross section can more than make up for small transmission! Conversely, the opposite is true: high transmission can be accompanied by low statistics. This is especially the case for extremely neutron-rich isotopes.

Experience is essential. First-time analyses may require significant iteration to converge on physically reasonable parameters and transmission uncertainties. As familiarity with the process grows, the procedure becomes faster and more intuitive.

6 Cross Section Calculation

With all requisite components determined, the cross section can now be calculated. An introduction to cross sections is presented to ground the reader and give important context and motivation for the quantities derived earlier in the manual.

6.1 From First Principles to Measured Yields

The production cross section for a given fragment is fundamentally a measure of the probability that an interaction between an incident beam and a target nucleus will produce that fragment. In nuclear physics, the cross section σ is defined such that it relates the number of observed reactions to the incident particle flux and target properties. At a conceptual level, the cross section has units of area and quantifies an effective “interaction area” presented by the target to the incoming beam; this is why nuclear cross sections are traditionally expressed in barns ($1 \text{ b} = 10^{-24} \text{ cm}^2$), corresponding roughly to the geometric area of a nucleus at typical energies.

In a differential form, the two-dimensional distribution of produced fragments in momentum and angle is expressed as the double-differential cross section

$$\frac{d^2\sigma}{d\Omega dP}(\theta, P), \quad (2)$$

which expresses the probability per unit solid angle $d\Omega$ at laboratory angle θ and per unit momentum interval dP that a reaction producing a given fragment will occur. The experimentally measured count rate for fragments in a particular kinematic bin is then related to this differential cross section by

$$N_{\text{meas}}(\theta, P) = N_{\text{beam}} N_{\text{tgt}} \frac{d^2\sigma}{d\Omega dP}(\theta, P) \Delta\Omega(\theta) \Delta P \varepsilon(\theta, P), \quad (3)$$

where

- $N_{\text{meas}}(\theta, P)$ is the number of fragments detected in a finite solid angle $\Delta\Omega$ around θ and momentum interval ΔP ,

- N_{beam} is the number of incident beam particles,
- N_{tgt} is the number of target nuclei presented per unit area (target thickness times number density),
- $\Delta\Omega(\theta)$ is the angular acceptance of the fragment detection system,
- ΔP is the momentum acceptance corresponding to the selected kinematic window, and
- $\varepsilon(\theta, P)$ is the overall detection efficiency (including geometrical acceptance, detector efficiency, and analysis cuts).

This expression shows that the number of observed counts in a detector is built up from:

- the available phase space for the fragment (solid angle and momentum interval),
- the intrinsic likelihood of the reaction (the differential cross section), and
- the experimental acceptance and efficiency.

In many practical production experiments – such as those described in this manual – the acceptance in momentum and angle is effectively integrated over by the ion optics, and the detector acceptance and gate efficiencies are folded into a single transmission factor ξ that corrects the observed yield for acceptance and loss effects¹⁵. Likewise, the momentum dependence is often not explicitly resolved in the final yield extraction. Mathematically, this means the extracted cross sections are partial cross sections; see subsection 7.1 and Appendix F for more details. Under these common experimental conditions, the integrated yield Y for a given fragment can be written approximately as

$$Y \approx N_{\text{beam}} N_{\text{tgt}} \sigma, \quad (4)$$

where σ now represents a suitably integrated production cross section. Rearranging this relation and expressing N_{beam} and N_{tgt} in terms of the beam integral B_{int} (the total number of incident particles) and the target thickness t (atoms/cm²) yields the simplified experimental formula used in this manual:

$$\sigma \approx \frac{Y}{\xi t B_{\text{int}}}. \quad (5)$$

Here, ξ absorbs the effects of finite acceptance, efficiency losses, and transmission corrections. The target thickness t is directly proportional to the number of available target nuclei per unit area, and B_{int} represents the total incident beam exposure during data collection.

The connection between the differential form and the simplified form can thus be understood as an integration over the experimentally accessible phase space, with detector acceptance and efficiency corrections folded into a transmission factor.

Before calculating the cross sections with Equation (5), ensure that the units on the parameters are consistent: yield and beam integral should be in counts (unitless), target thickness should be in atoms/cm², and transmission as a unitless decimal.

¹⁵This is the purpose of the LISE⁺⁺ transmission calculation in section 5

Cross sections are generally reported in millibarns (mb). If Equation (5) is evaluated using CGS units (cm^2), multiply the result by 10^{27} to convert to millibarns.

Each of these quantities carries an uncertainty, and propagating these uncertainties to the final cross section requires careful attention. It is important to distinguish between the three types of uncertainties that compose the cross section uncertainty: the universal systematic uncertainty, which affects all fragments the same (beam integral and target thickness), isotope-dependent systematic uncertainty (transmission), and statistical uncertainty (yield counts). These three sources of uncertainty arise from different places and are treated differently.

Assuming the uncertainties are independent and sufficiently small, standard error propagation gives:

$$\frac{\Delta\sigma}{\sigma} = \sqrt{\left(\frac{\Delta Y}{Y}\right)^2 + \left(\frac{\Delta\xi}{\xi}\right)^2 + \left(\frac{\Delta t}{t}\right)^2 + \left(\frac{\Delta B_{\text{int}}}{B_{\text{int}}}\right)^2}. \quad (6)$$

Once all components of the uncertainty have been evaluated, Equation (6) defines the uncertainty on the final calculated cross section. The `lise64` function `S_ArrayAsymErrorLog_Transpose_Quadrature` is used to compute the final cross section uncertainties in logarithmic space. The function takes as input the central (nominal) values of each cross section component, the corresponding positive and negative uncertainties for each component, and an optional weighting factor. For cross section uncertainty propagation, this weighting option should not be used.

This procedure ensures that all known sources of uncertainty – statistical, experimental, and model-dependent – are incorporated into the reported cross sections. Documenting each contribution separately is strongly recommended, as it facilitates transparent error analysis and meaningful comparison with previous measurements.

6.2 Uncertainty Redux

For convenience, the primary sources of uncertainty and the procedures used to estimate them are summarized below.

- **Isotope Yield Y**

- For $Y > 20$, the uncertainty can be approximated as $\Delta Y = \sqrt{Y}$.
- For $Y \leq 20$, Poisson statistics must be applied using the specialized corrections summarized in Table 1.

- **Transmission ξ**

- $\Delta\xi$ is estimated by recalculating the transmission after perturbing each adjusted simulation parameter (e.g., target thickness, wedge thickness, wedge defects) within physically reasonable bounds.
- For each isotope, the maximum deviation from the nominal transmission is taken as $\Delta\xi$, as illustrated in Figure 24.

- **Target Thickness t**
 - The uncertainty on the target thickness is defined by the defect parameter used in the baseline LISE⁺⁺ transmission calculation.
 - Physically, this uncertainty reflects manufacturing tolerances, non-uniformity across the target area, and measurement error.
- **Beam Integral B_{int}**
 - When the beam power integration procedure is correctly performed, the uncertainty on the integrated beam is typically small.
 - A conservative estimate treats the largest negative fluctuation in the measured power as the uncertainty on each data point; trapezoidal integration then propagates this uncertainty to the total beam integral.

7 Interpretive Analysis

At this point, a reader following this manual has successfully extracted production cross sections and quantified their associated uncertainties directly from experimental data. This marks the completion of the experimental analysis phase. The cross sections themselves represent the final experimental observables. All subsequent analysis consists of interpretation, comparison, and modeling.

The remainder of this manual focuses on placing the newly measured cross sections in a broader physical context. In particular, this section guides a first-time researcher through the two most common post-analysis studies performed with production cross sections: abrasion–ablation modeling [12] and Δ BE systematics [5]. For each, the theoretical motivation, analysis goals, and practical implementation are discussed in turn.

Before proceeding to model comparisons, a crucial first diagnostic step is to visualize the extracted cross sections. The recommended starting point is to plot the cross sections grouped by atomic number Z as a function of mass number A . Along a given isotopic chain, the cross sections should vary smoothly as a function of A . Sudden discontinuities or sharp kinks in these trends often indicate problems upstream in the data analysis, such as misidentified fragments, over-cleaning, or incorrect transmission corrections.

There are, however, important and physically meaningful exceptions to this expectation. Odd–even staggering effects are commonly observed, particularly in neutron- or proton-rich regions. Unbound nuclei produce genuine gaps in isotopic chains, and shell or deformation effects can introduce local structure in otherwise smooth trends. An example of such behavior is shown in Figure 26 for the heaviest isotopes in the titanium isotopic chain.

In some cases, individual cross sections may be excluded from further analysis. This is most often necessary for fragments with very large relative uncertainties, which typically arise from small transmissions or extremely low statistics. Exclusion should be conservative and well-documented; questionable points are often more valuable as diagnostics than as final results.

Once the internal consistency of the data set has been verified, comparisons to previous measurements should be performed. Plot the new results alongside published cross sections

for the same isotopes and perform an initial visual comparison. Agreement within quoted uncertainties is expected but not guaranteed. If significant discrepancies are observed, the affected isotopes should be flagged and the analysis carefully reviewed. Subtle errors – particularly in PID, transmission, or yield extraction – often only become apparent at this stage.

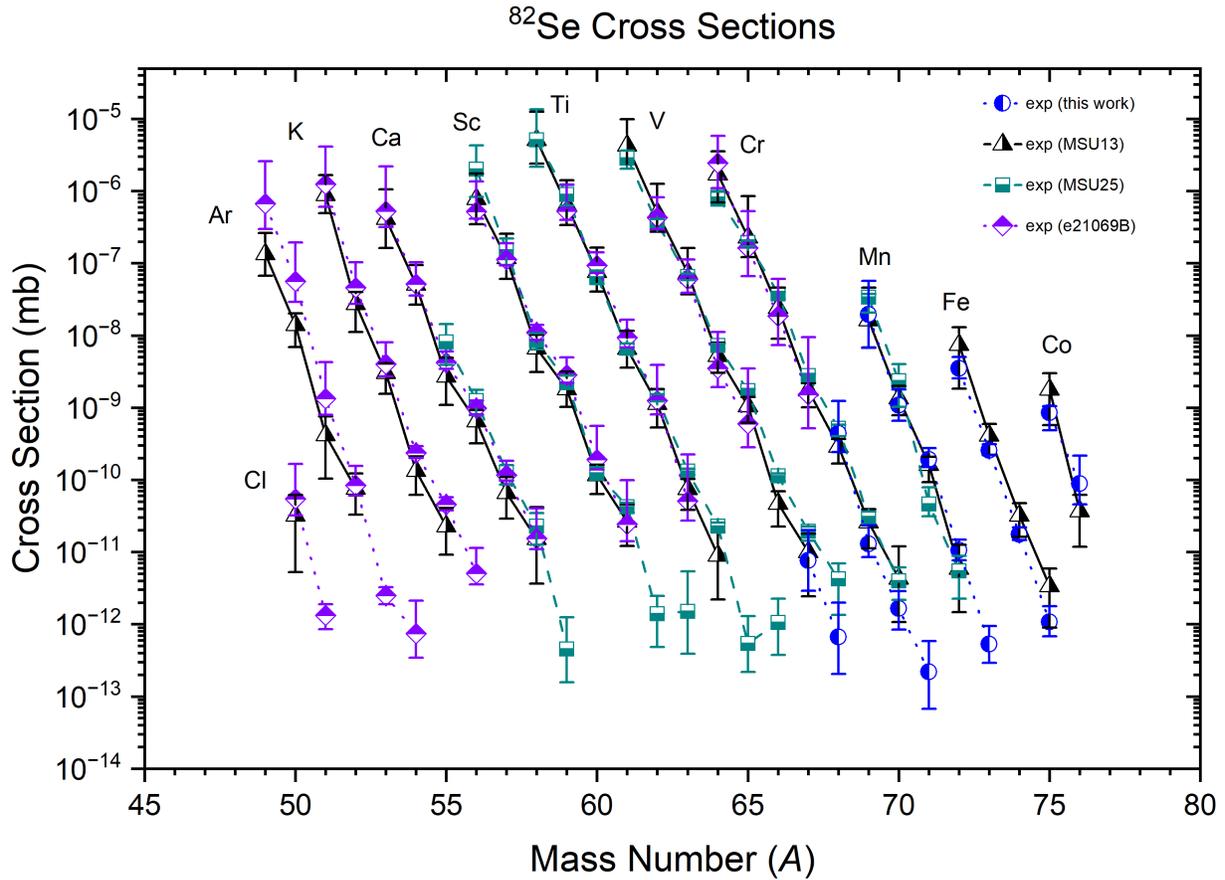


Figure 26: Cross sections from multiple experiments plotted together. From the reaction ^{82}Se on carbon. Data from Refs. [5, 12] and manuscript in preparation.

If no errors are identified, discrepancies may reflect genuine experimental differences. The most common sources include variations in primary beam energy, target material or thickness, and experimental acceptance. Such differences must be considered carefully before drawing physical conclusions from comparisons across experiments.

7.1 Partial vs Total Cross Sections

In projectile fragmentation experiments, the term cross section is used in several closely related but distinct ways. Clarifying this terminology is essential for correctly interpreting experimental results and for making meaningful comparisons to previous measurements and theoretical models.

Throughout this manual, the cross section calculated using Equation (5) represents a **partial cross section**. This quantity corresponds to the yield of a given fragment measured within a specific experimental acceptance and at a fixed magnetic rigidity. It reflects the fragment production rate sampled over a limited region of momentum space defined by the separator optics and analysis gates.

In contrast, the **total production cross section** for a fragment is, in principle, defined as the integral of the differential cross section over all emission angles and momenta. In practice, however, projectile fragmentation experiments rarely measure this quantity directly¹⁶.

For the purposes of this manual, and consistent with standard experimental practice, the total cross section is defined as the cross section integrated over the fragment's longitudinal momentum distribution. This is achieved by repeating the cross section measurement at multiple magnetic rigidities ($B\rho$), each of which samples a different portion of the fragment's momentum distribution [15]. The resulting partial cross sections are then combined – typically by fitting the momentum distribution and integrating – to obtain an estimate of the total production cross section.

Under typical intermediate- and high-energy fragmentation conditions, fragments are strongly forward-focused and their angular distributions are narrow. When the angular acceptance of the separator is sufficiently large and the transmission parameter of the partial cross section includes the momentum acceptance, integrating over the longitudinal momentum effectively captures the dominant contribution to the total yield. In this regime, the momentum-integrated cross section provides a reliable approximation to the formal total production cross section defined in Equation 3.

It is important to emphasize that a single cross section measurement, as derived in this manual, should always be interpreted as a partial cross section. Only by performing the analysis at two or more magnetic rigidity settings – and accounting for transmission and acceptance effects – can a total cross section be extracted. The procedures for momentum scans, fitting, and integration are experiment-specific and are therefore deferred to Appendix F.

This distinction is particularly important when comparing cross sections between experiments. Apparent discrepancies may arise if measurements from different locations along the momentum distribution are compared, or if differences in acceptance and transmission are not consistently treated.

7.2 Abrasion-Ablation

Abrasion–ablation is a widely used reaction model for describing intermediate- and high-energy projectile fragmentation [16]. In this framework, the reaction is separated into two conceptually distinct stages: **abrasion**, in which nucleons are rapidly removed from the projectile nucleus due to its geometric overlap with the target, and **ablation**, in which the highly excited projectile-like remnant de-excites through particle emission.

During the abrasion stage, the projectile and target pass through one another on a timescale of order 10^{-22} s. Nucleons located in the geometric overlap region are removed

¹⁶Experiments have been performed to explicitly measure the total cross section. See, for example, Refs.[13, 14]

from the projectile, while the remaining nucleons form a projectile-like pre-fragment. To first order, the number of abraded protons and neutrons is determined by the impact parameter of the collision and the nuclear density distributions of the projectile and target¹⁷. This stage is essentially a fast, non-equilibrated process, and the resulting pre-fragment retains approximately the forward velocity of the original projectile.

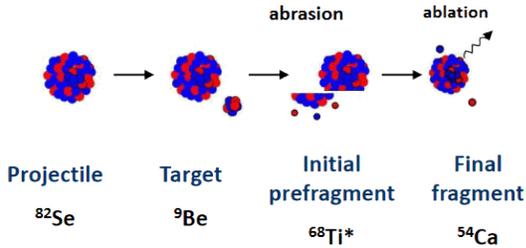


Figure 27: Schematic illustration of the abrasion–ablation process.

The sudden removal of nucleons leaves the pre-fragment in a highly excited state, with excitation energy that scales with the number of abraded nucleons. In the subsequent ablation stage, this excess energy is dissipated through the statistical emission of light particles –primarily neutrons, but also protons and light clusters at higher excitation energies. This de-excitation occurs on a longer timescale ($\sim 10^{-19}$ – 10^{-16} s) and drives the pre-fragment toward the final observed residual nucleus.

Following particle evaporation, the residual fragment may still populate long-lived excited or isomeric states. On even longer timescales, these states can decay via γ -ray emission, which can be observed experimentally using ancillary detector systems. While this electromagnetic decay does not alter the mass or charge of the fragment, it can provide valuable spectroscopic information and, in some cases, serve as an additional handle for particle identification.

A schematic overview of the full abrasion–ablation sequence is shown in Figure 27. In practice, this model provides an intuitive and effective framework for interpreting isotopic production cross sections and for connecting measured yields to underlying reaction mechanisms. Subsequent sections build on this picture to compare experimental cross sections with abrasion–ablation–based calculations and systematics.

7.2.1 Dissipation and Friction

Between the abrasion and ablation stages of the abrasion-ablation model, an additional interaction regime may contribute to the excitation and kinematic properties of the pre-fragment. Within the simplified framework described above, this intermediate stage is largely neglected. Nevertheless, it can play an important role by modifying both the excitation energy generated during abrasion and the subsequent particle emission during ablation.

During the collision of the projectile and target nuclei, the geometrical overlap region persists for a finite interaction time and may experience nuclear friction [17]. This friction is proportional to both the degree of geometrical overlap and the duration of the overlap. Microscopically, it arises from nucleon-nucleon interactions within the overlap region. These interactions dissipate kinetic energy, converting a portion of the relative motion of the projectile into internal excitation energy of the pre-fragment. As a result, the excitation energy entering the ablation stage is increased, enhancing the probability for light-particle emission and thereby modifying final fragment yields.

¹⁷Note that the impact parameter must be large. Small impact parameters usually lead to complete destruction of the projectile nucleus.

When the interaction is treated as non-instantaneous, this intermediate stage also provides a natural mechanism for nucleon pick-up processes from the target. This is seen experimentally, as in Figure 32. Such processes are difficult to reconcile within a purely sudden abrasion picture but arise naturally when finite interaction times and dissipative effects are included.

Experimental signatures of frictional dissipation extend beyond total production cross sections. In particular, the low-momentum tail of fragment longitudinal momentum distributions has been attributed, in part, to energy-loss arising from nuclear friction¹⁸. Increasing friction leads to a broadening of this tail, which in turn enhances production cross sections at momenta lower than those predicted by purely geometric abrasion models.

In LISE⁺⁺, the geometrical model of Wilson, Townsend, and Badavi [18] incorporates additional excitation energy arising from frictional dissipation. In this formulation, the total excitation energy of the pre-fragment is expressed as

$$E^* = (\gamma f \Delta S)_{\text{geom}} + E_{\text{friction}}, \quad (7)$$

where γ is the surface energy (MeV/fm²), f is a correction factor accounting for surface distortion effects, and ΔS is the excess surface area of the pre-fragment immediately following abrasion relative to a sphere of equal volume. The correction factor f depends on the number of abraded nucleons.

The additional excitation energy due to friction is parameterized as

$$E_{\text{friction}} = k_1 C_p + k_2 C_p C_t, \quad (8)$$

where C_p is the length of the longest chord at the projectile surface interface (green line, in Figure 28), C_t is the chord of intersection between the projectile and target (red line, in Figure 28), and k_1 and k_2 are phenomenological coefficients. This formulation provides an intuitive geometrical link between the collision geometry and the resulting dissipative excitation.

For typical FRIB beam energies ($\sim 150\text{-}400$ MeV/ u) and light production targets (e.g. carbon), the contribution from frictional excitation is expected to be small. The high beam velocity and limited geometrical overlap result in short interaction times and minimal dissipation. Nonetheless, the explicit inclusion of friction highlights a physical mechanism that is absent from many modern excitation-energy prescriptions. Modern rare isotope production increasingly relies on extrapolations across beam energy, target material, and reaction systematics, where small effects at high energy can evolve into non-negligible contributions under different conditions. Incorporating friction provides a more complete and internally consistent description of the collision dynamics, ensuring that excitation-energy models remain valid beyond the narrow regime in which they were tuned. Moreover, even weak dissipative effects can influence subtle observables such as the low-momentum tails of fragment distributions and the population of marginally bound species, which are particularly relevant for precision cross section predictions. Including friction therefore improves the predictive robustness of excitation-energy prescriptions in LISE⁺⁺ and strengthens confidence in their applicability across current and future experimental programs.

¹⁸https://lise.frib.msu.edu/paper/2006_09_acs_pf.pdf

7.2.2 Abrasion-Fission

As will be discussed in paragraph 7.2.4.3, there are several possible de-excitation channels available to an excited pre-fragment following the abrasion step. For sufficiently heavy primary projectiles, and for pre-fragments with large excitation energies, fission can become the dominant de-excitation mechanism. In this reaction pathway, the projectile first undergoes abrasion, producing a highly excited heavy pre-fragment, which subsequently fissions into two primary fragments. Each fission product may retain a significant amount of excitation energy and can undergo further particle evaporation before reaching its final bound state. This reaction mechanism is illustrated schematically in Figure 29.

In LISE^{++} , fission following abrasion is treated as a stochastic decay channel whose probability depends on the mass, charge, and excitation energy of the pre-fragment^{19,20}. Detailed descriptions of the fission modeling implemented in LISE^{++} , including the treatment of fragment yields and kinematics, can be found in Refs.[6, 8, 9]. While abrasion-ablation dominates the production of fragments close in mass to the projectile, abrasion-fission plays a crucial role in accessing neutron-rich nuclei where direct fragmentation cross sections may be vanishingly small.

From an experimental perspective, abrasion-fission introduces several important complications for particle identification and transmission. In the rest frame of the fissioning nucleus, the two primary fission fragments are emitted nearly isotropically, populating a spherical shell in velocity space. When boosted into the laboratory frame, this distribution becomes strongly elongated along the beam axis, forming an ellipsoidal shell in momentum space. As a consequence, a significant fraction of the fission fragments are emitted at large angles relative to the beam direction, in contrast to projectile fragmentation residues, which are strongly forward-focused. The momentum space shell reconstruction of krypton isotopes created in the fission of uranium are shown in Figure 30.

The first implication of this kinematic structure is that only a subset of fission products fall within the angular acceptance of a fragment separator. Experiments relying on abrasion-fission therefore require large transverse acceptance to efficiently transmit the desired fragments, and the measured yields are particularly sensitive to the optical settings and acceptance modeling of the separator.

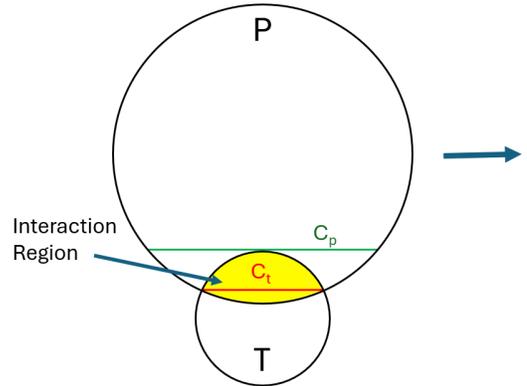


Figure 28: Geometry of a simple abrasion-ablation model at maximum overlap. The chords used for a simple friction calculation are shown.

Figure 29: Schematic illustration of the abrasion-fission process. The diagram shows a projectile nucleus (^{238}U) interacting with a target nucleus (^{238}U). The interaction results in a pre-fragment ($A_2 Z_2$) and a smaller fragment ($A_1 Z_1$). The pre-fragment then undergoes fission, producing two primary fragments ($A_3 Z_3$ and $A_4 Z_4$) and several smaller particles (A_3' and A_4').

Figure 29: Schematic illustration of the abrasion-fission process.

¹⁹https://groups.frib.msu.edu/nscl_library/nscl_preprint/MSUCL1299.pdf

²⁰https://lise.frib.msu.edu/7_5/lise++_7_5.pdf

A second complication arises from the broad longitudinal momentum distribution of the fission fragments. Fragments emitted forward and backward in the fissioning nucleus rest frame can possess identical proton and neutron numbers, but markedly different laboratory-frame velocities. Because magnetic rigidity and time-of-flight measurements are both velocity-dependent, fragments with the same Z and A can populate different regions of PID space. In extreme cases, backward-emitted fragments may be misidentified or fall outside the nominal PID gates defined for forward-emitted residues. Careful trajectory reconstruction, charge-state separation, and acceptance corrections are therefore essential when interpreting abrasion-fission data.

In spite of these experimental challenges, abrasion-fission remains a powerful production mechanism for rare isotopes. When properly modeled and corrected for acceptance effects, fission-based production enables access to regions of the nuclear chart that are otherwise unreachable with standard abrasion-ablation reactions, and it plays a central role in many current and planned rare isotope experiments at FRIB.

7.2.3 Abrasion–Ablation Calculations in LISE⁺⁺

The abrasion-ablation (AA) calculations implemented in LISE⁺⁺ apply the reaction model described above to predict fragment production cross sections for a given projectile-target system²¹. In practice, the calculation folds together several physical ingredients that govern (i) how a pre-fragment is formed during abrasion and (ii) how that pre-fragment subsequently de-excites through particle evaporation. Many of these ingredients introduce parameters that can be adjusted during fitting to better reproduce experimental data.

The most important model components and user-controlled inputs are summarized below, along with their physical motivation and typical role in the fitting process.

Nuclear Mass Model

The nuclear mass model provides tabulated values of nuclear masses and derived quantities such as binding energies, neutron and proton separation energies, and Q -values for particle emission. These quantities enter directly into the ablation stage, where they determine which evaporation channels are energetically allowed and how probable they are.

In LISE⁺⁺, experimentally measured masses are taken from the 2020 Atomic Mass Evaluation (AME2020) wherever available [19]. For nuclei beyond the reach of experimental

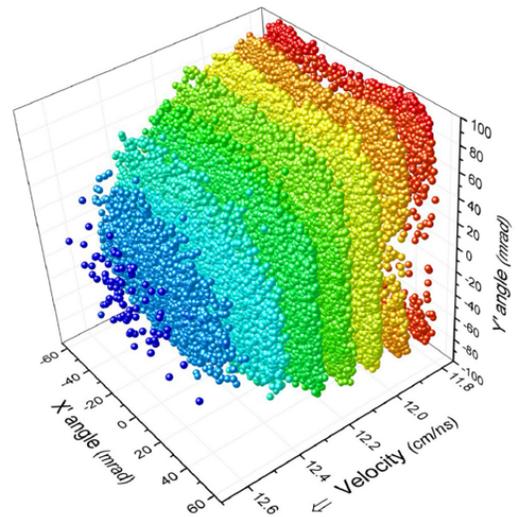


Figure 30: Reconstruction of momentum space as a function of transverse angle of krypton fragments from the fission of uranium. From [6].

²¹See <https://lise.frib.msu.edu/AA> for more details.

mass measurements, theoretical mass models are used to supply the missing values. Among these, the Hartree-Fock-Bogoliubov HFB22 mass model [20] is the most commonly employed due to its global coverage and reasonable performance across much of the nuclear chart.

Because separation energies strongly influence evaporation probabilities, the choice of mass model can have a noticeable effect on predicted cross sections, particularly for very neutron-rich or proton-rich fragments. This sensitivity underlies the connection between cross sections and nuclear binding systematics discussed in the ΔBE analysis.

Particle Evaporation Channels

The particle evaporation settings define which light particles are permitted to be emitted during the ablation stage. By default, LISE⁺⁺ allows evaporation of neutrons, protons, and α particles, which are the dominant de-excitation modes for most projectile-like pre-fragments at intermediate energies.

If the excitation energy of the pre-fragment is sufficiently high, additional channels may become relevant. For heavy systems, fission can compete with particle evaporation, while at very high excitation energies multifragmentation or breakup processes may occur. Enabling or disabling specific evaporation channels effectively constrains how the pre-fragment sheds energy and mass, and therefore influences the final isotopic distributions.

Excitation Energy of the Pre-Fragment

The excitation energy imparted to the pre-fragment during abrasion is one of the most influential inputs to AA calculations. Physically, this excitation energy arises from the sudden removal of nucleons, which disrupts the mean field and leaves the remaining nucleus far from equilibrium. To first order, the total excitation energy scales with the number of abraded nucleons.

Within LISE⁺⁺, several parameterizations are available to describe this excitation energy. The simplest is the Gaimard–Schmidt model, which assigns an average excitation energy per abraded nucleon [21]. More sophisticated treatments introduce statistical distributions of excitation energy, allowing event-by-event fluctuations around a mean value.

The most commonly used option in practical analyses is the exponential mean- T model [22], which characterizes the excitation energy distribution in terms of an effective temperature. Increasing the excitation energy generally increases the probability of emitting multiple neutrons, thereby shifting the predicted fragment yields toward lower mass numbers along a given isotopic chain.

Effective Coulomb Barrier

The effective Coulomb barrier parameter modifies the probability for charged-particle evaporation, particularly proton and α emission. Physically, it reflects the electrostatic repulsion experienced by charged particles attempting to escape the nuclear potential well.

In LISE⁺⁺, this parameter effectively sets a charge radius for the pre-fragment and therefore controls the height of the Coulomb barrier. Adjusting this parameter primarily impacts proton-rich fragments and can be important when reproducing yields on the proton-deficient side of the projectile.

Odd-Even Effect

Experimental cross sections often exhibit odd-even staggering, especially near the drip lines, where nuclei with paired protons and neutrons are systematically more bound than their odd- Z or odd- N neighbors. This behavior originates from nucleon pairing correlations.

The odd-even parameter in `LISE++` introduces a phenomenological correction to reproduce this staggering in calculated cross sections. While not derived from first principles, it provides a practical means of capturing pairing effects that are otherwise smeared out in purely statistical evaporation models.

Breakup Limiting Temperatures

At sufficiently high excitation energies, a pre-fragment may no longer be well described as a compound nucleus undergoing sequential evaporation. Instead, it may undergo complete breakup into its constituent nucleons or light clusters. To prevent unphysical evaporation chains, `LISE++` introduces limiting temperatures that cap the excitation energy available for standard ablation processes.

These breakup limiting temperatures act as free parameters during AA fits and effectively regulate the transition between evaporation-dominated de-excitation and more violent disassembly modes. While somewhat phenomenological, they play an important role in stabilizing calculations at high excitation energies and improving agreement with experimental yields.

7.2.4 Performing Abrasion-Ablation Fits

The next step is to perform a fit that optimizes the model parameters to reproduce the measured cross sections. Because the fitting procedure can modify many interconnected quantities, it is strongly recommended to perform these calculations in a fresh `LISE++` file to avoid unintended carryover from previous analyses.

Begin by selecting the nuclear mass model (`HFB22` is recommended unless a specific alternative is justified), importing the experimental cross sections, and choosing the excitation energy model (typically `mean-T`). When importing data, it is advantageous to combine newly measured cross sections with reliable literature values so that each isotopic chain spans as wide a range in neutron number as possible. Broad coverage helps constrain the excitation and evaporation parameters and reduces sensitivity to local statistical fluctuations.

To access the fitting interface, navigate to:

`Utilities` → `Reaction Utilities` → `User Cross Section Analysis` →
`Minimization of E^* Parameters.`

Upon opening this dialog, the `Fit` button may initially be disabled. This indicates that prerequisite internal quantities have not yet been initialized. These are generated automatically by selecting the `Make items 1-3` button located near the center of the dialog (Figure 31, **A**). Once these steps are completed, the fitting controls become available.

Underneath the `Make items 1-3` button is a `Pre-fragment Excitation Energy` button. This opens a dialog that allows you to define the excitation energy model.

The screenshot shows the LISE++ fitting dialog with several key regions labeled: **D** (AA minimization Settings), **C** (Local line to analyze), **B** (Weights), and **A** (Operations). The dialog includes various input fields, checkboxes, and buttons for configuring the fitting process.

Figure 31: The abrasion-ablation fitting dialog in LISE⁺⁺. Key regions discussed in the text are labeled.

A central aspect of the minimization procedure is the treatment of experimental uncertainties and the relative weighting of different isotopes (Figure 31, **B**). LISE⁺⁺ provides four objective-function components: global χ^2 , local χ^2 , and the corresponding logarithmic differences. The logarithmic terms are particularly important for fragmentation data, where cross sections often span many orders of magnitude. Without them, a fit driven purely by χ^2 would be dominated by the most abundant fragments and could fail to reproduce the rare, exotic isotopes that are often of greatest physical interest.

The *local line* selection, controlled by the button near Figure 31, **C**, determines which isotopic chain (or subset of chains) is emphasized during a given stage of the fit. LISE⁺⁺ also allows the local line to be cycled automatically: after a fixed number of iterations, the fit shifts its focus to neighboring isotopic chains (e.g., $16 \rightarrow 17 \rightarrow 15 \rightarrow 16$). This cycling helps prevent the minimization from becoming trapped in a local minimum dominated by a single chain and encourages a more globally consistent solution. Utilizing this feature is highly recommended. For small- and medium-sized data sets, a scanning range of ± 1 is sufficient; for larger sets, ± 2 or larger is recommended.

The parameters eligible for fitting are listed at the top of the dialog (Figure 31, **D**). In the example shown, the mean- T excitation energy model is selected, with its associated parameters exposed for adjustment. Although many parameters can be floated simultaneously, doing so is rarely optimal.

A practical and robust fitting strategy is to proceed incrementally:

- Begin by fitting only the base excitation parameter (k_1), which sets the overall scale

of the excitation energy.

- Introduce higher-order terms (k_2 and k_3) sequentially to refine the curvature and width of isotopic distributions.
- Add the odd–even parameter to reproduce pairing-related staggering once the gross trends are established.
- Include the effective Coulomb barrier parameter to fine-tune proton-rich yields.

The breakup limiting temperatures may also be floated, but in most cases the default values provided by LISE⁺⁺ are robust. Large deviations from these defaults often signal that the fit has settled into an unphysical local minimum rather than uncovering genuine physics.

The final parameter to adjust should be the overall amplitude factor. This multiplicative scale shifts all calculated cross sections uniformly and is best optimized only after the relative shapes of the isotopic distributions are well reproduced. Adjusting the amplitude too early can mask deficiencies in the underlying model parameters and lead to misleading agreement.

7.2.4.1 Physical Bounds for Fitting Parameters

Early in the abrasion-ablation fitting process, the numerical values of the fit parameters can appear abstract, and it is not always obvious what constitutes a *physically reasonable* solution. This subsection provides physical intuition and empirical guidance for interpreting the fitted parameters. The ranges quoted below are not strict limits, but rather typical values observed across many successful fits. Parameters that converge well outside these ranges should be treated as warning signs and examined critically.

- k_1 (**Base excitation energy per abraded nucleon**)

In the mean- T excitation energy model, k_1 sets the average excitation energy deposited per abraded nucleon. Physically, it reflects the typical energy cost of removing a nucleon from the projectile during the abrasion stage and embedding it into an excited pre-fragment.

This parameter depends on the mass and energy of the primary beam, as heavier projectiles generally generate more excitation per removed nucleon due to increased participant–spectator interactions. Across a wide range of intermediate-energy fragmentation experiments, acceptable values typically lie in the range

$$15 \lesssim k_1 \lesssim 20 \text{ MeV}$$

Values significantly below this range often underpredict particle evaporation, while values well above it tend to overbroaden isotopic distributions.

- k_2 (**Quadratic excitation term**)

The k_2 parameter introduces a quadratic dependence of excitation energy on the number of abraded nucleons (ΔA^2). Physically, this term accounts for the fact that large

overlaps between projectile and target produce disproportionately higher excitation energies than would be expected from a purely linear scaling.

Empirically, k_2 is typically small, reflecting a modest correction rather than a dominant contribution. Reasonable values usually fall in the range

$$0 \lesssim k_2 \lesssim 1$$

Values outside this range often indicate that the fit is compensating for deficiencies elsewhere in the model.

- **k_3 (Isospin-dependent excitation term)**

The k_3 parameter encodes an isospin dependence in the excitation energy, proportional to $(\Delta N - \Delta Z)$. This term reflects the differing nuclear interactions experienced by neutrons and protons during abrasion, including symmetry-energy effects.

Fits across the nuclear chart show that k_3 is generally negative, indicating slightly reduced excitation for neutron-rich abrasion. Typical values lie in the range

$$-1 \lesssim k_3 \lesssim 0$$

Positive values are rarely physical and usually signal an unstable or poorly constrained fit.

- **Effective Coulomb Barrier**

The effective Coulomb barrier parameter controls the suppression of charged-particle evaporation (primarily protons and α particles) during the ablation stage. Physically, it represents an effective charge radius and barrier height experienced by evaporating particles.

This parameter decreases with increasing projectile mass, as larger systems have more diffuse charge distributions and lower effective barriers. Empirical fits suggest typical values in the range

$$\text{Barrier} \approx 5 - 7$$

with heavier beams favoring the lower end of this range.

- **Odd-Even Parameter**

The odd-even parameter models the staggering observed in isotopic yields near the driplines, arising from nuclear pairing correlations. Its effect becomes most pronounced for nuclei where separation energies change rapidly with neutron or proton number.

This parameter generally scales with projectile mass, reflecting stronger pairing effects in heavier systems. Reasonable values typically fall in the range

$$\text{Odd-Even} \approx 6 - 12$$

Unusually large values may indicate that the fit is attempting to reproduce fine structure that is not well supported by the data.

- **Limiting Temperatures**

The limiting temperatures define the transition from sequential evaporation (ablation) to breakup or multifragmentation. Physically, they cap the temperature of the pre-fragment to prevent unphysical excitation energies.

The default limiting temperatures provided by LISE⁺⁺ have been shown to be robust across much of the nuclear chart. Large deviations from these defaults are strongly discouraged, as they often indicate convergence to a nonphysical local minimum rather than genuine sensitivity to the data.

- **Amplitude Factor**

The amplitude factor uniformly rescales all calculated cross sections. Ideally, this parameter should remain close to unity, indicating that the absolute normalization of the model is consistent with the experimental inputs.

In practice, values in the range

$$\text{Amplitude} \approx 0.8 - 1.2$$

are considered acceptable. Larger deviations may arise, but should be treated skeptically.

Taken together, these parameter ranges provide a practical framework for assessing the physical plausibility of abrasion-ablation fits. A fit that reproduces the data well but yields parameters far outside these ranges should be treated with skepticism and revisited before being used for interpretation or extrapolation.

7.2.4.2 Diagnosing and Correcting Unphysical Fits

It is not uncommon for abrasion-ablation minimization to converge to a statistically acceptable solution (e.g., low global χ^2) while producing one or more parameters that lie outside physically reasonable bounds. Such outcomes typically indicate that the minimizer has compensated for deficiencies in the dataset or model assumptions rather than capturing the underlying reaction physics. The following diagnostic checks provide guidance for identifying and correcting these situations.

An important visual aid is the χ^2 table that LISE⁺⁺ generates at the conclusion of a fit. To see the χ^2 and logarithmic differences plot, click the **Generate chi2-table** in the **Fitting** box at the center of the dialog.

- **Unphysical excitation-energy parameters:** Extreme values of k_1 , k_2 , or k_3 (e.g., $k_1 \ll 15$, $k_2 > 1$, or $k_3 > 0$) often indicate insufficient isotopic coverage or strong parameter correlations. This most commonly arises when only mid-chain isotopes are included or when exotic tails are underweighted.
 - Fix higher-order parameters (k_2 , k_3) at reasonable values and refit k_1 alone.
 - Expand or contract the dataset by including/excluding previous measurements or neighboring elements, even if uncertainties are larger.

- Verify that the chosen excitation-energy model is appropriate for the projectile mass and beam energy.
- **Instability under local-line cycling:** Strong dependence of fitted parameters on the chosen local line or cycling pattern indicates an underconstrained fit or the presence of multiple local minima.
 - Reduce the number of free parameters.
 - Improve isotopic coverage before refitting.
 - Favor solutions that are stable across multiple cycling configurations.
- **Misleadingly good global χ^2 :** A low χ^2 can mask systematic trends. Always inspect residuals along individual isotopic chains.
 - Look for systematic over- or under-prediction of neutron-rich or proton-rich tails.
 - Check for smooth reproduction of trends across neighboring elements.

In general, physically meaningful solutions are characterized by parameter values within established bounds, smooth reproduction of isotopic trends, and stability under reasonable variations of fitting strategy. When ambiguities persist, conservative exclusion of poorly constrained isotopes is preferable to propagating unphysical parameters into interpretation or extrapolation.

7.2.4.3 When to Enable Additional Evaporation Channels

The user-controlled evaporation settings in LISE⁺⁺ determine whether additional composite light-particle channels – specifically $2n$, $2p$, deuteron (d), and triton (t) emission – are permitted during the ablation stage. The decision to enable these channels should be guided by physical considerations, not solely by fit quality.

Under typical intermediate-energy projectile fragmentation conditions, de-excitation of the pre-fragment is dominated by sequential neutron evaporation, with proton and α emission contributing as excitation energy increases. Composite-particle channels such as $2n$, $2p$, d , and t represent higher-order decay modes that become relevant only when the excitation energy is sufficiently large to overcome their higher effective emission thresholds.

The need to study or enable these additional channels generally arises under the following conditions:

- The fitted excitation-energy parameters (e.g., k_1) imply excitation energies well above typical neutron-evaporation scales.
- Systematic discrepancies appear between measured and calculated cross sections for specific regions, particularly proton-rich isotopes or light fragments where cluster emission may compete with sequential evaporation.
- The fit reproduces neutron-rich isotopes well but consistently overpredicts yields closer to stability, suggesting that additional decay pathways may be draining population from those channels.

- The reaction involves light or medium-mass projectiles, where composite-particle emission is known to be more competitive at high excitation energies.

When these conditions are met, enabling composite evaporation channels should be treated as a controlled physics study rather than a default fitting step. Channels should be enabled incrementally, and their impact on both global fit quality and isotopic-chain behavior should be examined carefully. In particular, one should assess whether enabling a new channel leads to physically reasonable changes – such as selective suppression of specific isotopes – rather than a uniform rescaling of all cross sections.

It is important to recognize that opening additional evaporation channels can significantly alter the sensitivity of the fit to excitation-energy parameters. Apparent improvements in χ^2 may reflect compensating effects between channel competition and excitation energy, rather than a more accurate physical description. For this reason, enabling composite channels often necessitates revisiting previously fitted parameters to ensure they remain within physically meaningful bounds.

As a best practice, abrasion–ablation fits should first be stabilized using the minimal evaporation channel set (typically n , p , and α only). Composite channels ($2n$, $2p$, d , t) should then be enabled selectively and only when motivated by clear physical or systematic evidence. This approach preserves interpretability of the fit and avoids masking deficiencies in excitation-energy modeling or mass input with excessive decay freedom.

7.2.4.4 Determining a Global Minimum

The abrasion-ablation minimization procedure described above identifies a best-fit solution for a *single* choice of nuclear mass model and excitation-energy prescription. However, because the AA model is inherently multi-parameter and highly non-linear, a single minimization run does not guarantee convergence to the true global minimum.

Different nuclear mass models and excitation-energy parameterizations can lead to distinct local minima with comparable χ^2 values but different physical implications. To ensure that the extracted parameters and predictions are robust, it is therefore recommended to repeat the full minimization procedure using multiple mass models and, where appropriate, alternative excitation-energy models.

In practice, past analyses have found that the combination of the HFB22 mass model with the mean- T excitation-energy prescription provides the most consistent and physically reasonable description across large regions of the nuclear chart. Nevertheless, this outcome should not be assumed *a priori*. Different projectile masses, beam energies, or regions far from stability may favor alternative models.

The final abrasion-ablation prediction should be selected based on a combination of statistical quality (e.g., χ^2), physical plausibility of fitted parameters (see paragraph 7.2.4.1), and stability of the solution across repeated minimizations. Agreement among multiple fits provides strong evidence that a true global minimum has been reached rather than an artifact of initial conditions or model assumptions.

After a suitable parameter set has been selected, the predicted cross sections should be plotted directly against the measured values, as shown in Figure 32. This comparison provides a visual assessment of the quality of the fit across the full dynamic range of the data.

It is also useful to overlay predictions from the empirical parameterization EPAX3, a global fragmentation systematics model derived from a large body of experimental data [23]. EPAX3 is not expected to reproduce detailed isotopic structure or extreme neutron- or proton-rich yields, but it serves as a valuable baseline reference. Systematic deviations from EPAX3 – particularly those that are smooth and localized along isotopic chains – often reflect genuine nuclear structure effects or limitations of empirical systematics, whereas erratic or isolated discrepancies may indicate issues in the experimental analysis or the fitting procedure.

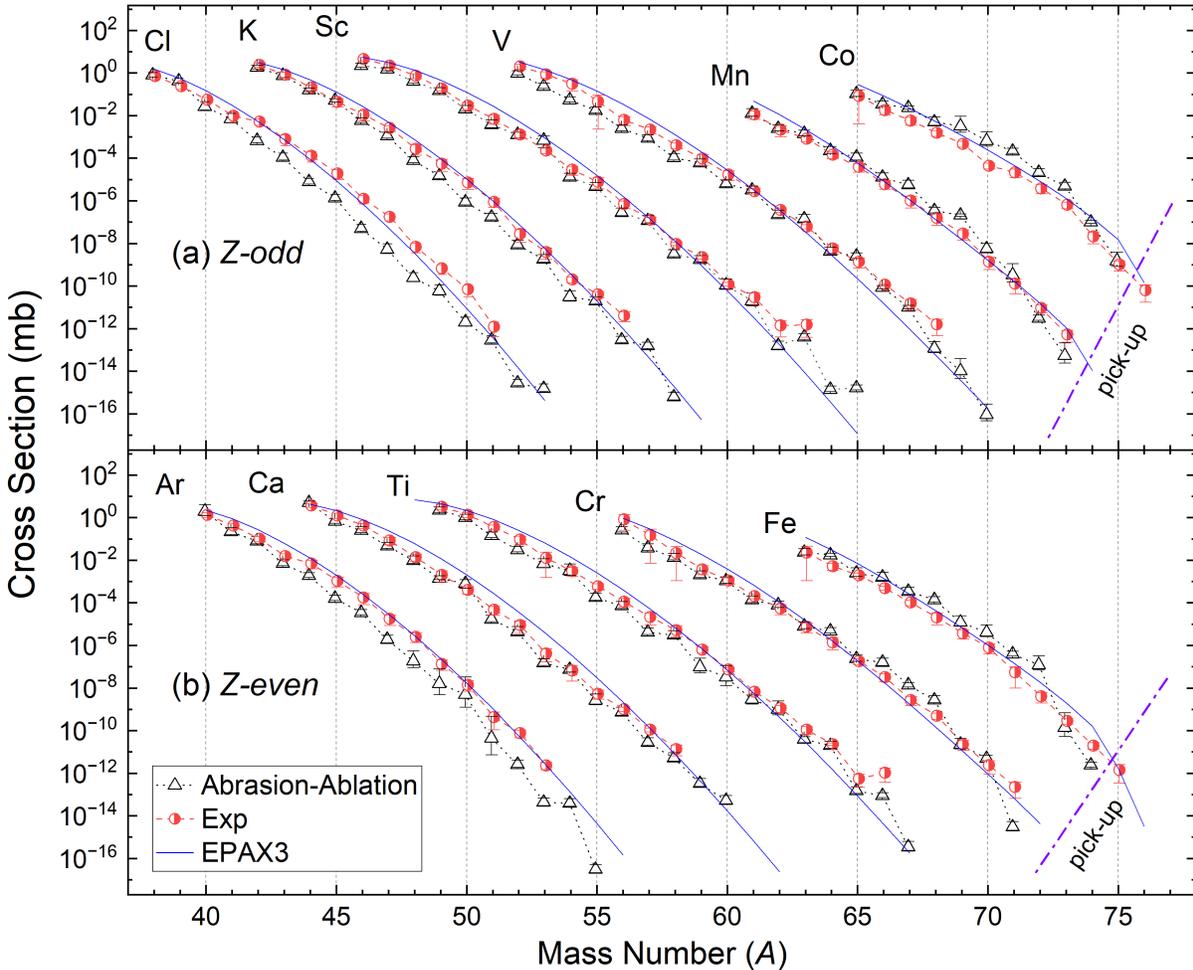


Figure 32: Abrasion-ablation and EPAX3 predicted cross sections vs experimentally measured cross sections. Pick-up reactions are labeled.

7.3 ΔBE Systematics

Cross sections for projectile fragmentation products exhibit a strong and remarkably smooth correlation with nuclear binding energy. This observation motivates the use of empirical ΔBE systematics, which relate measured production cross sections to differences in binding energy along isotopic chains. A commonly used parameterization is

$$\log(\sigma) = k_1 \ln(dBE + 1) + k_2 \ln\left(\frac{e^{k_0 S_{\min}} - 1}{k_0}\right) + b + dZ, \quad (9)$$

where

- S_{\min} is the minimum nucleon separation energy for the isotopic chain, typically taken as the separation energy of the last bound isotope,
- $dBE = BE_{\max}(Z) - BE(Z, N)$ is the binding-energy difference between the most bound isotope of a given element and the isotope of interest,
- $dZ = t(Z_{\text{beam}} - Z_{\text{frag}})$ accounts for systematic shifts between isotopic chains of different atomic number,
- b is an overall normalization constant, and
- $k_0, k_1, k_2,$ and t are empirical fit parameters.

The quantity dBE serves as a compact measure of how far a fragment lies from the most strongly bound configuration of a given element. As dBE increases toward the neutron or proton dripline, production cross sections decrease rapidly and approximately exponentially. Expressing the cross section in logarithmic form reveals a near-linear dependence on dBE , enabling systematic trends to be extracted even across many orders of magnitude in yield.

The term involving S_{\min} accounts for the finite termination of isotopic chains at the dripline and encodes sensitivity to separation energies, which govern both nuclear stability and evaporation pathways during the ablation stage of the reaction. The dZ term allows for a smooth interpolation between neighboring elements by compensating for changes in projectile–fragment mass asymmetry.

Although Equation (9) is purely empirical, its success reflects the underlying role of nuclear binding in fragment production: nuclei with lower binding energy are both harder to populate during abrasion and more likely to decay during ablation. Studying deviations from the smooth ΔBE trends can therefore reveal nuclear structure effects such as shell closures, odd–even staggering, or changes in separation energy behavior near the dripline.

A full derivation and physical motivation of these systematics are beyond the scope of this manual. Interested readers are referred to the foundational work by Tarasov *et al.* [5], which established and validated this formalism across a wide range of projectile fragmentation data.

7.3.1 Performing ΔBE Fits

ΔBE systematics fits are performed in LISE⁺⁺ via the dialog

```
1D-plot → Systematic Distributions → dBE(N, max) → dBE-systematics vs CS
→ dBE-systematics Settings.
```

The overall fitting philosophy closely mirrors that used for abrasion-ablation analyses. Fits should be built incrementally: begin with the minimum number of free parameters

and introduce additional parameters only as required to achieve stable convergence and a physically reasonable trend.

Experience has shown that the parameter k_0 is consistent with zero for essentially all ΔBE fits performed to date. Physically, this implies that the exponential cutoff term associated with the minimum separation energy does not introduce additional curvature beyond what is already captured by the logarithmic dependence on ΔBE . As a result, k_0 should be fixed to zero unless compelling evidence suggests otherwise.

Similarly, the parameter t , which introduces a shift between isotopic chains of different atomic number, is typically set to zero. This reflects the fact that ΔBE systematics are most robust when applied along individual isotopic chains rather than across large Z ranges. Allowing t to float can artificially absorb deficiencies in the mass model rather than reveal them.

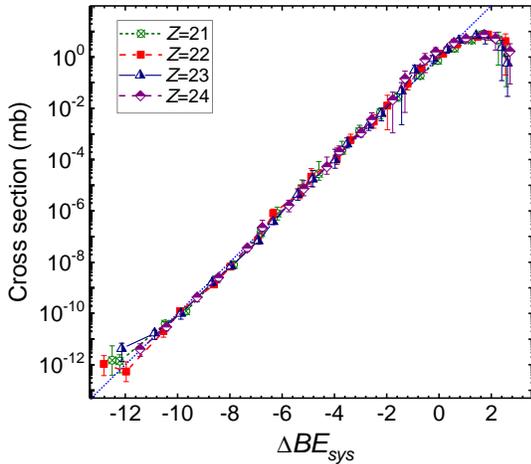


Figure 33: Measured cross sections for isotopes produced in the fragmentation of ^{82}Se plotted as a function of ΔBE using the HFB-22 mass table. The blue dotted line represents the fitted functional form $\sigma = 10^{\Delta BE_{\text{sys}}}$. From Ref. [5].

extracted trends. If a given mass model fails to reproduce a smooth or physically reasonable ΔBE dependence, alternative mass tables should be tested. Figure 33 shows an example of a high-quality ΔBE fit obtained using the HFB-22 mass model.

7.3.2 Mass Table Adjustments

The ΔBE formalism relies fundamentally on nuclear binding energies, which are well known near stability but increasingly uncertain far from it. In regions where experimental mass measurements are unavailable, theoretical mass models must be used, and different models can predict substantially different binding energies for the same exotic nucleus. This sensitivity underlies why changing the mass model can significantly alter both abrasion-ablation and ΔBE fit results.

Within the minimization dialog, an upper cross section cutoff may be specified. This option exists to avoid generating multiple reduced data files and reflects an important physical limitation of the ΔBE approach: the systematics are most applicable in the very neutron-rich (or proton-rich) regime, where binding-energy differences dominate the production probability.

Including high-cross-section isotopes near stability can degrade the fit quality, as these nuclei are influenced by reaction mechanisms and structural effects not captured by the ΔBE parameterization. Reducing the cutoff often improves convergence and fit quality; however, care should be taken not to remove too many data points, as this can artificially constrain the slope and inflate parameter uncertainties.

The choice of nuclear mass model plays a critical role in ΔBE analyses. Because dBE enters explicitly into Equation (9), any deficiencies or biases in the mass table directly affect the ex-

In this sense, the mass table itself acts as an implicit, high-dimensional parameter set within the ΔBE framework. The fit parameters described in the previous section primarily control the global trend of the cross section dependence, while the placement of individual data points along that trend is governed almost entirely by the calculated values of dBE , which are derived directly from the mass table.

Adjusting mass excess values for neutron-rich nuclei can therefore improve agreement between measured cross sections and ΔBE systematics. Importantly, such adjustments are not merely technical: they carry physical meaning, as changes in binding energy directly correspond to changes in nucleon separation energies and, by extension, nuclear stability.

At present, a formalized fitting procedure for mass excesses using ΔBE systematics is still under development. No dedicated Excel workbook or automated LISE⁺⁺ tool currently exists to perform this task in a systematic and reproducible way, and this remains an area of active research. Current approaches entail hand-tuning mass excesses on a case-by-case basis.

A natural extension of the analysis chain is to take a mass table modified through ΔBE considerations and reapply it to abrasion-ablation calculations. In principle, a mass model constrained by cross section systematics should lead to improved agreement in AA predictions as well. This feedback loop between ΔBE and AA analyses has not yet been fully explored and represents a promising direction for future work.

7.4 Complementarity of Abrasion–Ablation and ΔBE Analyses

Abrasion-ablation (AA) and ΔBE systematics provide two complementary lenses through which newly measured production cross sections may be interpreted. While both approaches seek to describe the same experimental observables, they emphasize different physical aspects of the reaction and carry distinct sensitivities and limitations. Used together, they form a powerful, self-consistent framework for diagnosing both reaction dynamics and nuclear structure effects.

The abrasion-ablation model is fundamentally reaction-driven. It explicitly models the geometry of the projectile-target interaction, the excitation energy imparted to the pre-fragment, and the subsequent statistical decay through particle evaporation or breakup. As a result, AA calculations are sensitive to assumptions about excitation energy generation, evaporation channels, Coulomb barriers, and odd-even effects. When tuned successfully, AA provides a global description of cross sections across wide regions of the nuclear chart and offers insight into reaction mechanisms and transport dynamics.

In contrast, ΔBE systematics are intentionally minimalist. They abstract away the details of the reaction mechanism and instead focus on a strong empirical correlation between production cross sections and nuclear binding energies. In this framework, deviations from smooth ΔBE trends are interpreted primarily as manifestations of nuclear structure – such as shell closures, changes in separation energies, or deficiencies in mass models – rather than as shortcomings of reaction modeling. As such, ΔBE analyses are particularly sensitive to the underlying mass surface and are most informative in regions far from stability.

Because of these differing sensitivities, discrepancies between AA predictions and experimental data may have multiple origins. A poor AA fit may arise from an inadequate excitation energy model, an inappropriate evaporation prescription, or an incorrect mass

input. The ΔBE analysis helps disentangle these possibilities: if the same data follow a clean ΔBE trend, the issue likely lies in the reaction modeling rather than the mass surface. Conversely, if both AA and ΔBE fail to reproduce the data using the same mass table, this strongly suggests deficiencies in the assumed binding energies or separation energies of the fragments.

This complementary behavior enables a productive feedback loop. ΔBE analyses can guide adjustments to mass tables by highlighting nuclei whose binding energies are inconsistent with observed cross sections. These adjusted masses can then be fed back into AA calculations, potentially improving their predictive power. Similarly, AA calculations can provide a global consistency check on whether mass adjustments inferred from ΔBE remain compatible with known reaction systematics.

In practice, the most robust interpretation of newly measured cross sections is achieved when both analyses converge on a consistent description. Agreement between AA predictions, ΔBE trends, and experimental data provides strong confidence that the essential reaction dynamics and nuclear structure inputs are being captured correctly. Persistent disagreements, on the other hand, are often the most scientifically valuable outcome, as they point directly toward regions where existing models or mass predictions require refinement.

8 Concluding Remarks

This manual has been written with the goal of guiding a first-time researcher through the complete process of rare isotope production cross section measurements, from raw experimental data to finalized, publishable results.

While the guide is designed to be as self-contained as possible, it is not truly stand-alone. Background knowledge is assumed in several areas related to rare isotope production, detector systems, and nuclear reaction modeling. Where appropriate, references to external sources have been provided to direct the reader to more detailed discussions.

In an effort to address the full scope of the analysis workflow, additional documentation has been appended to this manual. These appendices include practical discussions, expanded technical details, and a frequently asked questions section intended to address common points of confusion that arise during cross section analyses.

The focus of this manual has been on the how of cross section extraction, rather than the why. Nevertheless, production cross sections are valuable in two closely related contexts. First, they are essential for experimental planning in exotic regions of the nuclear chart. Production cross sections represent the ultimate limiting factor of an experiment: even with ideal transmission and high primary beam intensities, certain nuclei may only be produced on timescales of days or weeks. Quantifying production rates therefore informs the low-energy nuclear science community which experiments are currently feasible, and which may become feasible as next-generation accelerator facilities come online and reach full performance.

In this context, and particularly at FRIB, the extrapolation of cross sections using theoretical models such as abrasion–ablation and ΔBE systematics allows predictions of the facility’s ultimate reach. One of FRIB’s central scientific goals is the exploration of the neutron dripline, which is currently established experimentally only up to neon [24]. Rare isotope discovery and dripline determination are among the most compelling scientific objectives of

the Rare Isotope Research Group, and both are enabled by reliable cross section measurements. These measurements are often obtained parasitically, alongside a wide range of other cutting-edge experiments conducted at FRIB, underscoring their broad impact across the experimental program [25, 26].

Reliable cross section measurements therefore serve not only as bookkeeping quantities, but as benchmarks for reaction models and as guideposts for future facility development. Discrepancies between measured and predicted values highlight where existing excitation energy treatments, transmission modeling, or structural inputs require refinement. In this way, cross section work feeds directly back into improvements in theoretical descriptions and simulation tools such as LISE⁺⁺.

It is the author’s hope that this manual lowers the barrier to entry for new researchers while preserving the precision required for meaningful results. As facilities continue to push further from stability and into unexplored regions of the nuclear landscape, the role of careful, transparent cross section measurements will only grow in importance.

Appendix A Installing and Booting SpecTcl and SpecTk

This appendix provides a self-contained guide for installing and launching the SpecTcl and SpecTk software packages required for cross section analysis. It assumes a Windows desktop with access to the laboratory servers hosting experimental data.

1. Required Software

Two software packages must be installed on the local machine:

1. **MobaXterm**²²: Provides remote terminal access to Unix-based laboratory servers.
2. **Xming**²³: Enables the display of graphical Unix applications (X11) on the local desktop.

After installation, launch Xming through the search bar. It must be running prior to starting SpecTcl, as SpecTcl relies on X11 forwarding for all graphical output. The system tray icon confirms Xming is active.

2. Configuring Remote Access

Remote connections are established using PuTTY. It is recommended to create two saved sessions:

1. Personal account on the `fishtank` server:
`<yourFRIBUserName>@fishtank`
2. Laboratory account `arisdaq` on `fishtank`:
`arisdaq@fishtank`

²²The program can be downloaded from: <https://mobaxterm.mobatek.net/>

²³This program is available from the “Software Center” on your desktop.

For both sessions, configure:

- Connection type: SSH
- Port: 22
- X11 forwarding: enabled (under SSH → X11)

Assign descriptive names to each session and click **Save**.

3. Logging In

Open MobaXterm. The sessions defined in PuTTY appear under **User sessions**. Double-click the desired session to connect. After the laboratory banner, log in with the appropriate password: personal account passwords match your email password; the **arisdq** password is distributed separately.

Once logged in, the default directory is `/user/arisdq`. Navigate to the **SpecTcl** directory:

```
/user/arisdq/PID/User.
```

4. Preparing a Working Directory

Create a dedicated working directory for the analysis. At the start of each analysis:

1. Copy the most recent **SpecTcl** build from `/PID/Develop` into your working directory.
2. Clean and rebuild the program using the provided makefile in the **source** subdirectory.

Maintaining a single **SpecTcl** build per analysis greatly reduces potential errors and ensures reproducibility.

5. Starting the Buster Shell

Before launching **SpecTcl**, the Buster shell must be active. A script to start it is located at `/user/arisdq/scripts`. Execute the script using:

```
./busterdaq
```

Return to your working directory after the shell is running.

6. Launching SpecTcl

From the root directory of the **SpecTcl** build:

1. Open two additional terminals using `xterm &`.
2. Launch **SpecTcl** from one of these terminals.

Upon successful startup, the following windows appear:

- `info`: displays detector efficiency
- `treegui`: used for defining spectra and managing gates
- Two `SpecTcl` windows: one command-line interface, one file control GUI
- `Xamine`: not used in this workflow

If `SpecTcl` fails to load, MobaXterm will display an error message. Common causes include Xming not running, the Buster shell not active, or issues with the `steelhead` server.

7. Loading Definition Files

In the `treegui` window, load a definition file using the control in the upper-right corner (Fig. 34, **A**). Definition files are typically located in:

```
/<yourVersionName>/def/.
```

Successful loading is indicated by text appearing in the central display area of the window.

8. Launching SpecTk

In an `xterm` terminal, navigate to:

```
/user/arisdaq/SpecTk
```

and execute the program. Once `SpecTk` opens:

1. Select `Connect` to from the `SpecTk` drop-down menu.
2. Enter `localhost` as the server name.
3. Enter the four-digit port number displayed in the `SpecTcl` GUI under “`SpecTcl` server is up and listening on port `xxxx`”.

Upon successful connection:

- The `clients connected` counter in `SpecTcl` increments to one.
- The top of the `SpecTk` window reads: `Status: Connected to localhost on port xxxx` (Fig. 36, **A**).

`SpecTk` files may now be opened via `File` → `Open`, or new files may be created and saved.

9. Notes and Best Practices

- Always ensure Xming is running before launching `SpecTcl`.
- Use a dedicated working directory for each analysis to ensure reproducibility.
- Keep a single `SpecTcl` build per analysis.
- If errors occur, check Xming, the Buster shell, and server connectivity.

Appendix B Analysis Program Basics

This appendix summarizes the essential functionality of `SpecTcl` and `SpecTk` needed for cross section analysis. Advanced features are documented in the official references and are not covered here²⁴.

B.1 `SpecTcl`

Most user interactions with `SpecTcl` occur through the `treegui` or the main `SpecTcl` GUI.

Defining Spectra

The `treegui` is used to define spectra and to create, edit, and apply gates. To define a spectrum:

1. Select the spectrum type in the upper-left corner (Fig. 34, **B**). The most common types are one-dimensional and two-dimensional spectra.
2. Choose the parameter(s) from the central list (Fig. 34, **C**). For two-dimensional spectra, parameters are ordered as (x, y) .
3. Set axis limits and the number of bins. Note that excessively fine binning, especially for two-dimensional spectra, may exceed the shared-memory limit and should be avoided.
4. Click **Create/Replace** to define the spectrum.

Existing spectra can be edited by double-clicking their names in the list, modifying fields, and clicking **Create/Replace**. Controls adjacent to the list allow spectra to be deleted, duplicated, or cleared.

²⁴See https://lise.frib.msu.edu/work/SpecTclSpecTk_Tutorial_Haak.pdf for a guide on the structure and analysis pipeline.

Managing Gates

Gates are created and managed through the `treegui` Gate tab (Fig. 35):

- Display existing gates from a loaded definition file by clicking `Update gate list` (Fig. 35, **A**).
- Gates may be defined manually by entering coordinates or interactively via `SpecTk`. Manual definition is recommended only for simple cases.
- Gates can be edited directly in the `treegui`, though this becomes cumbersome for complex gates.

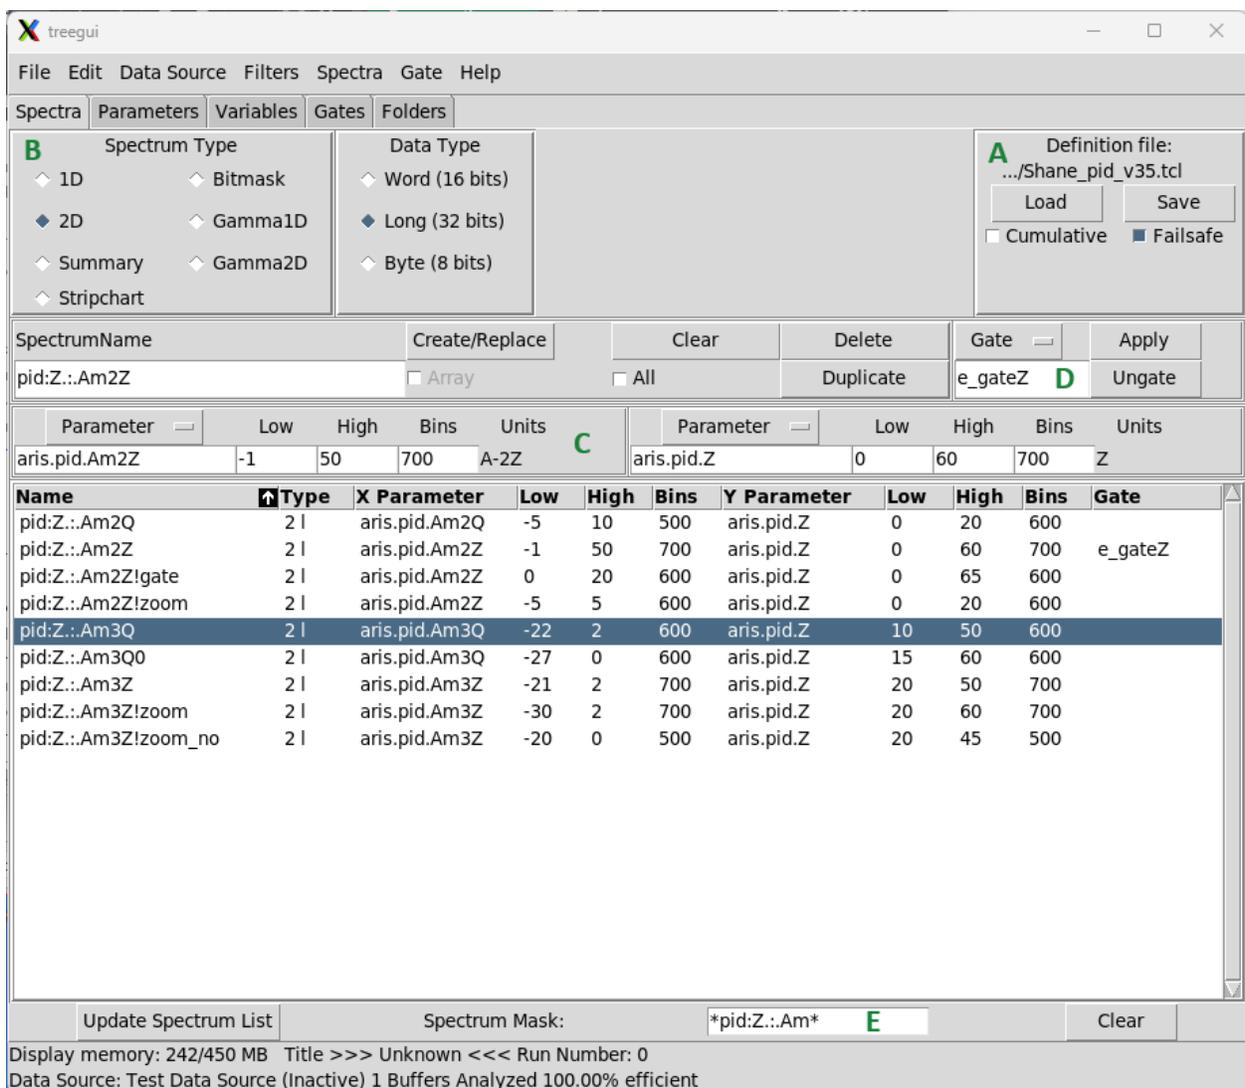


Figure 34: The `SpecTk1` `treegui` Spectrum tab interface.

Compound gates (AND, OR, NOT) can be created as follows:

1. Assign a name beneath the Create/Replace button.
2. Select the gate type from the drop-down menu (e.g., And) (Fig. 35, B).
3. Enter the names of gates to combine or select them from the Gate Select menu.

To apply a gate:

- Select one or more spectra.
- Click **Gate** and choose the desired gate.
- If a gate is already selected, clicking **Apply** assigns it to the highlighted spectra.
- Click **Ungate** to remove all gates from the selected spectra.

A search field at the bottom of the window (Fig. 34, E) allows locating specific spectra or gates using queries enclosed in asterisks, e.g., *SpectrumName*.

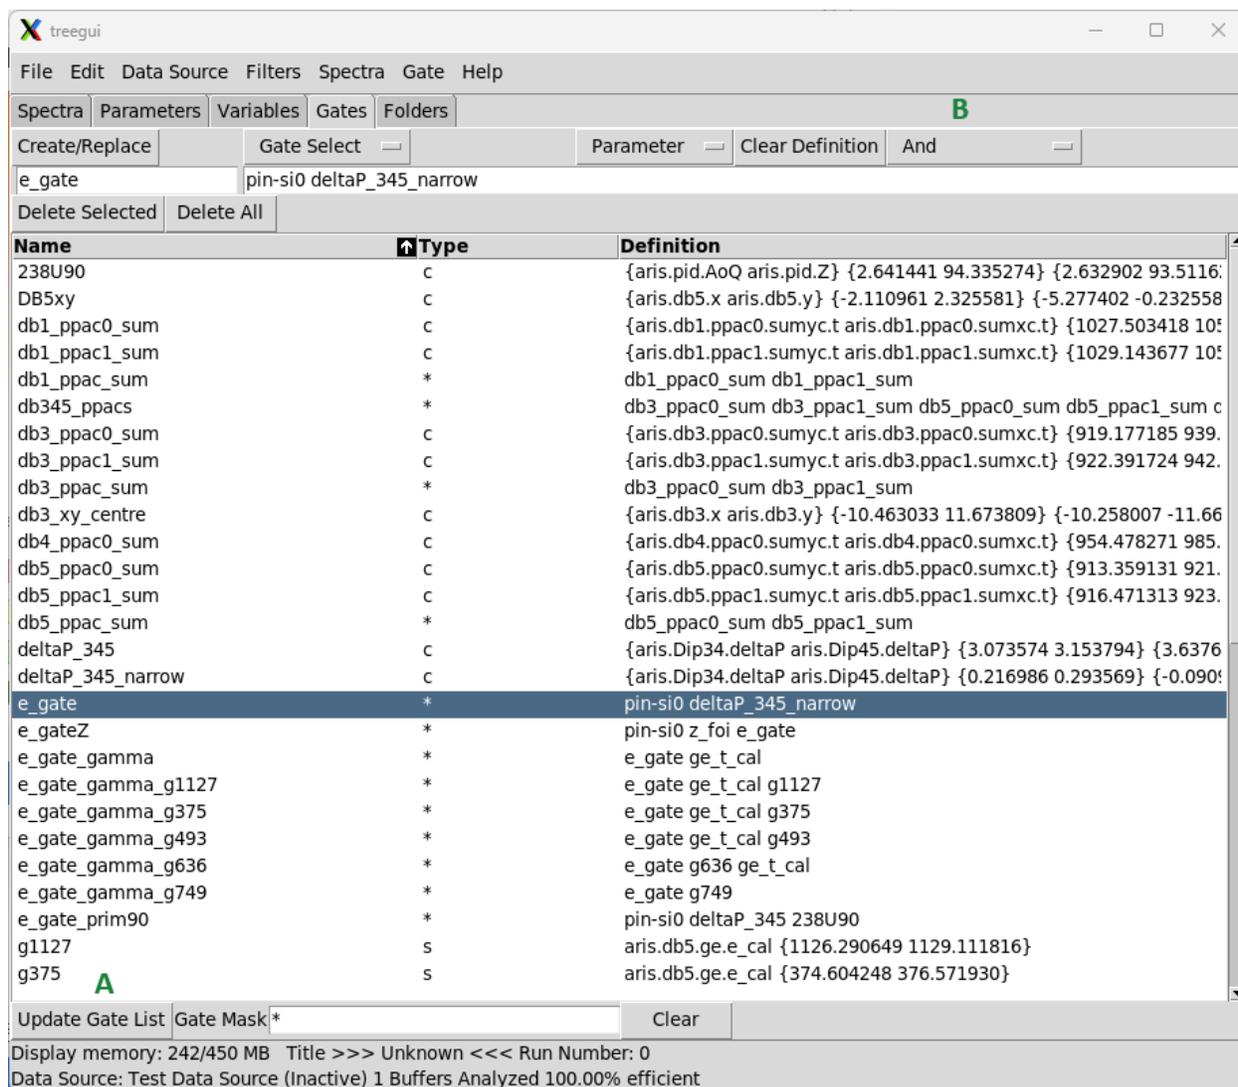


Figure 35: The SpectTcl treegui Gates tab interface.

File Operations

Data files are read using the `SpecTcl` GUI. Key functions include:

- **Clear Spectra:** removes all accumulated data from memory.
- **Pause Analysis:** halts file processing.
- **Read ArisVariable:** forces `SpecTcl` to reload calibration files.
- **Attach Online:** connects to live experimental data.
- **Attach to previously used file:** stops the current run, clears spectra, reloads calibration files, and reattaches the last file (useful during calibration).
- **Attach to next/prior runs:** increments or decrements the run number after reloading calibration files.
- **Attach list of files:** reads a cluster file containing a user-defined list of runs processed sequentially.
- **Detach:** stops reading all files.

B.1.1 `SpecTcl` Filtering Output

One of the core functionalities of `SpecTcl` is the ability to filter event data through gates, as described in the previous section. Within the `SpecTcl` framework, gated event data can be written out to a new file, allowing subsequent stages of the analysis to operate on a reduced and cleaner data set. For very large data sets, this reduction can result in substantial time savings during downstream processing.

Traditionally, `SpecTcl` writes filtered data in a native binary format intended for re-ingestion by `SpecTcl` itself²⁵. In the ARIS `SpecTcl` build, an extended capability has been implemented that allows filtered output to be written directly to a comma-separated values (CSV) file²⁶.

In this CSV format, each column corresponds to a selected spectrum (or parameter), and each row represents a single correlated event. If a particular spectrum is not defined for a given event, the corresponding cell is left empty. This preserves event-by-event correlations while allowing the data to be read by external analysis tools.

The primary purpose of filtering is to extract clean subsets of data that satisfy specific PID or quality gates. However, it is also possible to output essentially all calibrated event data if desired. The analysis tool `Ray Reader`²⁷ was originally developed for use with `LISE++` Monte Carlo data, but its functionality has since been extended to support `SpecTcl`

²⁵This functionality has not yet been successfully integrated into any working version of the ARIS `SpecTcl` analysis chain. Multiple attempts have been made to implement the required pipeline, including by the author of this manual.

²⁶See https://lise.frib.msu.edu/work/SpecTcl/SpecTcl_Filter_Output%20v3.pdf for more information.

²⁷<https://lise.frib.msu.edu/doc/ray-reader-documentation/>

CSV output. Once all calibrations have been finalized, `Ray Reader` can therefore serve as a lightweight substitute for `SpecTcl` for many downstream analysis tasks.

From the `SpecTcl` command line, the following sequence of commands defines and controls a filter:

- `filter -new <filterName> <gateName> {spectra}`
 Defines a new filter, specifying the gate to apply and the list of spectra to be written to the output file.
- `filter -file <fileName.csv> <filterName>`
 Specifies the output file for the filter. By default, files are written to the `SpecTcl` root directory, though an explicit path may be provided.
- `filter -format <filterName> csv`
 Changes the output format from the default XDR format to CSV.
- `filter -enable (-disable) <filterName>`
 Enables or disables the filter. Note that `SpecTcl` will overwrite existing files with the same name.

This filtering capability is required for the trajectory reconstruction procedure discussed in subsection C.3, as well as for the empirical transfer matrix regression described in subsection C.3.2.

B.2 SpecTk

`SpecTk` is a visualization extension for `SpecTcl`, providing functionality not available in `Xamine`.

Page and Spectrum Management

Open configuration options by clicking `Open Drawer` in the lower-right corner of the `SpecTk` window. Pages are created and managed under the `Page` tab, and spectra are assigned to page cells via the `Spectrum` tab (Fig. 36).

Defining ROIs and Gates

The `ROI` tab is used to define regions of interest (local to `SpecTk`) and gates, which are shared with `SpecTcl`. Toggle between ROI and gate creation using the buttons at the top of the window (Fig. 36, **B**).

To create a gate:

1. Select the gate type: `Slice` for one-dimensional plots, `Contour` for two-dimensional plots.
2. Define the region interactively; for contour gates, double-click to complete the shape.

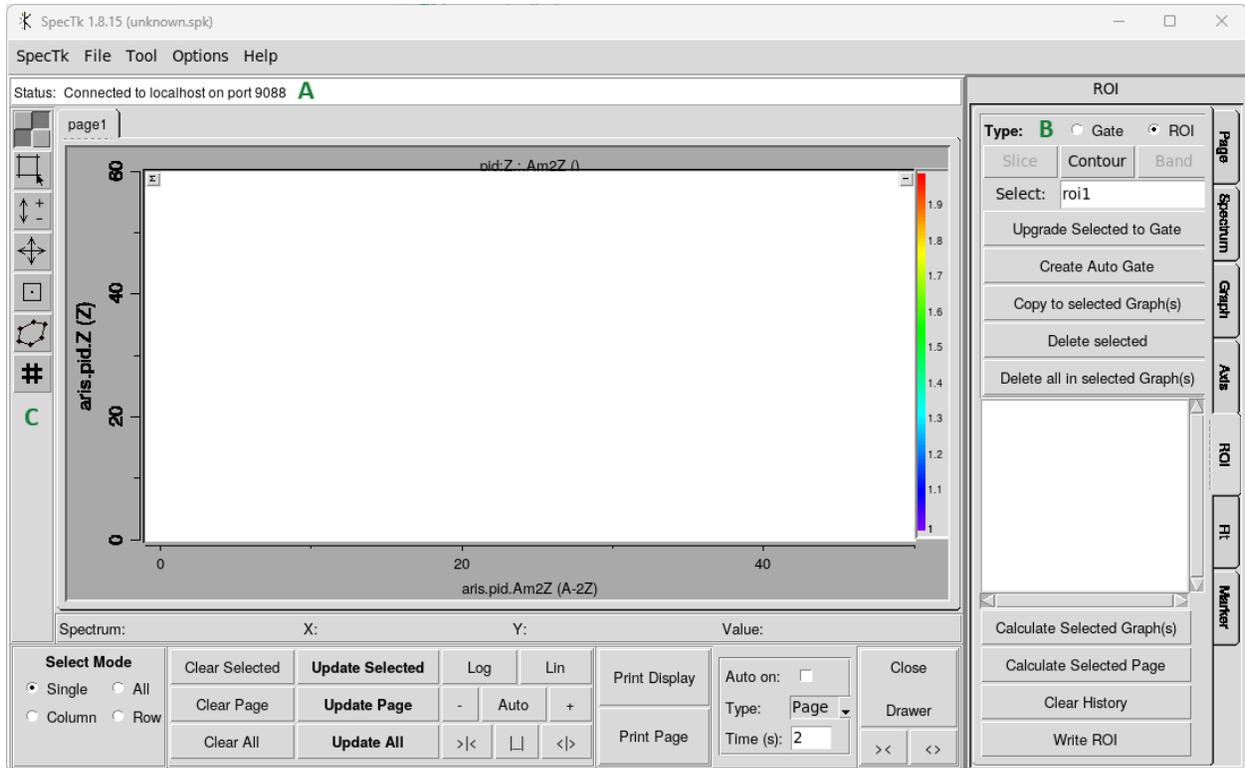


Figure 36: The SpecTk GUI.

3. Assign a name and click **Validate** to create the gate, or **Cancel** to discard.

Existing gates can be modified by:

- Selecting the contour-edit tool on the left side of the SpecTk window (Fig. 36, C) and dragging gate vertices.
- Creating a new gate with the same name. In both cases, the updated definition is synchronized between SpecTk and SpecTcl.

Appendix C Ion Optics and Trajectory Reconstruction Primer

The purpose of this appendix is to introduce concepts related to ion optics and trajectory reconstruction in the context of fragment separators and secondary beam purification. Throughout this appendix, and implicitly throughout the manual, we adopt the TRANSPORT notation of ion optical matrix elements [27].

As with other content in this manual, this discussion of ion optics is not intended to be complete. Rather, it aims to provide the reader with a practical working knowledge of optics as it pertains to cross section measurements, particle identification, and transmission corrections. Many textbooks and monographs have been written on charged-particle optics; interested readers are directed to these for a more formal and complete treatment [28, 29].

C.1 The Transfer Matrix and Other Optical Definitions

An optical system is a series of electromagnetic elements that are used to transport a beam of charged particles from one place to another. In particular, a fragment separator is an optical system designed to purify the secondary beam for delivery to an experiment. The Advanced Rare Isotope Separator (ARIS) consists of various magnetic elements: dipoles, quadrupoles, sextupoles, and octupoles. Additionally, materials such as wedges, position and energy detectors, and slits exist in the beamline to aid in the purification process.

A beam particle moving through an optical system can be represented as a six-dimensional phase-space vector, which contains its transverse position (x, y) , its angle relative to the reference trajectory in both transverse directions (θ, ϕ) , a relative path-length difference between an arbitrary particle and the reference particle (l), and the relative momentum deviation of the particle (δ):

$$\vec{V} = \begin{pmatrix} x \\ \theta \\ y \\ \phi \\ l \\ \delta \end{pmatrix} \quad (10)$$

Here, $\delta = \Delta p/p_0$ is defined with respect to the reference particle momentum p_0 . The path-length coordinate l represents the difference in distance traveled by an arbitrary particle compared to the reference trajectory and plays an important role in time-of-flight measurements and longitudinal beam dynamics.

The reference particle is defined as the ideal particle traversing the system: it follows the designed optical axis exactly and has the reference momentum and charge-state. All other particles are described relative to this trajectory. The trajectory definitions are shown in Figure 37.

Just as every particle has a vector associated with it, every element along the beamline has an associated transfer matrix. The transfer matrix calculates how the beamline element affects the particles as they traverse the system. The general matrix is defined as:

$$M = \begin{pmatrix} (x|x) & (x|\theta) & 0 & 0 & 0 & (x|\delta) \\ (\theta|x) & (\theta|\theta) & 0 & 0 & 0 & (\theta|\delta) \\ 0 & 0 & (y|y) & (y|\phi) & 0 & 0 \\ 0 & 0 & (\phi|y) & (\phi|\phi) & 0 & 0 \\ (l|x) & (l|\theta) & 0 & 0 & (l|l) & (l|\delta) \\ 0 & 0 & 0 & 0 & 0 & (\delta|\delta) \end{pmatrix} \quad (11)$$

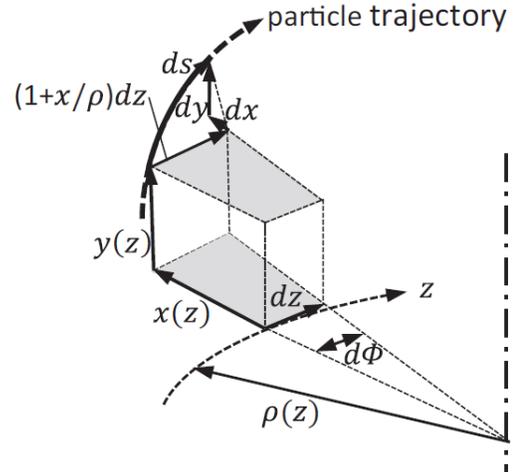


Figure 37: The curvilinear coordinate system used in the definition of the ion optical vector \vec{V} . The optical axis is \vec{s} .

where elements shown as identically zero arise from mid-plane symmetry, the assumed decoupling of horizontal and vertical motion, and the choice of coordinate system. ARIS is designed to be well described by this first-order, uncoupled approximation. Additionally, this matrix implicitly assumes that the x-plane is the dispersive plane. This is assumed in all further discussion.

Definitions of each matrix element are displayed in Figure 38, along with the definition of beam transport in the language of matrices. Elements along the diagonal represent the magnification in each coordinate. Highlighted cross terms represent the focusing of the beam, and the sixth column represents the dispersion of each coordinate.

$$\begin{pmatrix} x_f \\ \theta_f \\ y_f \\ \phi_f \\ l_f \\ \delta_f \end{pmatrix} = \begin{pmatrix} (x|x) & (x|\theta) & 0 & 0 & 0 & (x|\delta) \\ (\theta|x) & (\theta|\theta) & 0 & 0 & 0 & (\theta|\delta) \\ 0 & 0 & (y|y) & (y|\phi) & 0 & 0 \\ 0 & 0 & (\phi|y) & (\phi|\phi) & 0 & 0 \\ (l|x) & (l|\theta) & 0 & 0 & (l|l) & (l|\delta) \\ 0 & 0 & 0 & 0 & 0 & (\delta|\delta) \end{pmatrix} \begin{pmatrix} x_i \\ \theta_i \\ y_i \\ \phi_i \\ l_i \\ \delta_i \end{pmatrix}$$

Final Vector
Transfer Matrix
Initial Vector

Figure 38: Definition of principle elements in the first-order transfer matrix, along with matrix multiplication definition.

The explicit derivation of the transfer matrix for every element in the beamline is far beyond the scope of this manual. Detailed derivations are presented in Ref. [27]. An important distinction should be made here: the matrices above represent the *first-order* transport of a beam. First-order optics adequately describe drift regions, dipoles, and quadrupoles under small-angle and small-momentum-spread assumptions.

The global transfer matrix is defined as the matrix multiplication of all component parts of the transport system. It is important to note that the matrix multiplication is performed from the end of the system to the beginning:

$$M_{tot} = M_f \cdot M_{f-1} \dots M_2 \cdot M_1, \tag{12}$$

where M_i represents the transfer matrix of the optical element.

Higher-order effects – arising from intrinsic field nonlinearities, finite magnet apertures, or large momentum acceptance – are described by higher-rank tensors (second-order transport corresponds to rank-3 tensors, third-order to rank-4, and so on). These higher-order terms give rise to optical aberrations.

An aside on the types of magnets used for beam transport:

- **Dipoles:** magnets used to bend the beam trajectory. Dipoles are the primary source of dispersion in a beamline.
- **Quadrupoles:** first-order focusing elements used to control beam envelopes and imaging conditions. Quadrupoles introduce chromatic aberrations due to momentum-dependent focusing.
- **Sextupoles:** second-order correction magnets primarily used to correct chromatic aberrations introduced by quadrupoles and dipoles²⁸. Sextupoles introduce third-order geometric aberrations.
- **Octupoles:** third-order correction magnets used to compensate aberrations introduced by sextupoles. These introduce fourth-order aberrations.

These magnets are commonly used in beam transport to create the desired optical conditions. They are analogous to lenses in light optics. The focusing strength of a quadrupole magnet is related to the length of the magnet along the beam axis, the size of the aperture the beam passes through, and the amount of current flowing through its wires. For a set geometry, the only variable parameter is the current. Tuning the beam refers to the action of adjusting the focusing strengths (electric current) of the magnet to create the desired transport conditions. Mathematically, this is equivalent to changing the quadrupole matrix to achieve certain global transfer matrix elements.

Aberrations manifest themselves in various ways. No matter the manifestation, they are undesired and significant work has been done to correct them. In the simplest form, these aberrations impact the shape of the beam profile. This can enlarge the beamspot, creating issues for transport downstream. Aberrations can also conspire to *increase* the transmission of unwanted particles. This can be a large source of unwanted background for experiments and must be treated explicitly during the LISE⁺⁺ transmission analysis, as seen in section 5.

In the context of cross section measurements, aberrations are particularly important because they can distort the relationship between measured detector coordinates and the true particle rigidity or trajectory. If left uncorrected, these effects can bias transmission estimates and particle identification, especially for large-acceptance settings.

Magnets beyond the octupole can be implemented in a beamline if control over the high-order aberrations is desired (e.g., decapole (10 poletips) or dodecapole (12 poletips)), but these magnets are not implemented in ARIS.

An alternative approach to correcting high-order aberrations through the introduction of hardware (magnets) is to correct the optics through trajectory reconstruction. Trajectory reconstruction seeks to correct, on an event-by-event basis, for the effects of finite acceptance and optical aberrations using measured detector information and known optical properties of the separator. This approach can effectively extend the usable acceptance of the system without introducing additional hardware. This is discussed in the subsection C.3.

Before discussing the details of trajectory reconstruction, further groundwork must first be laid on ion optical terminology.

- **Phase Space:** the multidimensional space defined by particle coordinates and their conjugate momenta (e.g., $x-\theta$ or $y-\phi$). A beam occupies a finite volume in phase space.

²⁸Dipoles with shaped pole faces can deliberately introduce second-order terms.

- **Emittance:** a measure of the phase-space area occupied by the beam in a given plane. Emittance is conserved under linear, conservative (symplectic) transport.
- **Acceptance:** the region of phase space that an optical system can transport without loss.
- **Transmission:** the fraction of particles produced at the target that fall within the acceptance of the system and reach a downstream location.
- **Dispersion:** the dependence of transverse position or angle on momentum deviation. In a dispersive plane, particles of different rigidity separate spatially.
- **Momentum Resolving Power:** the ability of an optical system to distinguish particles of different momenta, often characterized by the ratio of dispersion to beam size at a focal plane.
- **Magnification:** the scaling of beam size between two locations, given by first-order matrix elements such as $(x|x)$ or $(y|y)$. Large magnification increases spatial separation but also amplifies detector resolution effects.
- **Chromatic Aberrations:** aberrations arising from momentum dependence of focusing elements.
- **Geometric Aberrations:** aberrations arising from finite beam size or angular spread, independent of momentum.
- **Focus:** the beam is focused in a given plane when the final position is independent of the initial angle in that plane $[(x|\theta), (y|\phi) = 0]$.
- **Achromatic:** the beam position and angle are independent of momentum deviation $[(x|\delta) = (\theta|\delta) = 0]$.
- **Aberration:** a higher-order optical term that distorts the image at a focal plane (e.g., chromatic aberrations such as $(x|\theta\delta) \neq 0$).
- **Beam Envelope:** the maximum transverse extent of the beam as it propagates through the system.
- **Parallel Beam:** a beam for which particle angles are small and narrowly distributed about zero (small angular divergence), often realized at a waist.
- **Beam Waist:** a location where the beam envelope reaches a minimum and the angular divergence is maximized.
- **Coordinate Correlation:** a statistical dependence between phase-space coordinates (e.g., x correlated with δ in a dispersive plane). Trajectory reconstruction exploits these correlations to infer unmeasured quantities.

These terms may be used in conjunction with each other to describe the optical properties of the system. For example, a focus does not need to be achromatic. Optics specific to ARIS are discussed in the next subsection.

C.2 The Optics of ARIS

ARIS is a third-generation fragment separator designed to efficiently collect, purify, and deliver rare isotope beams produced via projectile fragmentation. The system is composed of two major optical sections: a vertical pre-separator followed by a horizontal C-Bend. A schematic overview of the full system is shown in Figure 1. The following subsections describe the optical principles and operating modes that are most relevant for cross section measurements and trajectory reconstruction.

C.2.1 Momentum Compression

The pre-separator constitutes the first stage of separation in ARIS and plays a critical role in shaping the longitudinal phase space of the secondary beam. Of particular importance is its ability to perform *momentum compression*, a technique that reduces the momentum (or velocity) spread of the fragments and thereby increases the effective momentum acceptance of the separator.

The principles and formal description of momentum compression are detailed in Ref. [30]. In brief, momentum compression exploits the correlation between momentum and energy-loss in matter. Fragments with higher momentum lose more energy when passing through a wedge-shaped degrader placed at a dispersive focal plane, while lower-momentum fragments lose less. With appropriate optical matching, this differential energy-loss reduces the momentum spread of the beam downstream of the degrader.

The principal benefit of momentum compression is an increase in the number of fragments that fall within the downstream momentum acceptance of the separator. The trade-off is an increase in transverse phase space, both as a result of the momentum compression and because the wedge material introduces additional angular and spatial spread through multiple scattering and energy-loss straggling. A schematic illustration of the momentum compression process is shown in Figure 39.

After the pre-separator, the beam is transported into the C-Bend. The transition between these two sections is optically significant because it involves an effective rotation of the beam coordinates: the dispersive plane changes from the vertical direction (y) in the pre-separator to the horizontal direction (x) in the C-Bend. This rotation is dictated both by the physical layout of the facility – transporting the beam from below ground to ground level – and by optical considerations.

Following momentum compression, the transverse phase space in the dispersive plane of the pre-separator expands in proportion to the compression factor (three in the standard ARIS configuration). If left unaltered, this enlarged phase space would reduce downstream acceptance. By rotating the coordinate system, the dispersive plane is transferred to the orthogonal transverse direction, which was not affected by the momentum compression process. This preserves acceptance in the C-Bend while retaining the benefits of the reduced momentum spread.

C.2.2 Optical Modes of Operation

All particle identification measurements in ARIS are performed in the C-Bend section of the separator. The C-Bend contains five principal focal planes, labeled DB1 through DB5. DB1

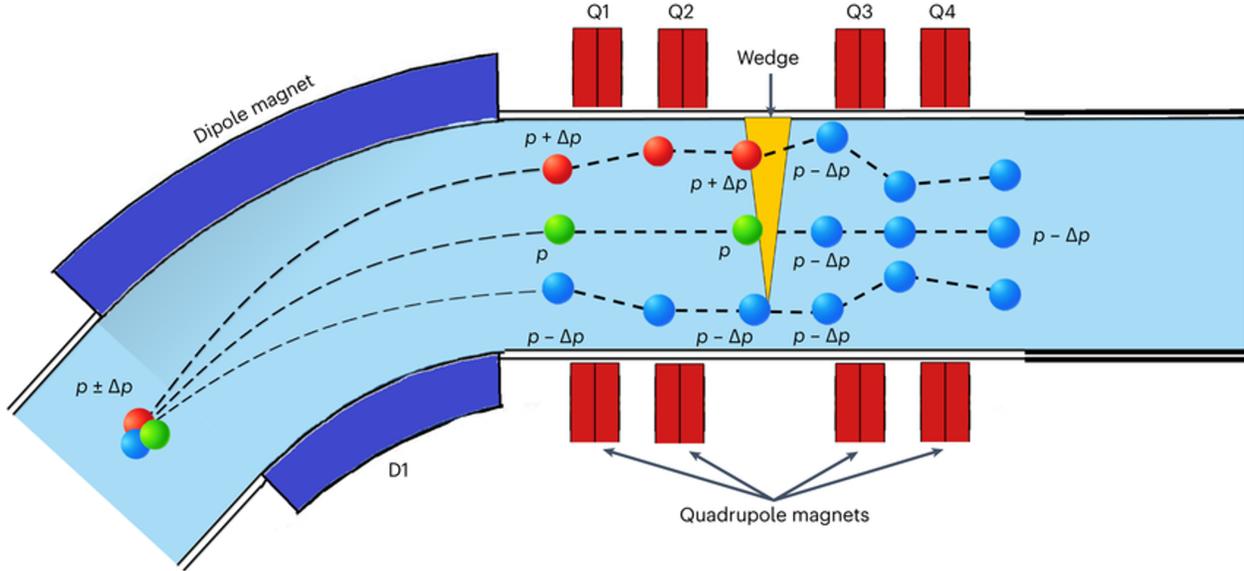


Figure 39: Schematic overview of momentum compression. Identical particles with different initial momenta (color scheme) are spatially separated at a dispersive focal plane and pass through a wedge-shaped degrader, resulting in a reduced momentum spread downstream. Paths of the particles are traced through the system. Taken from Ref. [31].

marks the entrance to the C-Bend and corresponds to the coordinate rotation point between the pre-separator and the horizontal transport section.

In the most commonly used optical configuration, known as the high-acceptance **CB2 mode**, DB2 and DB4 are dispersive focal planes, while DB3 and DB5 are achromatic. DB4 plays a central role in particle identification, as position measurements at this plane are used to correct the magnetic rigidity ($B\rho$) of individual fragments. DB5 marks the end of the fragment separator.

An alternative configuration, the high-resolution **CB1 mode**, may also be employed. In this mode, the optics are tuned such that only a single dispersive focal plane exists at DB3, prioritizing momentum resolution over acceptance.

While operating in CB2 mode, DB5 may be configured in one of two optical conditions. In the *focus* setting, DB5 serves as a focal plane for experiments terminating at the separator, such as dedicated new-isotope searches. In the *transmission* setting, DB5 is tuned to efficiently transport the beam to downstream experimental areas.

The pre-separator itself supports two primary modes of operation: k1 and k3. The k3 mode, which provides a momentum compression factor of three, is by far the most commonly used configuration. The k1 mode, in which no momentum compression is applied, is employed in cases where compression offers no benefit or is detrimental to transmission or resolution. The pre-separator and C-Bend modes may be combined arbitrarily, though the k3CB2 configuration is overwhelmingly dominant in experiments performed at FRIB.

The ideal first-order optics for the k3CB2 mode are shown in Figure 40. Magnifications are shown in the left plot, focuses are shown in the middle plot, and the spatial dispersion is shown in the right plot. The plotted global matrix elements illustrate the evolution of focusing and dispersion along the separator and clearly show the coordinate rotation between

the pre-separator and C-Bend.

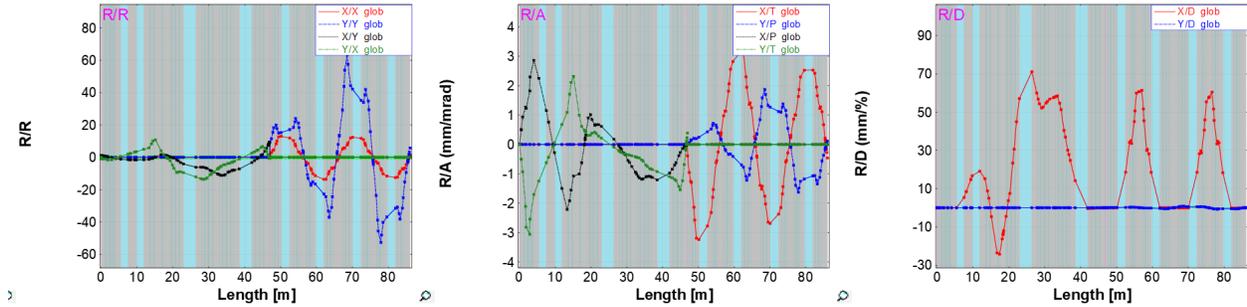


Figure 40: Ideal first-order optics of the k3CB2 mode of ARIS. The plotted quantities represent global transfer matrix elements as a function of position along the beamline. The rotation of the dispersive plane following the pre-separator is clearly visible.

C.2.3 Secondary Beam Purification

This section provides a brief primer on how secondary beams are purified for experiments at FRIB. Understanding this process is essential context for both rare isotope experiments and the production cross section measurements described throughout this manual.

Secondary beam purification is necessary because projectile fragmentation is an inherently stochastic process. In a high-energy collision between the primary beam and the production target, nucleons are randomly abraded from the projectile nucleus and the resulting excited pre-fragment subsequently ablates particles according to its excitation energy (see subsection 7.2). As a result, a single reaction produces a broad distribution of fragments spanning many atomic numbers, mass numbers, and charge-states. In principle, nearly any lighter combination of protons and neutrons may be produced²⁹. The role of the fragment separator is to isolate a specific subset of these fragments – the isotopes of experimental interest – from this overwhelming background.

In practice, secondary beam purification occurs through a combination of two complementary mechanisms: *optical separation* and *physical separation*. These mechanisms are interleaved along the separator and act together to progressively increase the purity of the transmitted beam.

C.2.3.1 Optical Separation

Optical separation exploits the dependence of particle trajectories on magnetic rigidity, $B\rho = p/q$, where p is momentum and q is the ionic charge. Dipole magnets are tuned to the rigidity corresponding to the fragment of interest. Particles with different A/q values experience different bending angles in the magnetic field and are therefore displaced from the reference trajectory. Fragments whose rigidity differs sufficiently from the tuned value will strike the beam pipe or magnet apertures and be removed from the beam.

²⁹Nucleon pickup reactions can also occur for large impact parameters, though with much smaller probability at intermediate energies.

This rigidity-based filtering defines one of the primary *selection lines* of a fragment separator. A schematic illustration of these selection lines is shown in Figure 41. While powerful, optical separation alone is insufficient: many fragments with different (A, Z) combinations can share similar A/q values and therefore follow nearly identical trajectories through the dipoles.

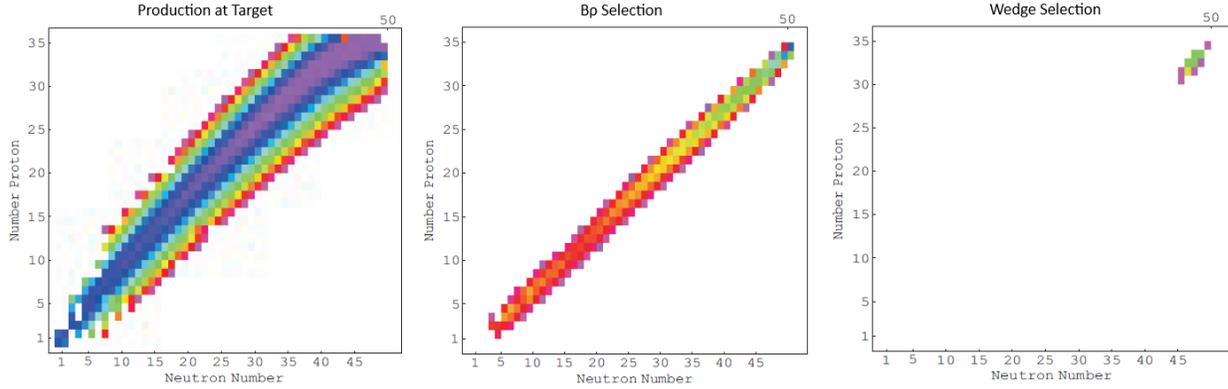


Figure 41: Illustration of selection lines and the resulting increase in beam purity achieved through magnetic rigidity ($B\rho$) selection and energy-loss degradation (dE). From Ref. [32]

C.2.3.2 Physical Separation

Physical separation relies on the interaction of beam particles with matter to further purify the secondary beam beyond what is achievable through magnetic rigidity selection alone. This is accomplished primarily through energy-loss mechanisms and geometric constraints imposed by the beamline.

The most important physical filtering element is the energy-loss degrader, typically implemented as a wedge-shaped material placed at a dispersive focal plane. As fragments traverse the degrader, they lose energy according to the Bethe–Bloch formalism, with the leading dependence scaling approximately as

$$\frac{dE}{dx} \propto Z_{\text{eff}}^2,$$

where Z_{eff} is the effective ionic charge of the fragment. As a result, fragments with different atomic numbers but similar magnetic rigidities will experience different energy-losses when passing through the degrader. The wedge geometry is chosen such that the fragment of interest emerges with the correct momentum to remain centered in the acceptance of the downstream dipole magnets, while contaminants are over- or under-degraded and subsequently bent out of the beam. This combination of dispersive optics and energy-loss defines a second, independent selection line in the separator.

A further complication in physical separation arises from the presence of multiple ionic charge-states. Following fragmentation and passage through materials in the separator, ions may not be fully stripped of their atomic electrons, especially for higher- Z fragments or at lower beam energies. Since the magnetic rigidity depends on the ratio $B\rho = p/q$, different

charge-states of the same isotope follow different trajectories through the separator. In some cases, partially stripped ions can possess nearly identical magnetic rigidities to fully stripped ions of a different isotope, allowing them to survive purely optical selection and contaminate the transmitted beam.

Energy-loss degraders help mitigate charge-state ambiguities because ions with different effective charges lose different amounts of energy in matter, even if their initial magnetic rigidities are similar. Additionally, operation at sufficiently high beam energies strongly favors fully stripped charge-states, reducing the overall complexity of the transmitted beam. Nevertheless, residual charge-state contamination is an intrinsic feature of in-flight separation and must ultimately be resolved during particle identification and transmission analysis. This consideration is particularly important near the limits of separator acceptance and in experiments aimed at discovering or characterizing the most weakly bound isotopes.

As mentioned previously in this manual, beam particles may interact with a wedge not only through energy-loss but also via nuclear reactions, undergoing further fragmentation within the material itself. These secondary reactions contribute to contamination of the transmitted beam and can reduce the yield of the isotope of interest. While optical separation can remove a portion of this contamination, additional physical separation stages are often required to maintain acceptable beam purity. This consideration motivates the use of multiple wedges in ARIS.

Depending on the experimental requirements, two wedges may be present in the system: the primary wedge in the pre-separator and a secondary wedge at DB2. The pre-separator wedge is typically thick and is responsible for momentum compression as well as initial physical separation. In contrast, the DB2 wedge is deliberately made thinner to minimize additional fragmentation while still providing incremental separation of contaminants that survive the first stage. The DB2 wedge does not perform momentum compression; rather, it serves as a refinement stage that improves beam purity at the cost of only a modest reduction in transmission.

The use of multiple wedges therefore represents a balance between purity and yield. For experiments targeting extremely rare isotopes, where background suppression is critical, the additional physical filtering provided by a second wedge can be essential. Conversely, for experiments where maximum transmission is prioritized, the use of the DB2 wedge may be avoided altogether.

In addition to degraders, physical apertures such as slits and collimators may be inserted at strategic locations to block regions of phase space known to be populated by unwanted species. While effective, slits necessarily reduce transmission and must be used judiciously, particularly in experiments targeting extremely rare isotopes. A simplified schematic of a two-dipole, one-wedge purification stage is shown in Figure 42.

C.3 Trajectory Reconstruction

As mentioned previously, trajectory reconstruction improves particle identification (PID) resolution by correcting for event-by-event kinematic variations that are not captured by nominal optical settings. In practical terms, it allows one to undo the effects of finite beam emittance, momentum spread, and higher-order optical aberrations by reconstructing the actual path each particle takes through the separator.

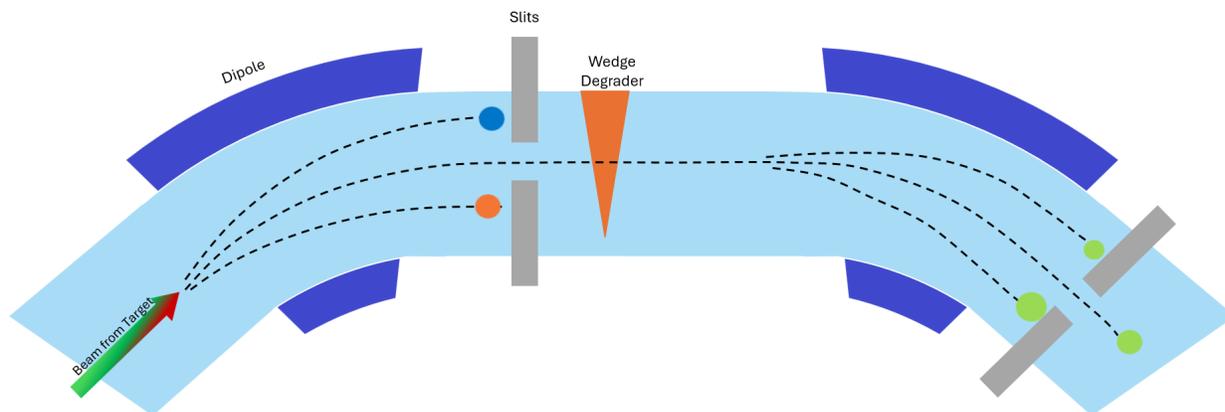


Figure 42: Schematic illustration of secondary beam purification using a wedge degrader between two dipole magnets. This demonstrates the combined action of optical ($B\rho$) and physical (energy-loss) selection lines shown in Figure 41.

In a broad sense, trajectory reconstruction is the process of inferring a particle's full phase-space coordinates at one location in the beamline (typically the production target) from measurements made downstream. This inference relies on precise measurements of final positions, angles, time-of-flight, and energy-loss, together with unambiguous particle identification and detailed knowledge of the magnetic settings and optical response of the system.

While trajectory reconstruction is most often introduced as a tool for improving PID, its utility extends well beyond kinematic corrections. By reconstructing particle trajectories across the acceptance of the separator, one can empirically characterize the optics, dispersion, and effective acceptance of the system. In this way, trajectory reconstruction serves both as a correction technique and as a diagnostic tool for validating optical models.

Work has been conducted to characterize ARIS using this approach. Data collected at DB5 have been traced backward through the system to the production target. In addition to improving PID resolution, these studies provide insight into the transport properties of the separator and the kinematic properties of the reaction products themselves. Preliminary trajectory reconstruction results for ARIS are shown in Figures 43 and 44.

The foundation of trajectory reconstruction as applied to PID lies in the fact that individual particles do not follow the reference trajectory. Each fragment experiences small but measurable deviations in momentum and path length as it traverses the separator. These deviations directly impact two key observables used in PID: time-of-flight and magnetic rigidity.

The path length traversed by a particle is directly tied to the measured time-of-flight. In standard analysis, a fixed path length – corresponding to the reference trajectory – is assumed when converting timing signals to velocity. However, particles that deviate from the reference orbit traverse slightly longer or shorter paths. These differences introduce a systematic broadening of the time-of-flight distribution that cannot be removed by detector resolution alone. By reconstructing the actual trajectory of each particle, event-by-event path length corrections can be applied, significantly improving the intrinsic ToF resolution.

Magnetic rigidity, $B\rho$, quantifies the bending strength required to transport a particle

through a magnetic field and is proportional to its momentum divided by its charge. As discussed in Appendix E, $B\rho$ is a critical input to the determination of the mass-to-charge ratio A/q . In analogy to the path length case, particles exhibit event-by-event momentum deviations relative to the reference particle. While the dipole magnets set a global rigidity, local deviations can be inferred from position measurements at dispersive focal planes. In ARIS, these corrections are applied using position measurements at DB4, which is a dispersive focus in the CB2 optics mode.

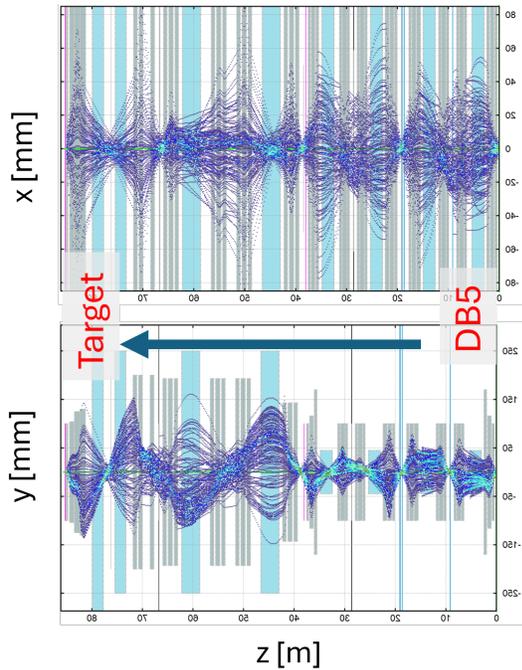


Figure 43: Calculations of particle trajectories traced backward from DB5 to the target position.

Corrections to the magnetic rigidity are particularly important when ARIS is operated in CB1 mode, where only a single dispersive focus is available. The impact of these corrections is illustrated in Figure 45. Prior to trajectory reconstruction, the same charge-state appears to populate different values of A/q at different magnetic rigidities – an unphysical result arising from uncorrected momentum deviations. After applying trajectory-based corrections, this correlation is removed, and the expected discrete charge-state structure is recovered.

Additional measurements performed during a ^{198}Pt experiment investigated the consistency of momentum reconstruction across different segments of the separator. The reconstructed momentum deviations for the DB1-3 and DB3-5 sections show a linear correlation with unit slope (Figure 46), indicating good agreement between the calculated dispersions and the measured transport. This agreement provides an important validation of both the optical model and the reconstruction procedure.

As discussed in paragraph 2.2.1.2, `SpecTcl` supports first- and second-order kinematic corrections for trajectory reconstruction. In principle, however, corrections of arbitrary order may be applied. Previous work with the S-800 spectrograph demonstrated trajectory reconstruction to fifth order using COSY Infinity³⁰ transport maps [6, 33]. These higher-order approaches offer a path forward for further refinement of reconstruction techniques in ARIS as detector resolution and statistics improve.

C.3.1 Detector Calibration and Connection to Transmission

Trajectory reconstruction is only as reliable as the detector information on which it is based. Accurate reconstruction therefore requires that all position-, timing-, and energy-sensitive detectors used in the analysis be well calibrated, stable over time, and correctly aligned within the optical coordinate system of the separator.

³⁰<https://www.bmtdynamics.org/cosy/>

Position-sensitive detectors at dispersive and non-dispersive focal planes must be calibrated to convert raw signals into physical coordinates with linear response and minimal distortion across the active area. Any nonlinearity or misalignment directly propagates into reconstructed momentum and path length, degrading both PID resolution and the physical interpretability of reconstructed quantities. Similarly, timing detectors must be precisely calibrated and synchronized, as residual offsets or drifts can mask the improvements gained from path length corrections. Energy-loss detectors used for PID and charge-state determination must also be stable, as miscalibration can introduce ambiguities that cannot be resolved through reconstruction alone.

Beyond its role in improving event-by-event kinematics, trajectory reconstruction provides a powerful mechanism for empirically determining the acceptance and transmission of the separator. By mapping reconstructed trajectories back to the production target, one can determine which regions of phase space are transmitted to the final focal plane and which are lost due to optical constraints, apertures, or higher-order aberrations. This information can be used to construct realistic, data-driven acceptance functions that complement or refine those obtained from ion-optical simulations.

In the context of cross section measurements, this feedback is particularly important. Transmission corrections are a dominant systematic uncertainty, especially for rare isotopes produced near the limits of acceptance. Trajectory reconstruction allows one to quantify these effects experimentally by correlating reconstructed phase-space coordinates with observed yields. In this way, reconstruction serves as a bridge between measured yields and absolute production cross sections, strengthening the connection between experiment and model.

C.3.1.1 Caveats and Limitations

While trajectory reconstruction is a powerful tool, it is not a panacea. The method relies on the accuracy of the underlying optical model, the stability of magnetic fields, and the quality of detector calibrations. Systematic errors in any of these inputs can lead to over-correction or the introduction of artificial correlations. In addition, finite detector resolution and incomplete phase-space coverage limit the order to which meaningful corrections can be applied; higher-order terms may fit statistical fluctuations rather than true physical effects.

Trajectory reconstruction also implicitly assumes that particles traverse the separator without undergoing additional interactions. Secondary reactions in degraders, detectors, or other materials introduce irreducible uncertainties that cannot be corrected kinematically. As such, reconstruction should be applied judiciously, validated against known benchmarks (e.g., charge-state behavior or rigidity scans), and interpreted in concert with simulation and physical intuition.

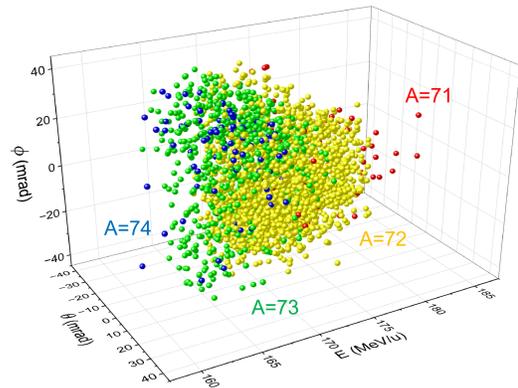


Figure 44: Reconstructed distributions of iron fragments at the target position as a function of energy and transverse angles.

When used carefully, trajectory reconstruction significantly enhances the resolving power and quantitative reliability of fragment separators like ARIS. When used uncritically, it can obscure as much as it reveals. As with all correction techniques, its greatest strength lies in informed application and rigorous cross-checking.

C.3.2 Empirical Transfer Matrix Determination

There exists another method to extract transfer matrices between two focal planes at which transverse position and angle measurements are available: regression analysis [34].

In its simplest form, regression is the determination of correlations among measured variables. In the context of ion optics, multivariate regression provides a natural and powerful framework for reconstructing the effective transfer matrix between two locations in the beam-line using experimentally measured phase-space coordinates. Given sufficiently precise measurements of transverse positions and angles at two focal planes, the elements of the transport matrix may be inferred directly from the data.

In this approach, each transfer matrix element corresponds to the slope of the correlation between a final coordinate and one of the initial coordinates. For example, the element $(x|x)$ represents the correlation between the final horizontal position and the initial horizontal position, while $(x|\theta)$ quantifies the sensitivity of the final position to the initial horizontal angle. When momentum information is included, dispersion terms such as $(x|\delta)$ can also be extracted.

As emphasized in the previous section, accurate detector calibration is essential. Even small systematic offsets in position or angle measurements can bias the extracted matrix elements.

This requirement places stringent demands on both the absolute calibration and the resolution of the position-sensitive detectors. Modern PPAC systems typically achieve millimeter-level position resolution, which is sufficient for this purpose. Although each individual measurement carries an intrinsic uncertainty, these uncertainties are statistically averaged down when a large ensemble of particles is used in the regression.

Assuming well-calibrated detectors, the measured phase-space coordinates encode the true transport of particles through the system, including contributions from optical aberrations of all orders. In practice, however, only first- and second-order matrix elements can be reliably extracted using regression, as higher-order correlations require prohibitively large

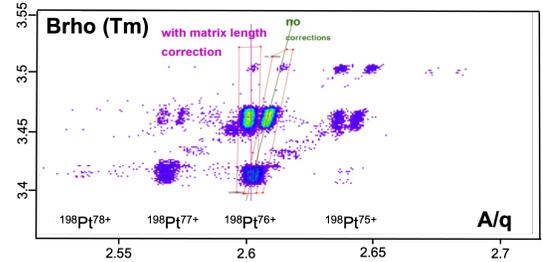


Figure 45: ^{198}Pt charge-states plotted as a function of A/q and $B\rho$ during a momentum scan in CB1 mode. Prior to trajectory reconstruction, an unphysical correlation is observed.

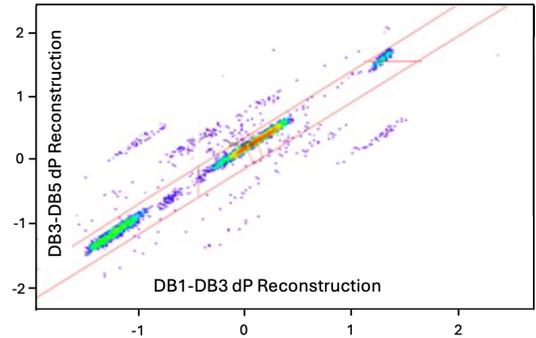


Figure 46: Reconstructed momentum deviations, $\Delta B\rho/B\rho$, for the DB1-3 and DB3-5 segments of the C-Bend for ^{198}Pt fragments.

data sets and rapidly become dominated by detector resolution and unmodeled material effects.

A key point in this analysis concerns the treatment of the momentum deviation δ . In ARIS, the momentum coordinate is commonly expressed operationally as $\delta = \Delta B\rho/B\rho$. For optical segments in which no wedge or significant material is present – such as the DB3–DB5 section of the C-Bend – the fragment momentum is conserved to good approximation. Under these conditions, the momentum deviation at the initial and final focal planes is effectively the same, $\delta_{\text{initial}} \approx \delta_{\text{final}}$. As a result, dispersion-related correlations extracted from regression are not distorted by trajectory reconstruction corrections, which act only as common-mode shifts to δ . This ensures that the regression-derived dispersion matrix elements reflect the true optical response of the system and are not circularly influenced by the simulated transport matrix already implemented in `SpecTc1`.

The principal advantage of extracting transfer matrices through regression is that it provides an empirical cross-check of the ion-optical simulations. Agreement between the simulated and experimentally derived matrices indicates that the optics are well understood and that the dominant transport properties are correctly modeled. Discrepancies between the two are expected to arise from finite detector resolution, energy and angular straggling in materials, and higher-order aberrations that are not explicitly included in first-order simulations. In this way, regression-based matrix extraction serves both as a validation of the optical model and as a diagnostic tool for identifying limitations in the simulated description of the beamline.

Currently, only a single PPAC is installed at DB4. As a consequence, while a complete first-order transfer matrix may be extracted between DB3 and DB5 – where both position and angle measurements are available – only partial transfer matrices can be obtained for the subsegments DB3–DB4 and DB4–DB5.

The planned installation of a second PPAC at DB4 will remove this limitation by enabling direct measurements of both transverse position and angle at that focal plane. This upgrade will unlock an additional and powerful optical consistency check. For a given particle, the transport through the system is deterministic and identical for each event. Consequently, the product of the experimentally extracted matrices for the two subsegments, should reproduce the independently extracted DB3–DB5 transfer matrix within experimental uncertainties:

$$M_{\text{DB3} \rightarrow \text{DB5}} = M_{\text{DB4} \rightarrow \text{DB5}} \cdot M_{\text{DB3} \rightarrow \text{DB4}} \quad (13)$$

Agreement between these matrices would provide a stringent validation of the measured optics, while systematic deviations would offer direct insight into unmodeled aberrations, detector systematics, or material effects localized to a specific segment of the C-Bend.

Appendix D Quantitative Gate Analysis

Quantitative gate analysis is required whenever absolute fragment yields are critical to the scientific outcome of an experiment. This situation most commonly arises when yields are low, such that small changes in gate placement produce large fractional changes in the

extracted counts. In these regimes, visual inspection of gates alone is insufficient, and a systematic, quantitative evaluation of gate performance is necessary.

This appendix supplements the qualitative cleaning philosophy described in subsection 2.2.3 of the manual. In that section, gates are introduced as tools to suppress unphysical events while preserving valid fragment counts. Here, the emphasis is on measuring how well those gates perform and on ensuring that the final yields are robust against reasonable variations in gate definition.

Quantitative gate analysis is particularly important when searching for very exotic isotopes or claiming first observations³¹, where statistical margins are small and the consequences of over- or under-cleaning are severe.

D.1 Conceptual Framework

The core idea of quantitative gate analysis is to partition PID space into two classes of regions:

- **Good regions:** gates that correspond to identified isotope blobs and contain physically meaningful events.
- **Bad regions:** gates placed between isotope blobs, where events are presumed to be dominated by background, mis-reconstructed particles, or unphysical combinations of observables.

Cleaning gates should reduce the population of bad regions as much as possible while minimally affecting the population of good regions. A gate that strongly suppresses both regions indiscriminately is not optimal, even if it appears visually clean.

An example of this partitioning is shown in Figure 47, where both isotope and between-blob regions are marked. The black line separates primary fragments (above) from wedge products and aberration-transported fragments (below). This determination was made using LISE⁺⁺, as outlined in subsection 5.2. In addition to well-populated isotope blobs, gates should also be placed on low-statistics or borderline isotopes, as these are often the most sensitive to over- or under-cleaning.

D.2 Procedure

The quantitative analysis proceeds as follows:

1. Starting from the final calibrated PID spectrum (with all PID variables validated), define gates around many identified isotope blobs. These constitute the initial set of good regions.
2. Define additional gates in the spaces between isotope blobs. These gates should follow the general topology of the PID spectrum and sample regions where no physical fragments are expected.

³¹As a fun aside, FRIB maintains a public database of isotope discoveries and ranks researchers by number discovered: <https://frib.msu.edu/public-engagement/learning-resources-and-programs/brief-history-of-rare-isotopes/rankings>

3. Record the initial event counts in every good and bad gate before applying any additional cleaning gates.
4. Apply one cleaning gate at a time (e.g., timing, energy-loss, position, or dispersion cuts), and record the change in counts in each region.

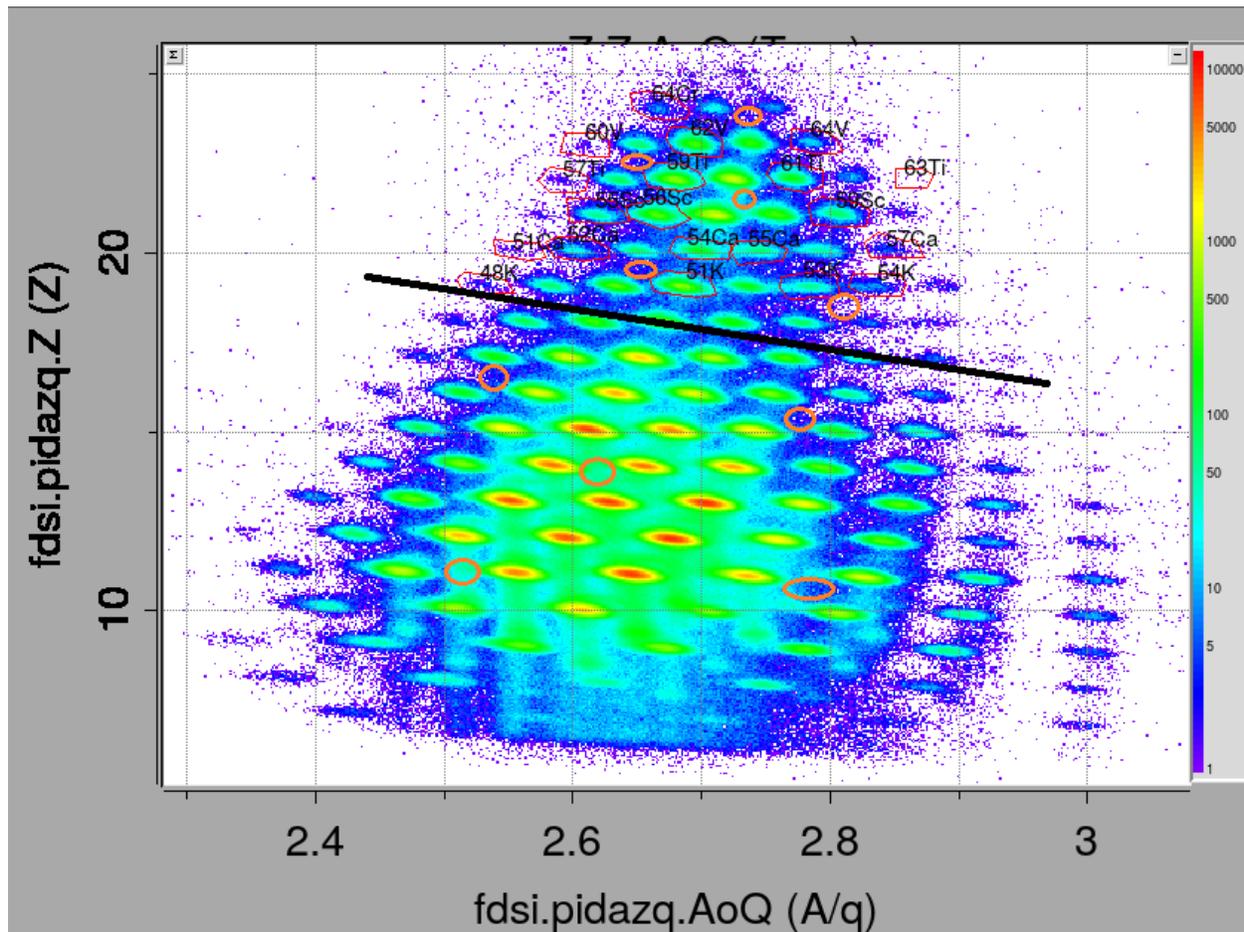


Figure 47: Final calibrated PID spectrum with fragment-identification gates and between-blob (“bad region”) gates defined. Orange circles indicate regions between isotope blobs used to quantify background suppression. The black line separates primary fragments (above) from wedge products and aberration-transported fragments (below).

A cleaning gate is considered effective if it significantly reduces counts in bad regions while producing only a small reduction in good-region counts. Because some unphysical events can mimic valid fragment observables, it is expected that all cleaning gates will reduce counts in both classes to some degree. The goal is not zero loss in good regions, but rather an acceptably small and well-understood loss.

D.3 Diagnosing Problematic Gates

If the application of a particular cleaning gate produces a large reduction in one or more good regions, further investigation is required. In such cases:

- Project the affected isotope gate onto the cleaning spectrum associated with the problematic cut.
- Examine whether valid events cluster near the gate boundary.
- If so, slightly loosen the gate and repeat the quantitative evaluation.

This procedure ensures that physically meaningful events are not being excluded due to overly aggressive gate placement.

Special care is required for gates involving energy-loss correlations. In energy-loss versus energy-loss spectra, valid fragments often extend along diagonal ridges corresponding to physically allowed trajectories. As illustrated in Figure 9, some legitimate events may lie far from the central density of the blob but remain aligned with the diagonal trend. Gates in these spaces must be sufficiently elongated to retain these events, particularly for rare isotopes near the limits of acceptance.

D.4 Best Practices

Quantitative gate analysis should be performed after final PID calibration and isotope identification have been completed, but before yield extraction and statistical uncertainty assignment. Performing this analysis earlier risks conflating PID calibration issues with gate-performance effects. Performing this analysis later risks invalidating the yield statistic derived in subsection 3.2.

The results of a quantitative gate analysis should be documented explicitly, including:

- The fractional change in good and bad region counts for each cleaning gate.
- Any gates that required loosening or modification.
- The final gate definitions used for yield extraction.

This documentation not only strengthens confidence in the reported yields, but also provides transparency and reproducibility for future re-analyses or comparisons to independent measurements.

Appendix E Particle Identification

This appendix familiarizes the reader with the physical principles, measured observables, and mathematical relationships underlying particle identification (PID) as implemented in this analysis. A complete first-principles derivation of the relevant equations is beyond the scope of this manual.³² Instead, the emphasis here is on clarifying how experimentally measured quantities are combined to infer intrinsic particle properties, and on documenting the assumptions and limitations implicit in this procedure.

³²More detailed explanations and derivations can be found in the appendix of Ref. [35]

E.1 Overview of Measured and Derived Quantities

In practice, PID proceeds by combining several directly measured experimental observables – magnetic rigidity, time-of-flight, and energy-loss – to infer intrinsic particle properties. Particle identification follows a strictly ordered chain of dependencies. Magnetic rigidity and time-of-flight determine the mass-to-charge ratio A/q . The total kinetic energy, together with A/q and β , determines the ionic charge-state q . Only after q is known can the atomic mass A be reconstructed. Finally, the atomic number Z is inferred from energy-loss measurements, which depend explicitly on velocity.

Errors introduced at any stage propagate forward and cannot be corrected downstream. For this reason, PID calibration must proceed in the order:

$$B\rho, \text{ ToF} \rightarrow \beta, \gamma \rightarrow A/q \rightarrow dE, \text{ TKE} \rightarrow q \rightarrow Z \rightarrow A.$$

It is therefore useful to distinguish explicitly between experimentally measured observables and quantities inferred through PID calculations:

- **Measured:** magnetic rigidity $B\rho$, time-of-flight, detector energy-loss signals, and total deposited energy.
- **Derived:** β , γ , A/q , ionic charge-state q , atomic mass A , and atomic number Z .

None of the intrinsic particle properties (A , q , or Z) are measured directly. Instead, they are reconstructed through calibrated combinations of measured quantities and empirical model relations. This distinction is critical when diagnosing PID artifacts, identifying failure modes, or assessing systematic uncertainties.

E.2 Mass-to-Charge Ratio and Mass Reconstruction

Throughout this manual, the quantity used for particle identification and mass reconstruction is referred to as the mass-to-charge ratio, denoted A/q . Strictly speaking, this is an approximation. What is measured experimentally is not the nuclear mass number A , but the total mass M of the ion, including any bound electrons, divided by its charge-state q . The physically correct quantity is therefore M/q .

For an ion with atomic number Z , mass number A , and charge-state q , the total mass can be written as

$$M = M_{\text{nucleus}} + (Z - q)m_e - E_{\text{bind}}/c^2 \quad (14)$$

where m_e is the electron mass and E_{bind} is the total electronic binding energy. In fragment separators operating at intermediate energies, ions are highly stripped, and $q \approx Z$ for most species of interest. As a result, the number of bound electrons is small, and the contribution of electronic binding energies is negligible compared to the nuclear mass.

Quantitatively, the nuclear mass is of order $A \times 931.5 \text{ MeV}/c^2$, while the mass contribution from even several electrons is at the level of a few MeV/c^2 . This represents a relative correction of order 10^{-4} or smaller, well below the resolution of in-flight particle identification systems and far below other dominant sources of uncertainty such as time-of-flight

resolution, energy-loss straggling, and optical aberrations. Consequently, for the purposes of isotope identification and cross section measurements, the distinction between M/q and A/q is numerically insignificant.

Despite this, the use of A/q rather than M/q is more than a casual shorthand. In experimental practice, the goal of mass reconstruction is to identify the integer mass number A associated with a given nuclear species, independent of its charge-state. Expressing results in terms of A/q makes this goal explicit and aligns naturally with how particle identification plots are interpreted: different charge-states of the same isotope appear as distinct loci in A/q , but they all correspond to the same underlying mass number A . This convention also facilitates direct comparison with reaction models, mass tables, and systematics, which are all formulated in terms of nuclear mass number rather than total ionic mass.

For these reasons, the notation A/q is retained throughout this manual. It should be understood as an experimentally reconstructed proxy for the true mass-to-charge ratio M/q , with the implicit assumption that electronic contributions to the ion mass are negligible at the level relevant for rare isotope production and cross section analysis.

The fundamental relations governing mass and charge reconstruction are:

$$\frac{A}{q} = \frac{B\rho}{3.1071 \beta \gamma}, \quad (15)$$

$$q = \frac{\text{TKE}}{(\gamma - 1) u A/q}, \quad (16)$$

$$A = \left(\frac{A}{q}\right) q, \quad (17)$$

where

- A/q is the mass-to-charge ratio (amu/ e),
- $B\rho$ is the magnetic rigidity (T·m),
- $\beta = v/c$ and $\gamma = (1 - \beta^2)^{-1/2}$ are the relativistic kinematic parameters,
- q is the ionic charge-state,
- TKE is the total kinetic energy of the ion,
- u is the atomic mass unit, and
- A is the atomic mass number.

Equation (15) follows directly from the relativistic momentum relation $p = \gamma m v$, rewritten in units appropriate for heavy-ion spectrometers. The numerical constant in the denominator represents the combination uc/e , evaluated such that $B\rho$ is expressed in T·m and A/q is obtained in amu/ e . This relation highlights the central role of both magnetic rigidity and time-of-flight measurements in determining A/q .

It should be noted that Equation (15) implicitly assumes that higher-order optical aberrations and large-acceptance effects have been adequately corrected. In fragment separators

operating at large momentum or angular acceptance, residual higher-order distortions can introduce correlations that broaden or distort the reconstructed A/q distributions if not properly accounted for. Such effects must be addressed through appropriate optical corrections and validated through PID consistency checks, as performed in subsection 2.2.1.

Equation (16) follows from the relativistic kinetic energy expression $T = (\gamma - 1)mc^2$. Determination of the ionic charge-state therefore relies on accurate measurements of A/q , β , and the total kinetic energy. This expression further assumes that the reconstructed total kinetic energy represents the full residual kinetic energy of the fragment after all upstream materials, including degraders and detectors, and that no significant unmeasured energy-loss occurs downstream of the final energy-loss detector.

Equation (17) is algebraically straightforward but physically important: the atomic mass number A cannot be uniquely identified without either a reliable charge-state determination or an explicit assumption about the charge-state of the fragment³³. This coupling of rigidity–timing information to energy-loss information motivates the emphasis placed on total kinetic energy calibration elsewhere in this manual.

In many intermediate- to high-energy rare-isotope experiments (typically $\gtrsim 100$ – 400 MeV/u), fragments are predominantly fully stripped at the final focal plane, such that $q \approx Z$. When this condition is satisfied, the reconstructed charge-state distribution collapses to a single dominant peak, and the assumption $q = Z$ may be safely applied for yield extraction.

For heavier fragments or lower velocities, however, multiple charge-states may be populated. In such cases, explicit charge-state reconstruction using total kinetic energy is essential. Failure to resolve or properly account for partially stripped charge-states can lead to incorrect mass assignments and distorted isotopic yields. The presence of multiple charge-states should therefore always be assessed directly in Z – q or equivalent PID spectra.

E.3 Atomic Number Reconstruction

The final PID observable is the atomic number Z , which is inferred from the measured energy-loss of the ion in silicon detectors. This procedure relies on heavy-ion energy-loss in matter and incorporates both theoretical guidance and empirical calibration. The equations presented here reflect the simplified, implementation-specific forms used internally by `SpectCl`.

A reduced, velocity-dependent energy-loss term is first defined as

$$dE_v = \frac{1}{\beta^2} \log \left(\frac{5930}{\beta^2 - 1} \right) - 1, \quad (18)$$

which captures the dominant velocity dependence of the stopping power. This expression is inspired by the Bethe-Bloch formalism but does not represent the full stopping-power equation. Rather, it provides a monotonic, velocity-dependent scaling that enables empirical linearization of Z in PID space. The numerical constants have been evaluated for silicon detectors and absorbed into an empirical parameterization optimized for PID calculations.

An intermediate variable is then defined as

³³For example, assuming fully stripped ions.

$$zv = \frac{e_{\text{reduced}}}{(dE_v)^{Z_{\log\beta}}}, \quad (19)$$

where e_{reduced} is the measured energy-loss signal and $Z_{\log\beta}$ is a user-defined parameter controlling the velocity dependence of the effective Bragg curve. In `SpecTcl`, the reduced energy-loss may be formed from a PIN detector, silicon stack segment 0, segment 1, or a weighted combination of multiple detectors.

Finally, the atomic number is calculated as

$$Z = zv Z_{\text{slope2}} + \sqrt{zv} Z_{\text{slope}} + Z_{\text{offset}}, \quad (20)$$

where Z_{slope} , Z_{slope2} , and Z_{offset} are user-defined calibration parameters that map the empirical energy-loss variable onto physical atomic numbers. These parameters are detector- and experiment-specific and must be determined through calibration. A discussion of this calibration can be found in subsection 2.2.5.

Equations (18)–(20) again highlight the central role of accurate time-of-flight measurements: the inferred atomic number depends explicitly on β , and timing errors will broaden or distort Z distributions if not properly controlled. Additionally, miscalibrated time-of-flight can also shift the Z distribution.

E.4 Assumptions and Validity

The PID formalism described here assumes intermediate- to high-energy heavy ions for which relativistic kinematics are valid and energy-loss follows smooth Bethe-Bloch-like behavior. It further assumes that charge-states are sufficiently well resolved and that detector response is linear and stable over the relevant dynamic range. At low velocities, for very light ions, or in the presence of unresolved multiple charge-states, the reliability of the reconstructed A , q , and Z may degrade. In such cases, additional data cleaning, alternative gating strategies, or more restrictive assumptions may be required.

E.5 Common PID Failure Modes

Several characteristic features in PID spectra indicate underlying problems:

- **Tilted or curved PID bands** typically indicate timing or magnetic-rigidity miscalibration.
- **Broad or asymmetric Z distributions** often reflect poor energy-loss calibration or unresolved charge-states.
- **Multiple parallel mass bands** are usually the result of multiple charge-states being populated.
- **Discontinuities across detector boundaries** suggest mismatched gain or offset calibrations between detector segments.
- **Fragment-dependent offsets** appearing only for specific Z or A values often indicate uncorrected non-linearities or saturation effects in energy-loss detectors.

Diagnosing and correcting these issues at the PID level is essential before proceeding to yield extraction or cross section analysis.

E.6 Identification of Blobs

Correctly assigning physical meaning to clusters (“blobs”) in particle-identification (PID) spectra is a critical step in rare isotope analysis. A PID blob corresponds to a specific combination of atomic number Z , ionic charge-state q , and atomic mass A , as defined by Equations (15)–(17). Misidentification at this stage propagates directly into incorrect yield extraction and cross section determination.

No single observable uniquely identifies all fragments in an in-flight fragmentation experiment. Instead, blob identification relies on a hierarchy of methods, combining definitive physical anchors with supporting consistency checks. These methods are summarized below and should be applied collectively whenever possible.

E.6.1 Primary Identification Methods

The following methods provide the most robust and model-independent identification of PID blobs and should be prioritized whenever available. Note that not every PID method is available for every experiment. Therefore, knowledge of, and ability to perform, each type of particle identification is an invaluable skill both for online and offline analysis.

PID Holes (Unbound Nuclei)

Certain nuclei are unbound and therefore cannot be produced as transmitted fragments. Their absence manifests as persistent gaps (“holes”) in otherwise smooth isotopic chains in PID space. Because these holes arise from nuclear structure rather than experimental conditions, they provide powerful topological anchors for isotope identification. Once a hole is identified, neighboring blobs may be assigned relative to it with high confidence.

An example of this type of identification is shown in Figure 48. The isotopic line for fluorine is identified. The two isotopes ${}^9\text{B}$ and ${}^{16}\text{F}$ may be unambiguously identified due to their absence and their relative location. It should be noted that this type of identification requires knowledge of the region in which the fragment separator is tuned (near the boron proton dripline, in this case). This method further assumes that transmission and acceptance are sufficiently smooth that experimental inefficiencies do not mimic physical holes.

Isomer Tagging via γ -Ray Coincidences

When germanium detectors are available, isomeric decays provide unambiguous isotope iden-

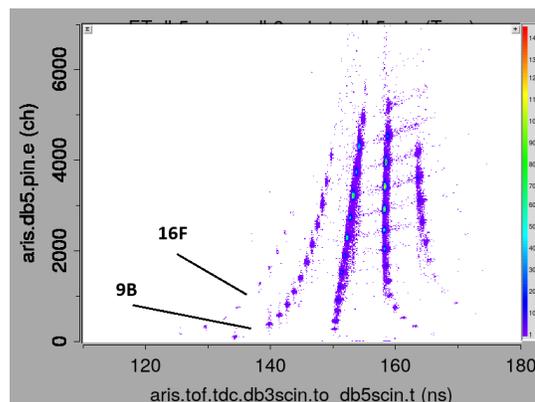


Figure 48: Example of an unbound nucleus PID (${}^9\text{B}$ and ${}^{16}\text{F}$).

tification. Coincidences between PID-selected fragments and characteristic γ rays – combined with known half-lives and decay schemes – directly link a PID blob to a specific nuclide. This method is independent of optics modeling and serves as a definitive validation of blob assignments.

Figure 49 shows an example workflow for isomer identification. The analysis starts with a well-resolved PID spectrum (A/q vs Z , Fig. 49, **A**) that has preliminary cleaning gates applied to it. It is important to note that positive isomer identification will *not* identify charge-states, as the isomer is indifferent to the number of electrons orbiting the nucleus³⁴.

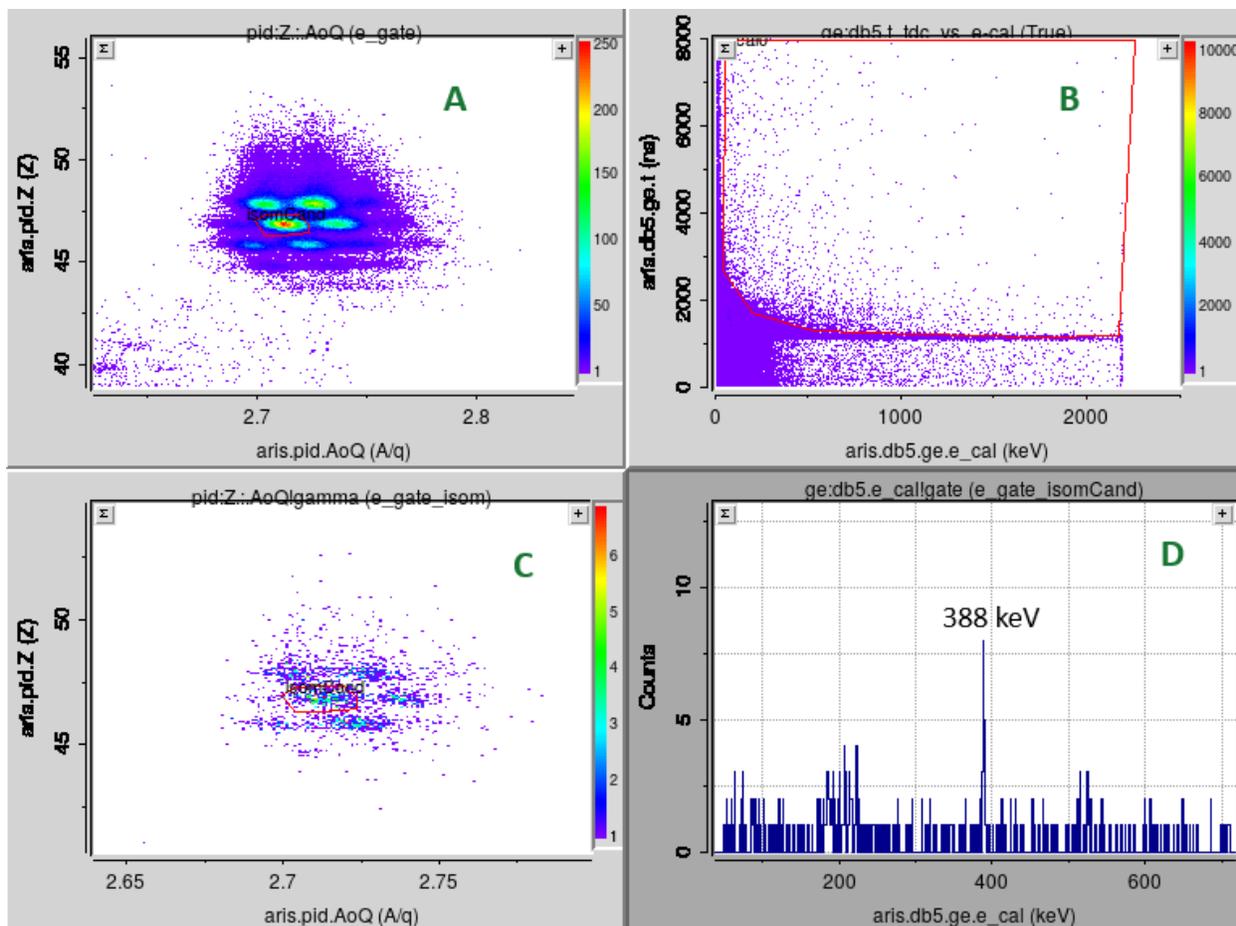


Figure 49: Example of the isomer PID process. Clockwise from top left: (A) finalized and cleaned PID. (B) Germanium detector’s registered energy vs time. (C) Isomer gate projected onto PID spectrum. (D) One-dimensional γ -energy spectrum with cleaning and blob gate applied.

After a high-quality PID spectrum has been made, create a gamma-ray gate (Fig. 49, **B**). This ungated spectrum plots the germanium detector’s registered energy vs time. The gate is drawn such that prompt gammas and background are excluded.

³⁴This isn’t strictly true: for isotopes with a large branching ratio for internal conversion, higher charge-states may impact gamma-ray statistics and mask the identity of the fragment. Consequently, isomer tagging confirms the nuclear identity (A , Z), but does not, by itself, resolve the ionic charge-state.

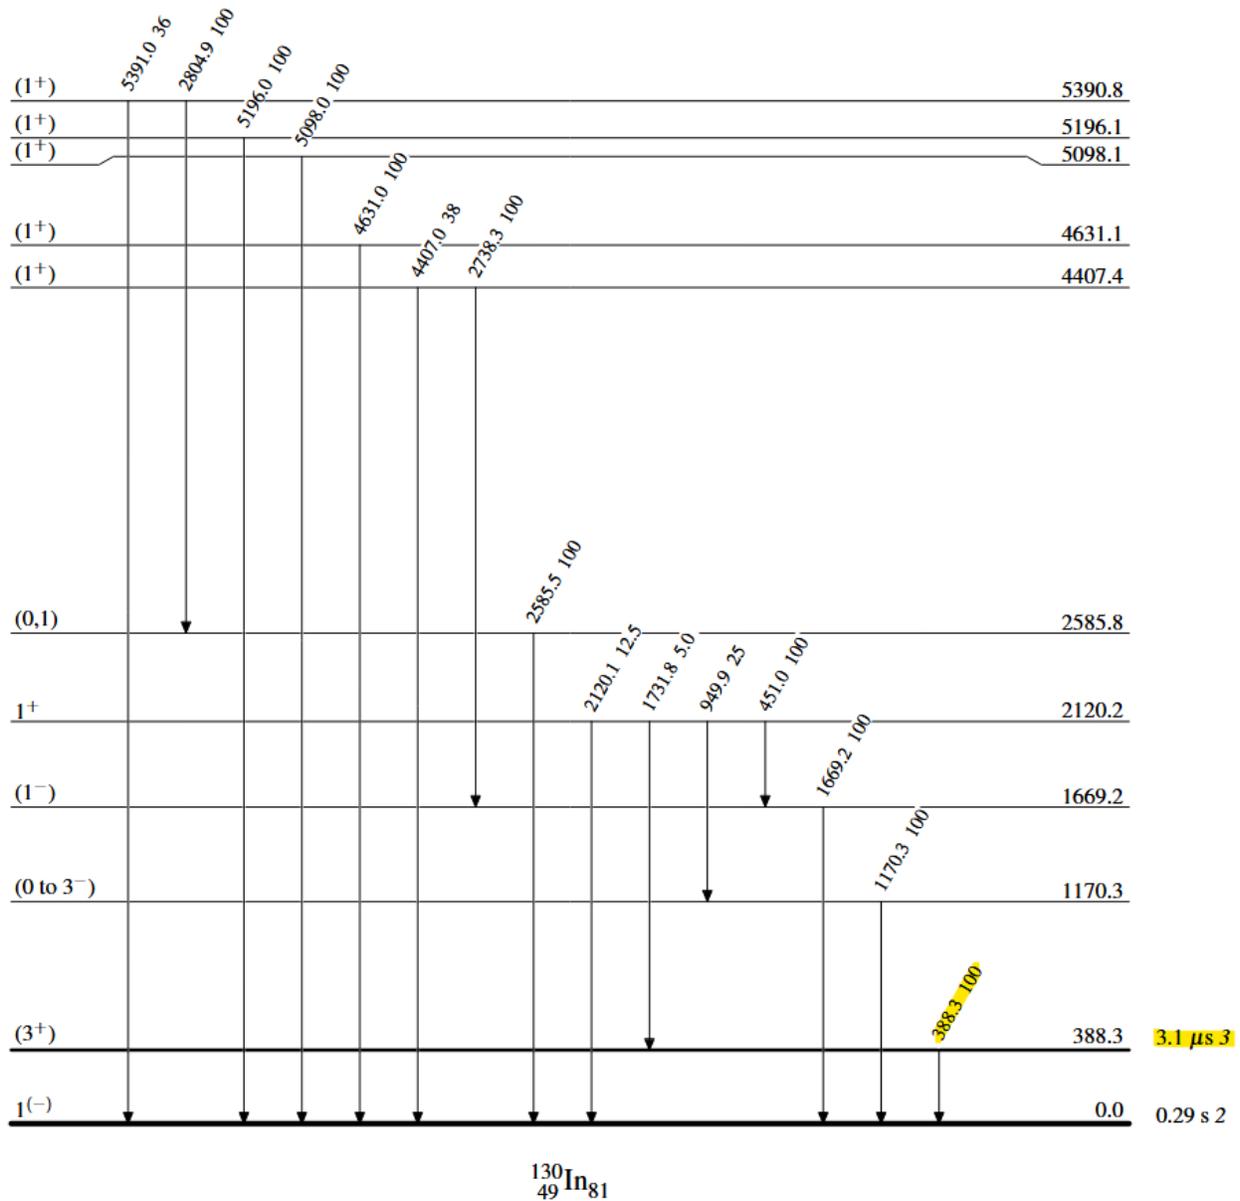


Figure 50: ^{130}In gamma-ray level scheme from the NNDC ENSDF. The detected isomeric state is highlighted.

This isomer gate is then projected back onto the PID plot (Fig. 49, C). Candidate isotopes are the remaining bright spots. Gate these with a temporary name (e.g., “isomCand”) and project the corresponding gamma-rays onto a one-dimensional energy spectrum. Sharp peaks with more counts than background are likely the isomeric decays. These gamma energies are compared to known isomers³⁵ in the region of interest. The isomer detected has an energy of 388 keV. The spectrometer was tuned for ^{130}In , which has a known isomer of that energy. The gamma-ray energy levels are shown in Figure 50. The highlights show the

³⁵It is recommended to use the LISE isomer database or the the NNDC ENSDF: <https://www.nndc.bnl.gov/ensdf/>

energy and half-life of the state. With this gamma-ray detection, we have confirmed the identity of the gated blob, and may fix the remaining PID accordingly.

The process may also be run in reverse: start with the one-dimensional gamma energy spectrum, gate on large peaks, project those gates onto the PID spectrum and see which, if any, blobs are strong emitters. This method is usually used in regions where isomers are known to exist and be populated (e.g., the strong ^{88}Br 110.9 keV gamma ray).

There is a fine art to isomeric identification; it takes time to master the craft. Important considerations when looking for potential isomeric states is to be aware of the half-life and angular momentum of the decaying state. Isomer tagging is only effective for states with lifetimes comparable to or longer than the flight time from the target to the focal plane (~ 550 ns). Similarly, isomers with very large angular momenta are not expected to be populated in great numbers during the fragmentation.

If isomers prove hard to identify, some recommended steps are to change the number of bins in the one-dimensional energy spectrum to potentially resolve multiple overlapping peaks that might conspire together to populate multiple blobs or combine peaks to boost the counts. Additionally, one may change the placement of the isomer gate – if there is too much background, bring the gate up, if there is too few counts, place the gate closer to the prompt gamma blob.

A final note is that in the extremely exotic region of the nuclear chart, we may detect isomers that have yet to be published. As such, these “extra” peaks in the spectrum may initially throw off identification; remember to stay alert for these possibilities.

Once positive isotopic identification has been made, the remaining calibration of the PID may proceed.

Primary and Pilot Beam Calibration

Fragments with well-known properties, including the primary beam itself or reference cocktail components, are used to establish absolute calibration of energy-loss, time-of-flight, and magnetic rigidity, including any kinematic or dispersion-dependent corrections applied in the optics model. These anchors validate the PID reconstruction defined by Equations (15)–(20) and ensure that at least one location in PID space is known with certainty. All other assignments are made relative to this calibrated reference.

Figure 51 shows a PID calibration centered around charge-states of the primary beam (^{238}U). The calculated A/q and Z values must match the known cocktail. Once the detectors are calibrated on the primary beam, the secondary fragments may be transmitted and PID may be confirmed.

In a similar manner, a pilot beam, which ideally has characteristics similar to the desired rare isotope beam may be transmitted and calibration performed based on it.

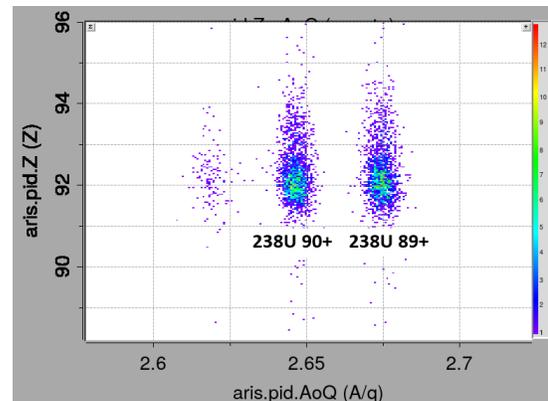


Figure 51: Example of detector calibration and PID using the charge-states of the primary beam (^{238}U).

In the example presented in Figure 52, the pattern of the relative intensity is compared between experiment and LISE⁺⁺ predictions. The cross checking of the observed intensity pattern can help verify that PID is correct.

Charge-State Behavior Consistency

Correct isotope assignments exhibit physically plausible charge-state distributions following total kinetic energy calibration. Fully stripped charge-states dominate at sufficiently high energies, while partially stripped states form predictable, lower-intensity branches. Unphysical behavior (e.g., $q > Z$ or isolated charge-states) signals misidentification or calibration failure. This check is particularly powerful when multiple charge-states of the same isotope are visible simultaneously and directly relies on the charge determination described in Equation (16).

Transport and Optical Consistency

Fragments correctly identified in PID space should behave consistently throughout the beamline. Correlations between position, angle, dispersion, and transmission across focal planes should align with expectations from ion-optical transport. Anomalous transmission or strong higher-order aberrations associated with a particular blob may indicate secondary reactions or incorrect identification.

E.6.3 Failure Modes and Best Practices

Purely visual alignment of PID blobs, unconstrained Gaussian fitting, or blind reliance on simulation-predicted centroids are insufficient for reliable identification. Assumptions such as $q = Z$ must always be validated experimentally, especially for high- Z beams where incomplete stripping is common.

Best practice is to establish at least one primary identification anchor, then apply multiple corroborative methods to confirm the full isotopic assignment. When ambiguities remain, conservative exclusion of questionable isotopes is preferable to propagating uncertain identifications into cross section results.

E.7 Summary and Recommended Workflow

Reliable particle identification requires a structured and conservative approach. Based on the principles outlined in this appendix, the recommended workflow is as follows:

- Establish and validate timing and magnetic-rigidity calibrations before attempting PID.
- Anchor at least one isotope using a primary identification method (primary beam, pilot beam, PID holes, or isomer tagging).
- Verify charge-state behavior and confirm that reconstructed q values are physically plausible.

- Use secondary methods (isotopic continuity, yield systematics, transport consistency) to corroborate assignments.
- When ambiguities remain, conservatively exclude questionable blobs rather than propagating uncertain identifications into yield or cross section results.

Following this workflow minimizes the risk of systematic PID errors and ensures that subsequent yield extraction and cross section analyses are physically meaningful and reproducible.

Appendix F Momentum Integrated Total Cross Section

As mentioned in subsection 7.1, a distinction must be made between partial and total production cross sections. This appendix describes the motivation for, and methodology behind, the determination of total production cross sections from experimentally measured partial cross sections.

A *partial cross section* is defined as the production cross section measured at a single magnetic rigidity setting. In practice, a given $B\rho$ selects a narrow slice of the fragment's longitudinal momentum distribution. For a fixed fragment and charge-state, higher $B\rho$ values correspond to fragments that have retained more longitudinal momentum – i.e., fragments for which less energy has been dissipated during the projectile fragmentation process.

The *total production cross section* of a fragment is a fundamental observable that encapsulates the full probability for its creation in a given reaction system. Unlike partial cross sections measured at a single magnetic rigidity setting, the total cross section integrates over the entire longitudinal momentum distribution of the fragment and is therefore independent of the specific optical tuning of the separator. For rare isotopes, particularly those produced at the limits of stability, the total cross section is often the most meaningful quantity: it sets the absolute production rate, governs experimental feasibility, and provides a stringent benchmark for reaction models.

From a practical standpoint, experiments rarely transmit the full momentum distribution of a fragment in a single optics setting. Fragment separators such as ARIS are tuned to a narrow range of magnetic rigidity $B\rho$, corresponding to a limited slice of the fragment momentum space. As a result, the yield measured at any one setting represents only a fraction of the total production. Reconstructing the total cross section therefore requires sampling the fragment yield over multiple $B\rho$ settings and integrating the resulting momentum distribution.

The magnetic rigidity, $B\rho = p/q$, provides a direct connection between the separator settings and the fragment kinematics. For a given charge-state q , changing $B\rho$ corresponds to selecting fragments with different longitudinal momenta. A $B\rho$ scan thus acts as a controlled probe of the momentum distribution generated in the reaction. Low- $B\rho$ settings preferentially transmit fragments from the low-momentum tail of the distribution, while higher- $B\rho$ settings sample the high-momentum tail. Near the centroid of the distribution,

yields are largest and acceptance effects are typically most benign; toward the tails, transmission decreases and sensitivity to optical aberrations, energy-loss effects, and charge-state contamination increases.

The shape of the momentum distribution itself carries physical meaning. In projectile fragmentation, it reflects a convolution of several processes: the Fermi motion of nucleons in the projectile, recoil from nucleon abrasion, evaporation during ablation, and energy-loss in the target. Deviations from simple Gaussian behavior – such as asymmetric tails or broadening – may indicate changes in the reaction mechanism, increased evaporation, or contributions from pickup or deep-inelastic processes. By measuring yields across a range of $B\rho$ settings, these features can be observed experimentally and compared directly with model predictions.

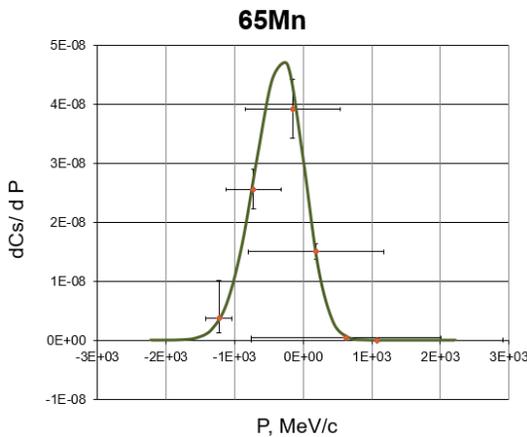


Figure 53: Example of a fitted momentum distribution from six partial cross section measurements.

For most fragments, the momentum slits are initially assumed to select the central region of the momentum distribution, which can often be approximated as Gaussian. During dedicated momentum scans, the tails of the distribution are intentionally sampled to study deviations from this simple behavior. This process is shown in Figure 53, where cross sections were measured at different momenta and a distribution has been fit to the data points. The vertical error bars are the error in the cross section measurement, which have been the point of discussion throughout this manual. The horizontal error bars represent the width of the momentum acceptance of the spectrometer.

For total cross section determination, the role of the $B\rho$ scan is therefore twofold. First, it enables a numerical integration of the experimentally measured momentum distribution, correcting for acceptance and transmission at each setting. Second, it provides a diagnostic of the reaction dynamics themselves, offering insight into how the fragment was produced and how well the underlying physics is captured by models such as abrasion–ablation. In this sense, total cross section measurements are not merely bookkeeping exercises for yield normalization, but a critical bridge between experimental observables and the microscopic mechanisms governing rare isotope production.

LISE⁺⁺ includes several parameterizations for the momentum distributions produced in projectile fragmentation reactions.

The most commonly used is the Universal Parameterization [15]. It should be noted that during the calculation of partial cross sections, the assumed momentum distribution implicitly governs the LISE⁺⁺-predicted transmission and yield. As a result, partial cross sections carry an inherent model dependence³⁶. One of the explicit goals of total cross section analysis is to assess whether the chosen momentum distribution model reproduces the experimentally measured yields.

For most fragments, the momentum slits are initially assumed to select the central region of the momentum distribution, which can often be approximated as Gaussian. During dedicated momentum scans, the tails of the distribution are intentionally sampled to study deviations from this simple behavior. This process is shown in Figure 53, where cross sections were measured at different momenta and a distribution has been fit to the data points. The vertical error bars are the error in the cross section measurement, which have been the point of discussion throughout this manual. The horizontal error bars represent the width of the momentum acceptance of the spectrometer.

³⁶Advanced users may indeed vary the momentum distribution model or modify the parameters of the Universal Parameterization during partial cross section analysis.

In practice, individual partial cross section measurements are fit using a dedicated Excel workbook to reconstruct the shape of the momentum distribution. The fitting procedure incorporates the uncertainties deduced during partial cross section analysis. This reinforces an important experimental consideration: only fragments with sufficiently high transmission are suitable candidates for total cross section determination. Fragments with very small transmission are extremely sensitive to small changes in the simulated optical conditions and therefore tend to carry large uncertainties. A reliable total cross section is one that minimizes these propagated errors.

During the fitting process, it is recommended to fix either the mean or the width of the distribution using systematics from established models – such as those of Morrissey [36], Goldhaber [37], or the Universal Parameterization – and to fit the remaining parameters. This approach imposes physically motivated constraints and enables direct, quantitative comparisons between experimentally extracted distributions and model predictions.

Row Labels	Sum of dv	Sum of index	Sum of Runt	Sum of Y
0.9685	4.29E-03	22540	2	1
0.9804	9.22E-03	22540	3	151
0.9973	1.65E-02	42540	4	1820
1.0143	2.44E-02	52540	5	868
1.0309	3.55E-02	62540	6	61
1.0626	4.92E-02	72540	7	1

FLAG	w0	dw0	dCS/dw0	dCS/dE	dCS/dA	pclean	MemBrio	MemC	dMemC	w0	dFM_mid	CS	dE	dA	dMemCa	cs_w0_cal	df	COEF	angle	area	USE test	ACS+CS	
TRUE	0.9685	0.0043	1.63E-04	1.04E-04	2.86E-04	4.43E+04	3.24E+04	3.38E+04	-1374	0.96851	1.90E+02	3.68E-09	2.35E-09	6.46E-09	-374	1.38E-04	4.1E+02	1.69E+01	1	97.3	0.35	3.67E-03	175%
TRUE	0.9804	0.0092	1.11E-03	1.44E-04	1.47E-04	4.32E+04	3.24E+04	3.32E+04	-839	0.98043	3.99E+02	2.56E-08	3.34E-09	3.40E-09	-839	1.08E-03	4.6E+00	1.21E+01	1	97.1	1.70	4.28E-03	13%
TRUE	0.9973	0.0165	1.64E-03	2.05E-04	2.10E-04	4.19E+04	3.24E+04	3.25E+04	-115	0.99726	6.92E+02	3.92E-08	4.90E-09	5.00E-09	-115	1.74E-03	5.3E+01	1.45E+01	1	96.8	1.98	4.37E-03	13%
TRUE	1.0143	0.0244	6.09E-04	5.13E-05	5.39E-05	4.06E+04	3.24E+04	3.18E+04	585	1.01434	9.91E+02	1.50E-08	1.26E-09	1.33E-09	585	5.88E-04	8.4E+00	7.34E+00	1	96.7	0.46	5.86E-03	9%
TRUE	1.0309	0.0355	1.54E-05	2.89E-06	2.91E-06	3.90E+04	3.24E+04	3.09E+04	1457	1.03090	1.38E+03	3.96E-10	7.38E-11	7.47E-11	1457	2.55E-05	7.2E+02	1.71E+00	1	96.3	0.01	3.85E-03	19%
TRUE	1.0626	0.0492	1.43E-07	9.91E-08	2.53E-07	3.73E+04	3.24E+04	3.00E+04	2394	1.06258	1.84E+03	3.92E-12	2.66E-12	6.78E-12	2394	1.17E-07	1.4E+01	5.03E-01	1	100.0	0.00	4.45E-03	177%
FALSE	1.2000	0.0000	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A

fit	calc	err	%	use sqrt 0	MODEL	save	coef	max	min	meanP	sigmaP	areaP	gamma	gmr_gamma	sig_mor	sig_gold	sig_gold_gamma	sig_gold
0.9928	0.992	8.11E-04	0.1%	0.1%	0.9928	0.992	40060.1	2394	-1374	-289.1	359.5	4.9E-08	1.140	87.2	97.3	97.3	97.3	97.3
0.0090	0.010	-8.34E-04	-6.6%	-6.6%	0.0087	0.010		0.969	0.969	359.5	359.5	4.9E-08						
1.91E-03	1.74E-03	1.73E-04	10.0%	10.0%	1.92E-03	1.74E-03		1.063	0.969	-289.1	359.5	4.9E-08	1.140	87.2	97.3	97.3	97.3	97.3
4.30E-05	4.21E-05	-5.13E-04	-92.3%	-92.3%	5.56E-04	4.30E-05		1.063	0.969	-289.1	359.5	4.9E-08	1.140	87.2	97.3	97.3	97.3	97.3
13.54	13.5							1.063	0.969	-289.1	359.5	4.9E-08	1.140	87.2	97.3	97.3	97.3	97.3

Figure 54: Excel workbook used to fit the momentum distribution from the partial cross section measurements.

A note should be made about the target thickness during this analysis: just as each partial cross section depends on the areal density of the target, so too does the total cross section. Changing the target thickness inherently alters the shape of the momentum distribution through energy-loss, straggling, and reaction-depth effects. As such, a target thickness scan should be treated as a separate phase space variable along which production cross sections may be measured.

Calculating the momentum distribution and total cross section is related to the discussion in section 8 about quantifying the reach of FRIB and the feasibility of certain experiments. In Figure 53, the measured partial cross sections do not align with the peak of the distribution. Empirically determining the location of the peak helps inform the experimental settings which should be chosen to maximize the production of the isotope.

FAQ

- **Why does transmission uncertainty dominate the cross section uncertainty?**

In nearly all fragmentation-based cross section measurements, the uncertainty associated with the transmission ξ dominates the total uncertainty budget. This is not an

artifact of poor analysis practice; rather, it is a direct consequence of both the physics of fragment production and the practical limitations of modeling particle transport through a complex beamline.

First, the transmission is not a directly measured quantity. Unlike the yield Y or the beam integral B_{int} , which are derived from experimentally recorded signals, ξ must be inferred from simulations. These simulations rely on an idealized description of the experimental setup, including material thicknesses, reaction models, charge-state distributions, and ion-optical properties. Even when carefully tuned, such models cannot fully capture the true, event-by-event complexity of fragment transport from the target to the final focal plane.

Second, transmission is intrinsically sensitive to many coupled parameters. Small changes in target thickness, wedge thickness, material defects, slit positions, or optical settings can alter angular acceptance, momentum acceptance, or charge-state survival probabilities. While each parameter may be known to within a few percent (or better), their combined effect on transmission is often nonlinear. As a result, even modest, physically reasonable parameter variations can lead to percent-level or larger changes in ξ for individual fragments.

Third, transmission effects are fragment dependent. Fragments near the center of the acceptance typically exhibit relatively stable transmission, while fragments near acceptance edges – such as those with extreme A/Z , large momentum offsets, or unusual charge-state populations – can show strong sensitivity to higher-order optics, material interactions, and secondary reactions. Consequently, the transmission uncertainty varies from isotope to isotope and must be evaluated individually, further increasing its contribution to the overall error.

In contrast, the other terms in the cross section equation are generally better constrained. Statistical uncertainties on the yield Y decrease rapidly with increasing counts and are often subdominant except for the rarest fragments. The target thickness t is usually measured with high precision, and its uncertainty is bounded by manufacturing tolerances and direct metrology. Similarly, the beam integral B_{int} benefits from long integration times and redundancy in power measurements, resulting in relatively small fractional uncertainties.

Finally, because the transmission enters multiplicatively in the denominator of Equation (5), any fractional uncertainty in ξ directly propagates to the same fractional uncertainty in the cross section. When $\Delta\xi/\xi$ is at the 5–15% level, it will exceed the combined contribution of all other uncertainty sources.

For these reasons, the transmission uncertainty sets the ultimate precision limit of most cross section measurements. A careful, physically motivated transmission analysis, coupled with conservative but reasonable uncertainty estimates, is therefore the most critical component of a credible cross section result.

- **Kinematic corrections don't seem to enhance the A/q resolution. What's going on?**

The PPACs are very sensitive to the rate at which particles pass through the gaseous volume. When too many traverse at once, the ability of the electronics to distinguish one cascade from another is hampered. This is also related directly to dead time. The degradation of PPAC signals is usually correlated to a large number of light- Z fragments which do not register well in the detector and can clog the electronic readout. If kinematic corrections are not improving the A/q resolution, then retuning the beam to remove the light contaminants may be the answer.

- **Why does unit slope of reconstructed momentum between optical segments indicate good reconstruction?**

References

- [1] M. Hausmann et al., “Design of the advanced rare isotope separator ARIS at FRIB,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **317**, 349–353 (2013).
- [2] M. Portillo et al., “Commissioning of the advanced rare isotope separator ARIS at FRIB,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **540**, 151–157 (2023).
- [3] S. Di Carlo and M. Cortesi, “Parallel-plate avalanche counters for heavy-ion beam tracking: history and mysteries,” *Phys. Rev. Accel. Beams* **27**, 044801 (2024).
- [4] H. Kumagai et al., “Development of parallel plate avalanche counter (PPAC) for BigRIPS fragment separator,” *Nuclear Instruments and Methods in Physics Research Section B* **317**, 97–107 (2013).
- [5] O. B. Tarasov et al., “Discovery of new isotopes in the fragmentation of ^{82}Se and insights into their production,” *Phys. Rev. C* **112**, 034604 (2025).
- [6] M. Bowry et al., “Abrasion-fission reactions at intermediate energies,” *Phys. Rev. C* **108**, 034604 (2023).
- [7] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Phys. Rev. D* **57**, 3873–3889 (1998).
- [8] O. Tarasov and D. Bazin, “Development of the program LISE: application to fission–evaporation,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **204**, 174–178 (2003).
- [9] O. Tarasov and D. Bazin, “LISE++: radioactive beam production with in-flight separators,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **266**, 4657–4664 (2008).
- [10] O. Tarasov and D. Bazin, “LISE++: exotic beam production with fragment separators and their design,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **376**, 185–187 (2016).

- [11] O. Tarasov et al., “LISE cute++, the latest generation of the LISE ++ package, to simulate rare isotope production with fragment-separators,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **541**, 4–7 (2023).
- [12] O. B. Tarasov et al., “Production cross sections from ^{82}Se fragmentation as indications of shell effects in neutron-rich isotopes close to the drip-line,” *Phys. Rev. C* **87**, 054612 (2013).
- [13] M. Mocko et al., “Projectile fragmentation of ^{40}Ca , ^{48}Ca , ^{58}Ni , and ^{64}Ni at 140 MeV/nucleon,” *Phys. Rev. C* **74**, 054612 (2006).
- [14] O. Tarasov et al., “A new approach to measure momentum distributions and production cross-sections of neutron rich nuclei using fragment separators,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **620**, 578–584 (2010).
- [15] O. B. Tarasov, “Analysis of momentum distributions of projectile fragmentation products,” *Nuclear Physics A* **734**, 536–540 (2004).
- [16] J. D. Bowman, W. J. Swiatecki, and C. F. Tsang, *Abrasion and ablation of heavy ions*, tech. rep. (Lawrence Berkeley Laboratory, 1973).
- [17] L. F. Oliveira, R. Donangelo, and J. O. Rasmussen, “Abrasion-ablation calculations of large fragment yields from relativistic heavy ion reactions,” *Phys. Rev. C* **19**, 826–833 (1979).
- [18] J. W. Wilson, L. W. Townsend, and F. Badavi, “A semiempirical nuclear fragmentation model,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **18**, 225–231 (1986).
- [19] M. Wang, W. Huang, F. Kondev, and S. Naimi, “The AME 2020 atomic mass evaluation (ii). tables, graphs and references,” *Chinese Physics C* **45**, 536–540 (2021).
- [20] S. Goriely, S. Hilaire, M. Girod, and S. Péru, “First gogny-hartree-fock-bogoliubov nuclear mass model,” *Phys. Rev. Lett.* **102**, 242501 (2009).
- [21] J.-J. Gaimard and K.-H. Schmidt, “A reexamination of the abrasion-ablation model for the description of the nuclear fragmentation reaction,” *Nuclear Physics A* **531**, 709–745 (1991).
- [22] L. Audirac et al., “Evaporation-cost dependence in heavy-ion fragmentation,” *Phys. Rev. C* **88**, 041602 (2013).
- [23] K. Sümmerer, “Improved empirical parametrization of fragmentation cross sections,” *Phys. Rev. C* **86**, 014601 (2012).
- [24] D. S. Ahn et al., “Location of the neutron dripline at fluorine and neon,” *Phys. Rev. Lett.* **123**, 212501 (2019).
- [25] A. Gade and T. Glasmacher, “In-beam nuclear spectroscopy of bound states with fast exotic ion beams,” *Progress in Particle and Nuclear Physics* **60**, 161–224 (2008).
- [26] B. A. Brown et al., “Motivations for early high-profile FRIB experiments,” *Journal of Physics G: Nuclear and Particle Physics* **52**, 10.1088/1361-6471/adb449 (2025).

- [27] K. L. Brown, *A first- and second-order matrix theory for the design of beam transport systems and charged particle spectrometers*, Stanford Linear Accelerator Center (1982).
- [28] H. Wollnik, *Optics of charged particles* (Academic Press, 2021).
- [29] D. C. Carey, *The optics of charged particle beams* (Harwood Academic Publishers, 1987).
- [30] L. Bandura et al., “Fragment separator momentum compression schemes,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **645**, 182–186 (2011).
- [31] V. Maradia et al., “Demonstration of momentum cooling to enhance the potential of cancer treatment with proton therapy,” *Nature Physics* **19**, 10.1038/s41567-023-02115-2 (2023).
- [32] D. J. Morrissey and B. M. Sherrill, “In-flight separation of projectile fragments,” in *The euroschool lectures on physics with exotic beams*, vol. i. (2004).
- [33] M. Berz et al., “Reconstructive correction of aberrations in nuclear particle spectrographs,” *Phys. Rev. C* **47**, 537–544 (1993).
- [34] N. Fukuda et al., “Identification and separation of radioactive isotope beams by the BigRIPS separator at the RIKEN RI beam factory,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **317**, 323–332 (2013).
- [35] O. B. Tarasov et al., “Production of very neutron-rich nuclei with a ^{76}Ge beam,” *Phys. Rev. C* **80**, 034609 (2009).
- [36] D. J. Morrissey, “Systematics of momentum distributions from reactions with relativistic ions,” *Phys. Rev. C* **39**, 460–470 (1989).
- [37] A. Goldhaber, “Statistical models of fragmentation processes,” *Physics Letters B* **53**, 306–308 (1974).

Suggested Further Reading

The following list of references are recommended reading which supplements and augments the content of this manual. Note that this is in no way meant to be complete; if one desires a complete summary of low energy nuclear physics research, begin with L. Meitner and O. Hahn, “Products of the Fission of the Uranium Nucleus,” *Nature* **143**, 471–472 (1939) and proceed chronologically!

- [38] A. K. Anthony, C. Y. Niu, R. S. Wang, J. Wieske, K. W. Brown, Z. Chajecski, W. G. Lynch, Y. Ayyad, J. Barney, D. Bazin, and S. Beceiro-Novo, “Beam particle identification and tagging of incompletely stripped heavy beams with heist,” *Review of Scientific Instruments* **93**, 013306 (2022).
- [39] D. Bazin et al., “Radio frequency fragment separator at nscl,” *Nuclear Instruments and Methods in Physics Research Section A* **606**, 314–323 (2009).

- [40] J. R. Cummings et al., “Determination of the cross sections for the production of fragments from relativistic nucleus-nucleus interactions. i. measurements,” *Physical Review C* **42**, 2508 (1990).
- [41] M. de Jong et al., “Fragmentation cross sections of relativistic 208Pb projectiles,” *Nuclear Physics A* **628**, 479–504 (1998).
- [42] W. A. Friedman, “Heavy ion projectile fragmentation: a reexamination,” *Physical Review C* **27**, 569 (1983).
- [43] H. Geissel et al., “Ion-optical layout of a powerful next-generation pre-separator for in-flight separation of relativistic rare isotopes,” *Nuclear Instruments and Methods in Physics Research Section B* **247**, 1–15 (2006).
- [44] H. Geissel et al., “Ions penetrating through ion-optical systems and matter — non-liououvillian phase-space modelling,” *Nuclear Instruments and Methods in Physics Research Section A* **282**, 247–253 (1989).
- [45] M. Giacomelli et al., “Projectilelike fragment emission angles in fragmentation reactions of light heavy ions in the energy region 200MeV/nucleon: modeling and simulations,” *Physical Review C* **69**, 064601 (2004).
- [46] A. S. Goldhaber, “Statistical models of fragmentation processes,” *Physics Letters B* **53**, 306–308 (1974).
- [47] R. Grzywacz et al., “Identification of μ s-isomers produced in the fragmentation of a 112Sn beam,” *Physics Letters B* **355**, 439–444 (1995).
- [48] K. Haak et al., “Production and discovery of neutron-rich isotopes by fragmentation of 198Pt,” *Physical Review C* **108**, 034608 (2023).
- [49] C.-W. Ma et al., “Nuclear fragments in projectile fragmentation reactions,” *Progress in Particle and Nuclear Physics* **121**, 103911 (2021).
- [50] P. N. Ostroumov et al., “Acceleration of uranium beam to record power of 10.4 kW and observation of new isotopes at facility for rare isotope beams,” *Physical Review Accelerators and Beams* **27**, 060101 (2024).
- [51] R. Pfaff et al., “Projectilelike fragment momentum distributions from 86Kr + Al at 70 MeV/nucleon,” *Physical Review C* **51**, 1348 (1995).
- [52] H. Suzuki et al., “Production cross section measurements of radioactive isotopes by BigRIPS separator at RIKEN RI beam factory,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **317**, 756–768 (2013).
- [53] H. Suzuki et al., “Discovery of 72Rb: a nuclear sandbank beyond the proton drip line,” *Physical Review Letters* **119**, 192503 (2017).
- [54] O. B. Tarasov et al., “New isotope 44Si and systematics of the production cross sections of the most neutron-rich nuclei,” *Physical Review C* **75**, 064613 (2007).
- [55] H.-L. Wei et al., “Cross-section prediction for isotopes near neutron drip line in 70, 80Zn projectile fragmentation reactions,” *Chinese Physics C* **43**, 074103 (2019).

- [56] J. W. Wilson et al., “A semiempirical nuclear fragmentation model,” *Nuclear Instruments and Methods in Physics Research Section B* **18**, 95–134 (1987).
- [57] D. J. Morrissey, W. R. Marsh, R. J. Otto, W. Loveland, and G. T. Seaborg, “Target residue mass and charge distributions in relativistic heavy ion reactions,” *Phys. Rev. C* **18**, 1267–1274 (1978).
- [58] L. W. Townsend et al., *An abrasion-ablation model description of galactic heavy-ion fragmentation*, tech. rep. (NASA, 1984).

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics and used resources of the Facility for Rare Isotope Beams (FRIB) Operations, which is a DOE Office of Science User Facility under Award No. DE-SC0023633, and by the US National Science Foundation under Grants No. PHY-20-12040 and 23-10078 “Windows on the Universe: Open Quantum Systems in Atomic Nuclei at FRIB”.

I would like to thank Brad Sherrill and Oleg Tarasov for guiding me through my first cross section analysis, which enabled me to commit to paper the process from start to finish. I am also indebted to them for all of their support throughout my graduate career.

I would also like to thank the ARIS Group and Rare Isotope Research Group with whom I’ve worked and learned much from. In particular, I would like to acknowledge Isaiah Richardson for his help in writing Appendix F.

Finally, I would like to thank ChatGPT, which served as the primary copy editor for this absurdly long manual. I am certain that the algorithm despises me.