FRIB ECR Ion Source: An Operations Guide

By J. W. Stetson, stetson@msu.edu, revision 24-Mar-2021

Purpose

The primary goal of this document is to provide guidance in the operation of the ion source(s) supplying beam into the low energy beam lines for subsequent acceleration in the FRIB LINAC. At present, this is done solely with the 2nd generation, 14 GHz *Electron Cyclotron Resonance Ion Source* (ECRIS), called ARTEMIS-B. For this reason, plus a measure of simplicity, specific hardware discussion will be of this particular source.

The intended audience are those tasked with delivering stable, high-quality beams on a routine basis who may or may not be ion source specialists. It is not intended as a formal publication to an expert audience, nor solely as an A-B-C checklist for any random technical person. As such, the written tone will take the general form of a control room discussion. Certain inaccuracies and over-generalities are hence, in the service of clarity and brevity, contained within. Familiarity with engineering-level physics is assumed, but considerable attempt has been made to introduce terms that may be unfamiliar beyond that level or used in peculiar ways within the ECRIS specialist community. (And the ECRIS community is widely considered to be, if nothing else, peculiar.)

While, mechanically and electrically, ECR sources are simple and have few adjustable parameters, it is a mistake to assume that their operation is correspondingly simple. Seldom is it the case where simply setting these parameters precisely to values successfully used on a previous Tuesday will give the exact same results on a subsequent Friday. The previous values certainly serve as a useful starting point, but immediate results can vary considerably in practice. The operational need then becomes to change the initial turn-on beam condition into a good one in a reasonably efficient way. Another such need is to compensate for inevitable tune shifts which develop over time. A major tool in doing both is to have a useful mental model of what is happening inside the plasma chamber and during the extraction process. (Useful, not necessarily correct in every detail). Hopefully this guide will provide those newly exposed to this technology with such a framework.

A Base ARTEMIS Model – Mechanical and Magnetic

To begin, here is a 3-D sketch of the ARTEMIS ion source and first segment of the *low-energy beam transport* (LEBT) system:

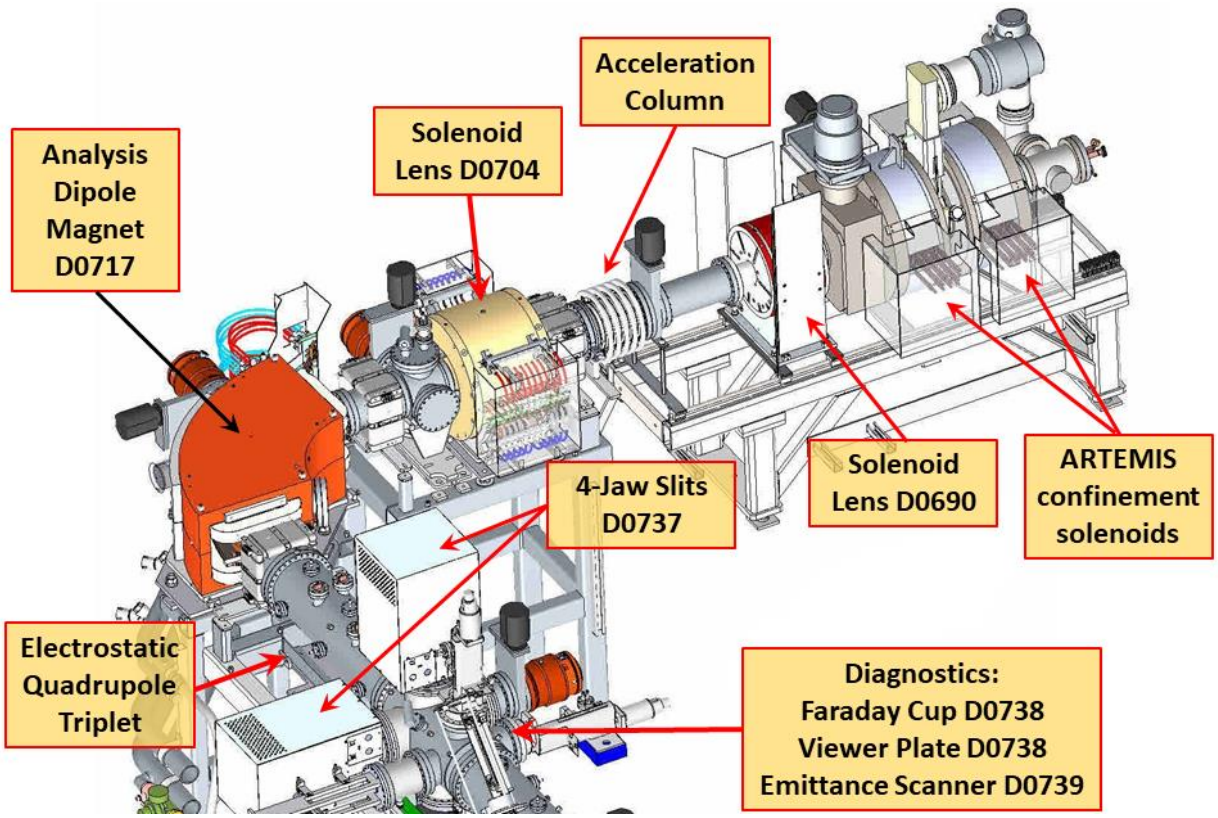

All of these labelled elements will be discussed in detail at some point in the subsequent text. Not shown is the high-voltage platform, electrically-isolated from ground beneath by supporting insulators and from the downstream beam line by the acceleration column. ARTEMIS is itself insulated from the platform.

At the core of ARTEMIS is a high vacuum (below about 2 x 10-7 Torr) cylindrical volume 75 mm, 3”, in diameter and about 300 mm long, ~12”, (working volume), which serves as the *plasma chamber*. All surfaces forming this container are aluminum. (Other materials have been tried but work either poorly or not at all). Here is a view looking inside and downstream, that is, in the direction of the extracted beam path:

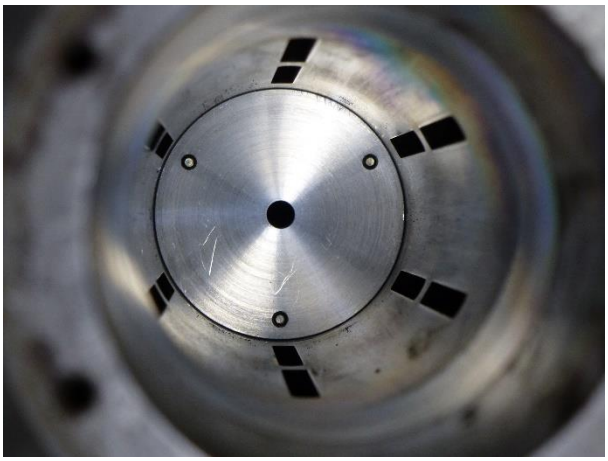

There are 6 lines of radial slot openings machined into this cylinder whose purpose is to provide paths for both injection of gas to be ionized and for vacuum pumping.

At the downstream, “extraction” end is a replaceable aluminum insert with a central (normally) 12 mm diameter hole in it, named *plasma aperture*, flat on the plasma side and conical on the exit side. Ions for subsequent acceleration pass through this hole. This ‘burn’ pattern caused by ions and electrons impacting the surface is typical.

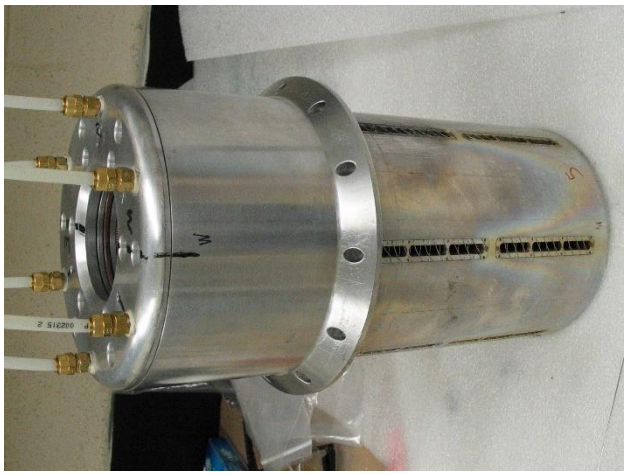

Inside of the complete *plasma chamber assembly* (shown above), exactly offset from these slots and between the inner surface and the outer surface, are provisions for mounting, and firmly holding against strong magnetic forces, a set of long, permanent bar magnets arranged to provide a sextupole (hexapole) field.

On the upstream end of the plasma chamber, a complex *injection assembly* is bolted in place.

A view looking in the upstream direction from the plasma chamber center shows the aluminum *RF baffle* (or *plasma baffle*) which is electrically connected to the chamber and closes the end opposite the plasma aperture. Note the star-shaped marks caused by ion and electrons that have escaped confinement.

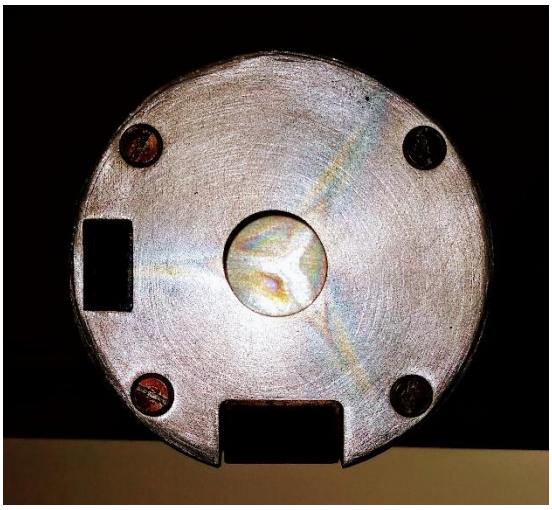

There are several different types of baffles and corresponding injection assemblies. For “solid feed” beams, that is when the material to be ionized is solid rather than gaseous, an oven, or something similar to “boil” the solid sample, special features are required. All types will have means for providing microwaves into the chamber. The pictured example shown above has two rectangular waveguide ports, the larger one for 14 GHz, the smaller, for 18 GHz. Also common is an axially-mounted ~14 mm diameter aluminum disc electrically isolated from the plasma chamber which can be biased with an independent power supply, and is hence called a *bias disk*. A profile view makes this clearer.

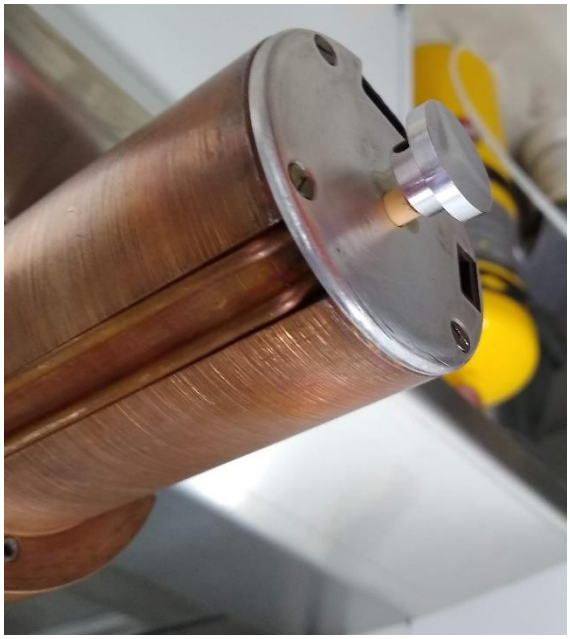

The plasma chamber assembly, all of the attached hardware (shaded yellow in the following figure) and connected power supplies are electrically insulated (floating) relative to the surrounding equipment on the ECRIS platform. This group of devices is attached to a HV power supply [FE_ISRC1:HVP_D0679] and generally biased to 15 - 20 kV over the platform potential. *ECR voltage*, *source bias voltage*, and *extraction voltage* are equivalent terms referring to this same thing. The insulation between the plasma chamber and the platform-potential iron of the solenoid bores is a somewhat fragile acrylic cylindrical sleeve inserted between them. (The sleeve is exposed to bremstraalung X-rays from the high-energy tail of hot electrons which creates cracks. These can allow leakage current and sparks from the chamber to ground which results in failure. The process is greatly sped up in the presence of moisture, either from condensation or outright leaks).

It's important to note that ARTEMIS itself, as well as extracted ions in the first ~3 m of beam line, does not "see" the platform potential of 0 – 120 kV relative to laboratory ground created by a separate power supply [FE_ISRC1:HVP_D0698].

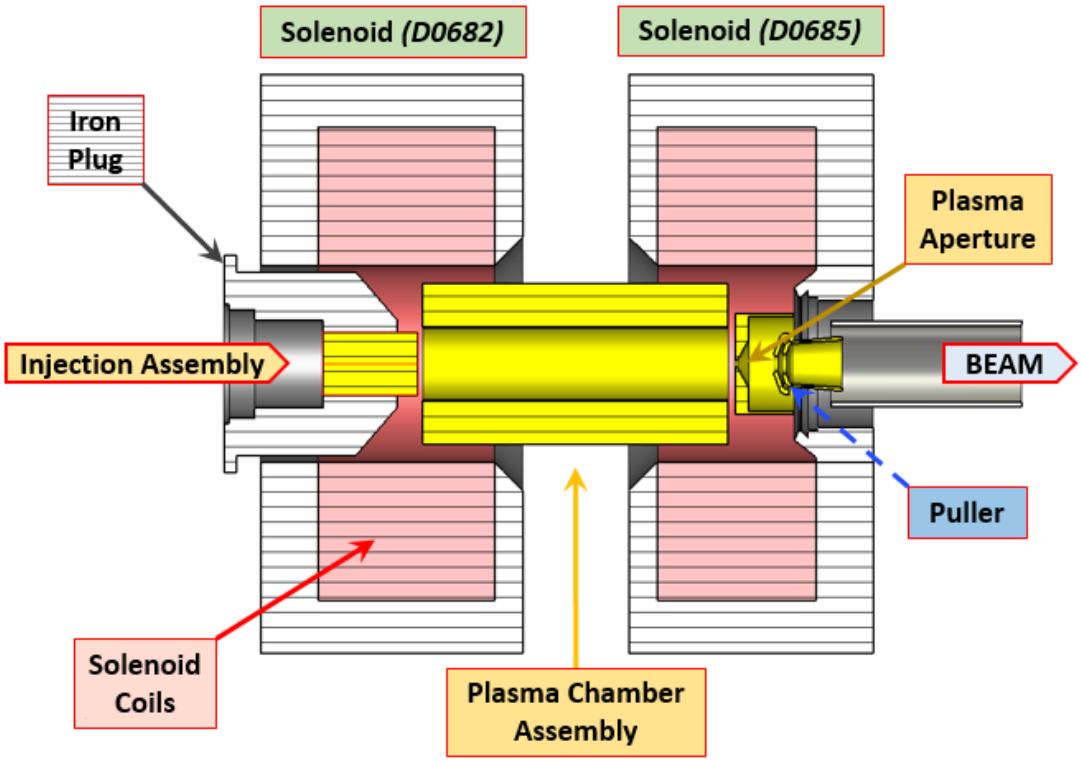

The two solenoids have their on-axis fields oriented in the same direction. In the case of ARTIMIS-B, these fields point downstream. The field shape and strength provide both the primary plasma confinement and a region near the center where the scaler field strength provides a resonance condition for free electrons with 14 GHz microwaves. Field strength peaks near the center of each solenoid. If the intent of this “magnetic bottle” configuration were to maximize internal plasma density, both peaks would be equally high. However, the intent is to provide a source for ions, so the downstream solenoid has a much weaker field than upstream, thus encouraging ion “leaks” in that direction. While the two solenoid power supplies are both run within 10% of the same current values, their resulting field peaks differ by about a factor of two. The injection-side solenoid *D0682* coils have more turns and its *iron plug* is more substantial than is the case for the extraction-side *D0685* solenoid. Both solenoids are built with less iron near the gap between them to provide a lower minimum field strength there compared to the peak fields.

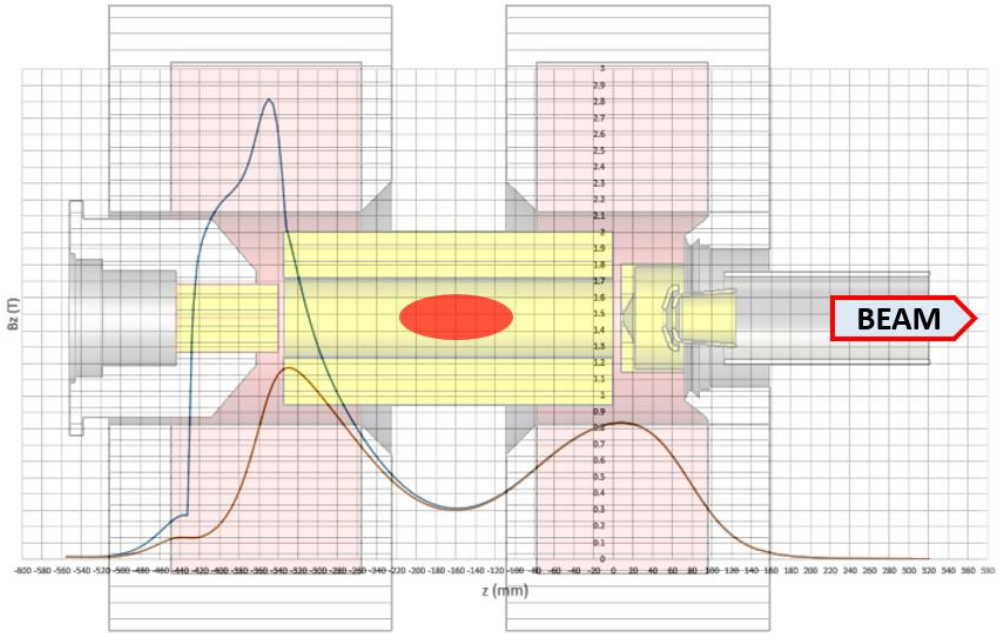

In the figure above, the strength of the downstream component of the solenoid-created magnetic field is shown relative to its on-axis longitudinal position. The plasma aperture and baffle/bias-disk are placed at or near field maxima.

The scaler magnetic field strength matches the 14 GHz free-electron resonance requirement at the surface of a 3-dimesional, approximately ellipsoidal shell shown in red above. The free-electron resonance frequency relative to magnetic field is an oddly simple straight line relationship with 28 GHz requiring a 1.0 Tesla field. Thus, 14 GHz requires 0.5 Tesla. (Choices of frequency for new ECR sources has been historically largely determined by the commercial availability of Klystron microwave transmitters used by the telecommunications and aerospace industry). The rising field at each end relative to the center creates a longitudinal trap or magnetic bottle for hot electrons when their transverse velocity component exceeds that of their longitudinal. Electrons and ions not meeting this condition are quickly lost on the plasma chamber ends. Those satisfying the condition are reflected back toward the center.

The large difference between the two very different injection-end field strengths shown is both illustrative and represents a situation unique to Artemis-B. The higher field strength is the designed “normal” value and is achieved by having an iron end piece on the injection assembly which closes the magnetic return circuit, thus enhancing the field seen within the chamber. The lower profile shows the situation with the iron plug replaced with one made of copper. Artemis-B seems to run well in both cases, but best case parameters are somewhat different. With the iron, current value of the extraction coil is usually greater than the injection coil. With the copper, this relationship tends to be reversed.

To create a rising-field trapping condition for electrons in the radial (transverse), a six-fold symmetric hexapole magnetic field is superposed upon this solenoidal main field. This is done with 12 permanent magnetic bars of NdFeB (N42SH) arranged in pairs forming a Halbach array, https://en.wikipedia.org/wiki/Halbach_array.

The field created is typical of any sextupole (but the downstream-looking field orientation of ARTEMIS-B, as mounted on the platform, is rotated 30 degrees counterclockwise with respect to this diagram, making it a skew sextupole):

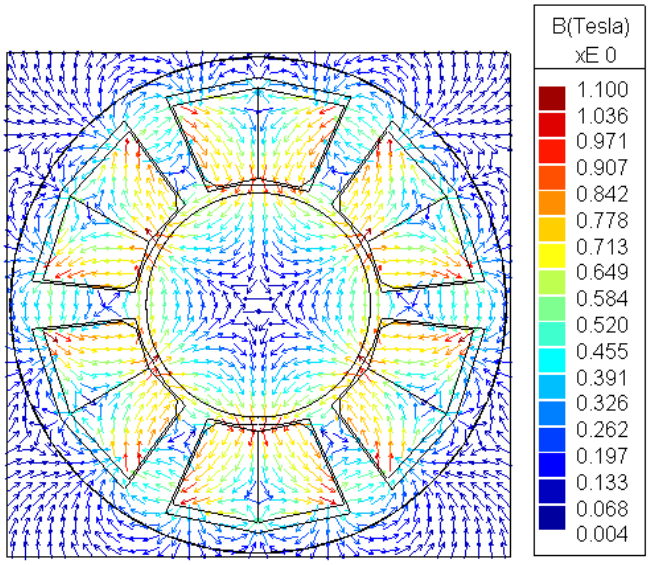

Starting from the axis, an outward-moving particle generally sees a rising field in every direction. However, ions/electrons moving toward the center of any pole are not doing so with a sufficient angle relative to the field lines to satisfy the mirror condition and thus can reach the chamber walls with ease. Hence, radial losses of energetic particles are not distributed evenly over the plasma chamber surface, but are concentrated at these six *loss lines*. This fact has a big influence on ECRIS design constraints.

While rare-earth permanent magnetics can create high fields in a compact size (see figure below). They have a critical weakness.

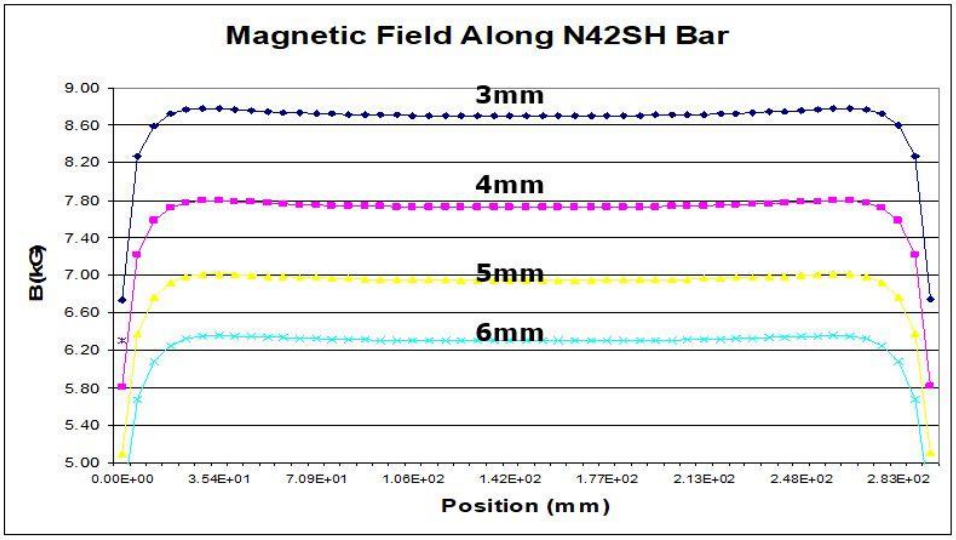

This shortcoming is a low critical temperature, that is, a temperature above which magnetization is lost. For N42SH, temperatures above 80 deg C will result in at least partial de-magnetization. Unfortunately, the radial loss lines concentrate heating from leaking plasma and are in exactly the wrong place: the magnet bars. Hence, sufficient water cooling becomes extremely important. Another complication is that iron rusts and rust is non-magnetic. Therefore, the bars must also be kept dry within water-tight sleeves.

Some construction details:

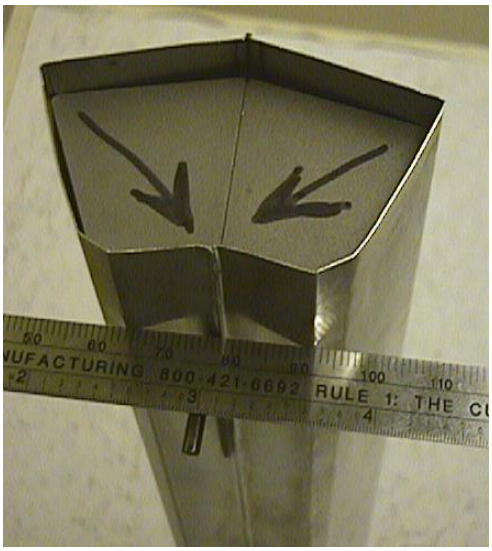

Since of course, permanent magnets are always fully-energized and they very much wish to move away from each other, they must be held in a robust structure. Insertion into this structure against strong non-linear forces is an engineering challenge and is also a definite hazard to those doing the assembly. Hence, any damage to the permanent magnets is very time-consuming to repair.

This concentration of power loss also represents a risk to the plasma chamber walls as well. Were the ~1kW of injected RF power lost over evenly over the entire inner surface, risk would be minimal. This, however, is not the case. A wall burn through would result in an ingress of water, disabling the source, and forcing a subsequent major repair of long-term duration.

So far in this exposition, the radial and axial fields have been discussed independently. They of course are not, and that has significant consequences. The vector sum of the two changes the 6-fold symmetry of the hexapole field into two 3-fold symmetries. One visualization of this can be seen in a slice 3 cm upstream of the plasma extraction aperture showing the radial component of total field strength. Another can be seen in a layered sketch of the resultant surface showing the condition in space where the ECR-resonant scaler field strength condition is satisfied. Both are shown below.

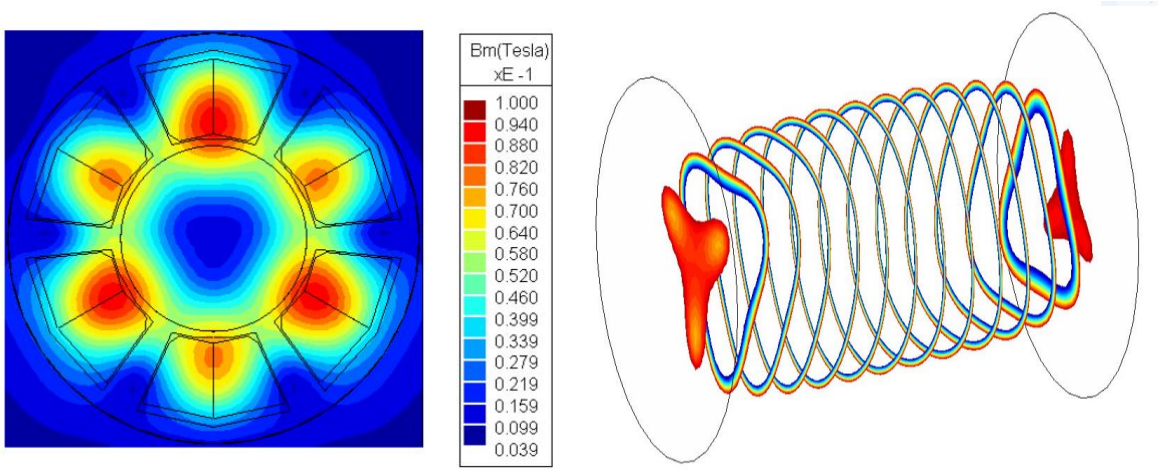

A Base ARTEMIS Model – RF

Energy to heat electrons in the source at their resonant condition is provided by a mechanism which can be thought of as depicted in this block sketch:

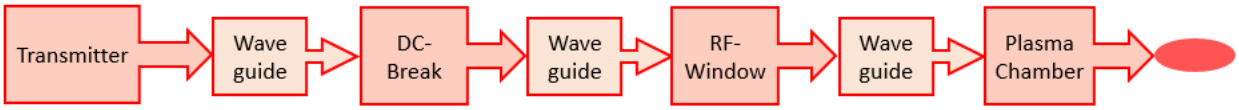

Microwaves are provided (primarily) by a *Klystron* RF transmitter located on the HV platform. (A brief description of Klystron function and some uses relative to accelerators can be found here: https://insights.globalspec.com/article/12732/what-is-a-klystron).

While both efficient and able to create high power (1 – 2 kW with the Artemis setup), Klystron-based amplifiers work only within a narrow region of their design frequency. It's been shown that dual-frequency heating can result in better source performance. That is, having one feed at the ECR resonant frequency and a second one at both lower frequency and lower power. (We've run 18 + 14 GHz with SUSI). However, it's also been shown that an ability to vary the lower "support" frequency over a wide range is of great advantage. Hence, a *TWT (Traveling Wave Tube)* amplifier and connecting waveguides have been added to the Artemis platform. While having a relatively low maximum output power of ~400 W, the frequency is adjustable over a range of ~10 – 14 GHz.

Generated microwaves are fed into the plasma chamber via a rectangular cross-sectioned copper waveguide, dimensioned to carry them with minimal losses. For 14 GHz transmission as in Artemis, we use a WR75 standard size of 0.750" x 0.357". For 18 GHz, it's WR62 with dimensions of 0.622" x 0.311", see http://www.rfcafe.com/references/electrical/waveguide-chart.htm for a reference. A more general overview can be found here: https://www.microwaves101.com/encyclopedias/waveguide-primer#components. Efficient microwave transmission is a messy subject.

Since the copper waveguide is electrically conducting, and the transmitter is at platform potential and the plasma chamber is 15-20kV above that, a piece called the *DC-break* must be inserted into the transmission line.

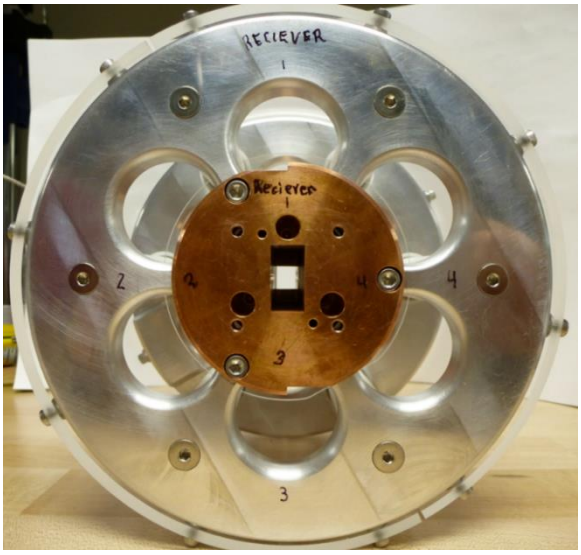

The incoming and outgoing waveguides bolt to these two large chunks of copper which essentially serve as transmitter and receiver antennas for passing microwaves across an insulating glass disk sandwiched between them. Dimensions inside these pieces are identical to the waveguide until reaching the ends

where tips form a carefully-sized opening intended to allow transfer of power from one side to the other with minimal loss across a non-conductive gap. However, this break into which the glass is slid, presents an opportunity for RF leakage which can be quite significant at high transmitter powers. Therefore, while in operation, the unit should be surrounded with a wire mesh for shielding. Care should be taken that this mesh doesn't cause a short between the source and platform ground.

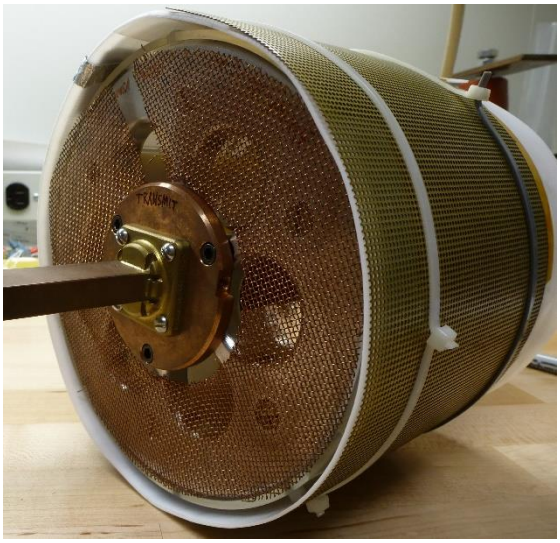

Unfortunately, while the waveguide passes power efficiently over a fairly wide frequency range, the DC-break does not, which can be seen below in a plot from an in-situ test done in 2017. Here, a shift of 150 MHz from the best-tuned frequency results in a factor of 10 increase in measured reflected power. Adjustment (filing) the size of the tip opening shifts where this minimum happens, but so far it has proven difficult to accomplish such adjustments due to discrepancies in bench and installed results.

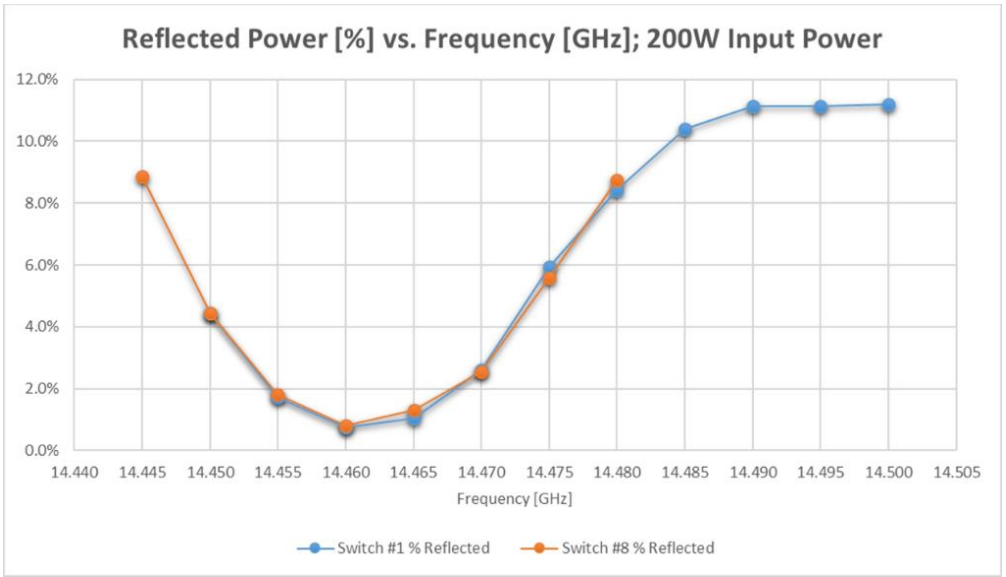

The ‘tune’ of the transmission line between transmitter and source depends on the behavior of the system as a unit, including the plasma chamber itself. While the nominal impedance of the line is 500 Ω, the plasma chamber is not impedance-matched to the incoming line and worse, the plasma itself presents a loading that is both unknown and varies according to conditions.

Partially as a consequence of this situation, the fundamental Klystron frequency will be made (slightly) adjustable, that is, fine-tune the frequency to the line rather than the line to the frequency. This can already be done to some degree by mechanical switches located within the transmitter assembly as shown (for Artemis-A) here:

Additionally, the present fixed-frequency crystal driving oscillator will be replaced in the near future with a remotely-controlled frequency synthesizer. This will allow fine-tuning for minimum line losses and matching to plasma loading.

Why do line losses matter? To first approximation, it doesn’t, unless power losses happen in bad places and the transition across the insulating glass is a bad place. While the tips and copper pieces can be water-cooled, the glass cannot and any loss here is concentrated within a small area. Yet while failure of this glass piece would interrupt operations, it would not be disastrous, per se. Losses at the next-in-line *RF-window* which caused heating to the point of failure, would be such a disaster.

The RF-window (https://www.cpii.com/docs/related/1/Pressure%20Windows.pdf) is a precision-made very thin piece of quartz (SiO_2) sized to match the waveguide dimensions. This piece, supported by a frame with an indium sealing gasket, separates the air in the upstream waveguide from downstream vacuum. While extremely hard, quartz is also brittle. Mechanical or thermal stress, caused by heating from power loss, could cause it to crack and allow a blast of air into the plasma chamber and following beam line. Recovery time would be considerable.

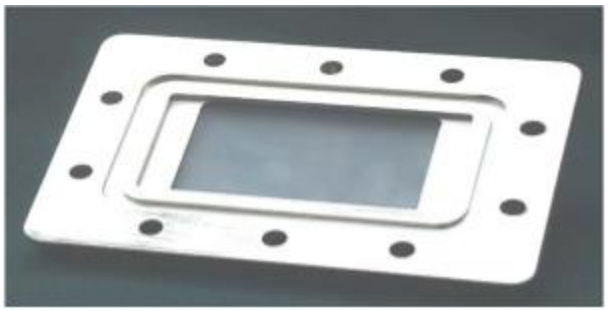

Power loss here depends of course on the total applied microwave power, but more importantly on the tune of the transmission line. Standing waves and constructive interference between forward and reflected waves can concentrate onto a small area of the window surface. Unfortunately, at present, the only diagnostics available are the forward and reverse power sensors located in the microwave transmitter and the source ion beam itself. We assume that a low ratio of measured reflected power to power transmitted is good, but it can also mean that the transmitted and/or reflected power is being absorbed somewhere in the line. Power taps are already installed on the waveguide segment between the window and source, but are not yet readable remotely. That information would be a useful addition in detecting hazardous power losses. A partial substitute is the beam itself. This chart shows why:

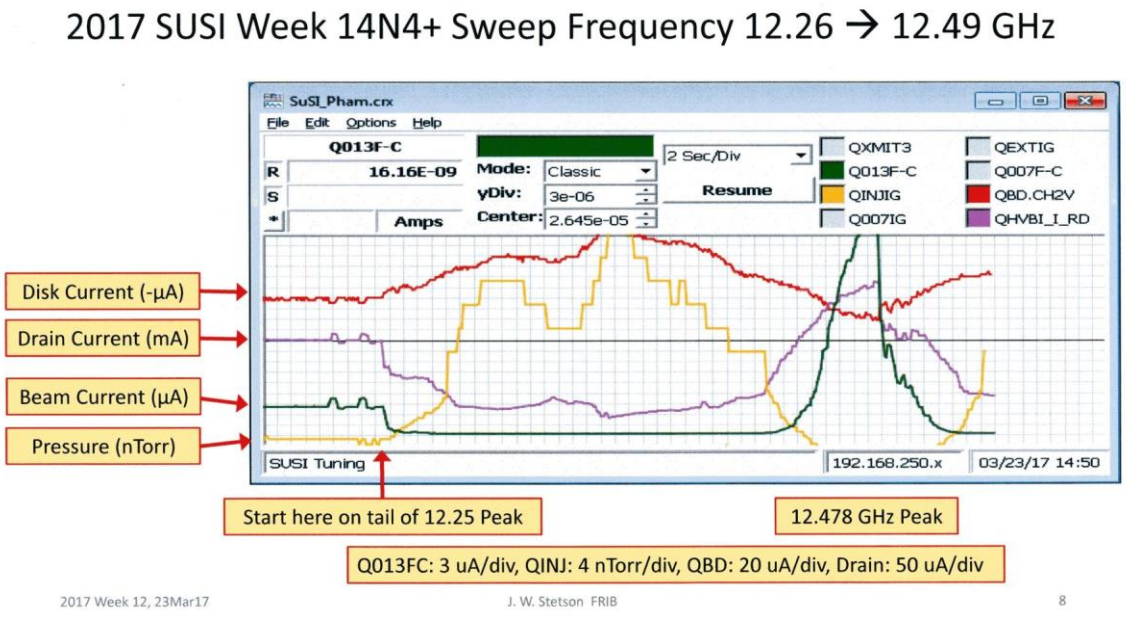

This was taken while using a wideband TWT microwave transmitter. Beam output varies with frequency, which was expected. However, the periodic sudden drops to near-zero were not expected, nor the accompanying rise of injection side pressure. It was later determined that the window used during this test was opaque to microwaves at such frequencies. In other words, the beam output dropped because incoming power was being absorbed by the window and not reaching the plasma. Not a good situation. Suddenly disappearing beam during any frequency tuning should be corrected immediately.

A Base ARTEMIS Model – Gas Feed

Gas for making plasma is introduced into the plasma chamber through the radial ports via an external manifold system and a pair of leak valves. The valves and associated control mechanism are attached to the plasma chamber in such a way that they are at the ~15 kV source potential when running. Connected to these valves are a "rat's nest" of plastic "Polyflo" lines leading to a stainless steel manifold "switchyard" routing gas flow from large tanks of high-purity helium, oxygen, nitrogen and argon sitting outside the lead shielding at platform potential. One should note however, that while chemically pure, these gasses are not isotopically pure. In particular, ^{18}O, ^{15}N and ^{36}Ar are present in their natural abundances and may be noticed as a component of the extracted beam. These oft-used gasses are cheap, even in a very pure form, so can be economically purchased in quantity and the volume lost because of long lines leading to the valves, therefore, doesn't matter.

Leak Valves

Large Gas Cylinders and Main Manifold

Secondary Manifold

In contrast, isotopically-enriched gasses (and solids) are very expensive. For example, a tiny 500 mL bottle of "cheap" ^{84}Kr costs ~$2,500 with ^{86}Kr and any form of xenon considerably more. In order to minimize loss of such expensive gas, a secondary manifold is located within the lead shielding as close as possible to the first leak valve, so that it can feed either from the remote big cheap tanks or from local little expensive cylinders.

In general, operators will not manipulate any part of the gas feed system other than the remotely controlled leak valve settings. Mistakes, which are easy to make, can be very costly so those tasks are left to the source specialists.

[Aside: Our most expensive source material to date is solid ^{48}Ca, which is also one for which there is much demand in the experimental program. At ~$300 per milligram, the cost of even a light oven load is over $50,000, so extreme care in preparation, loading and operation is warranted.]

A fundamental problem with feeding gas into the source is the transition down a dozen orders of magnitude in pressure and, at the same time, controlling precisely a tiny mass flow rate, estimated to be at most 1 x 10^{-3} cc/minute. The present valves are delicate and the present remote stepper motor drives and controllers on ARTEMIS, being decades old, are cantankerous – which is a bad combination. If driven too far out, source vacuum will go quickly into the Torr range, shutting everything down. If driven too far in, most often caused by incorrectly re-establishing the "0", fully-closed position, a thin sapphire diaphragm within the valve will break.

Adequate Physics Model – Inside the Chamber

Please read https://en.wikipedia.org/wiki/Plasma_(physics) for some basics. In terms introduced in the article, an ECR plasma is:

1. Magnetized (~ 0.5 T).
2. High Vacuum (~ 100 nTorr = 1 x 10-7 Torr).
3. Cold ($E_{(electron)}$ ~ 5 keV & $E_{(ion)}$ ~ 10 eV, T_e ~ 60 million deg K & T_i ~ 100,000 deg K). Not mentioned as such, but non-thermal applies: https://en.wikipedia.org/wiki/Nonthermal_plasma
4. Much richer in instabilities than mentioned, due mostly to the high magnetic field.

For an ECRIS, the plasma begins as a low-density, electrically-neutral gas. Yet, for short moments, free electrons, from cosmic ray interaction and quantum mechanical statistical fluctuations, exist. Microwaves applied from the transmitter interact with these electrons, causing mostly random motions. However, if the polarization is correct and the interaction occurs at the cyclotron resonant frequency for the scalar magnetic field at the particular location, those electrons rapidly (on millisecond time scales) gain energy. A nice graphic is shown in this article: https://en.wikipedia.org/wiki/Electron_cyclotron_resonance. (The gain in energy from this interaction is ultimately limited by the relativistic mass increase of the electrons, causing them to fall off-resonance. But, since the energy optimum for collisional ionization of neutral ions and further ionization into higher charge states is ~ 2 keV, this limitation is inconsequential.) Collisions with nearby neutral atoms set other electrons free to interact with the microwaves as well, resulting in a cascade effect. The gas, quite literally as in a fluorescent tube, lights up and heats up, driven by power absorbed from the microwave transmitter. In cases where the plasma doesn't form immediately, reflected power will read a much higher percentage of applied power than is normal. When it does light up, the reflected power will drop, showing that the ionization process is occurring. Since the goal is a steady supply of ions, as opposed to say ions supplied in a flash of thunderstorm lightning, a substantial number of balancing acts come into play. This time duration requirement implies overall charge neutrality, that is, the total number of free electrons will match the total charge represented by the resulting ions. *Electron density*, with the "free" implicit, puts a fundamental cap on ion source output and hence is considered and discussed as a primary parameter in dealing with ion sources. To produce lots of ions requires a lots of free electrons.

So, with the source lit and power being added, what factors limit electron density?

Starting from a too-low pressure and too-low power condition, adding feed gas creates more opportunities for ionizations, but that process absorbs energy so more microwave power is also

required. At the same time ions are being created, they are also being lost. More added gas increases the probability that an ion will re-combine with an electron. [Note: In many discussions and publications when the topic of *recombination* arises, the concept of *mean free path*, that is, the average distance traveled by one particle before it collides with another, is also discussed. (See https://en.wikipedia.org/wiki/Mean_free_path for some basic information.) While the concept can be useful, analogies to an ideal gas can be misleading. Particles within an ideal gas are assumed to move freely in 3 dimensions. This assumption is incorrect in the case of magnetized plasmas, where charged particle motion is limited by the external field, hence restricting dimensionality to less than three. Also, unlike an ideal gas, charged particles within an ECRIS are not uniformly distributed.]

A limit to adding gas for "more beam" is also time. Maximum plasma density can happen at quite high pressures (~100 mTorr) if one only wishes to create singly-charged ions. Ionization to higher charge states however, proceeds in a step-wise fashion, +1 -> +2 -> +3 and so on. Each step up in charge is a statistical process which takes time and each step up creates another rate equation with the probability of a recombination to step down. Also, as the positive charge of an ion increases, its affinity to nearby electrons also increases and the chances for capture increase. Likewise, with increasing pressure, the chances for ion-ion collisions resulting in loss of the desired charge condition increase; the ion must live long enough after creation to makes its way out of the plasma aperture. Ion lifetime is most often referred to as *confinement time*, typically around a few 10s of milliseconds. Optimum pressure for extraction of high charge-state compared to singly charged ions as a consequence of time-related processes, therefore drops by about 6 orders of magnitude to ~100 nTorr, that is, from 10^{-1} to 10^{-7} Torr.

[It should be noted that the source vacuum as measured by ion gauges or cold-cathode gauges becomes somewhat ambiguous when the source is running. The gauges measure atoms, not ions! For clarity of meaning, these readings are often referred to as *neutral gas pressure*. It's not unusual to see, with a well-conditioned source, an indicated pressure drop when RF power is added, implying that a substantial portion of gas is being ionized, getting trapped by the magnetic field, and not making it to the gauge to be measured.]

Having a long confinement time also requires having good containment. Limiting losses to the walls is a task for the magnetic field. If the initial velocity of a just-created free electron or ion aligns exactly with the magnetic field line at that point, odds are high that it will be lost. If however, this velocity has an angle to the field, called the *mixing angle*, the particle will spiral around that field line with a net helical motion. For electrons, the *gyroradius*, aka *Larmor radius*, of such motion is very small, usually considerably less than 1 mm. [For an overview, see https://en.wikipedia.org/wiki/Gyroradius]. Ions are much more massive, but, since they are not heated by the microwaves, also much colder (slower), with the two effects resulting in a gyroradius of up to a centimeter or so, depending on M/Q, the estimate of ion energy and the local magnetic field strength.

Electrons, therefore, are essentially stuck to the magnetic field lines and the basic axial magnetic field of a solenoid provides good confinement in the radial direction, but none axially. A partial fix is provided by the rising field strength at each end of the plasma chamber forming a *magnetic bottle*. If the particle's component of velocity perpendicular to the magnetic field is greater than its aligned component, the particle is reflected back toward the lower field region. If not, it is not contained and it exits into a region referred to as the *loss cone*, striking an end of the plasma chamber. Perhaps

surprisingly, because of much larger mass and lower velocity, ion confinement is achieved with equal effect. For both, the effect is proportional to the ratio of highest to lowest field in the structure, called the *mirror ratio*. Please see https://en.wikipedia.org/wiki/Magnetic_mirror for a more thorough description. [An aside: the references to plasma fusion devices within this article suggest a connection with ECR ion sources that is accurate; the two are closely related. However, fusion requires hot ions, sources prefer cold ions, so the two differ fundamentally in this way. The ECRIS was a spin-off from plasma fusion experiments in France done in the late 60's and early 70's driven by Richard Geller, initially using cast-off components from other projects.]

Mirror confinement enhances electron density by reducing electron losses, but also contains ions. Since the purpose of a source is to produce ions, the extraction side field is designed to be lower than the injection side field. Extracted ion beam can be considered as an intentional leak from a magnetic bottle.

Because the Larmor radius of ions is rather large compared to the plasma chamber diameter, radial confinement of them is inherently poor. A multipole field, usually sextupole, is added to the solenoidal field in order to provide a (very imperfect) rising-field geometry in the radial direction. This added field complicates things considerably. The loss cone on the ends becomes a tri-star "Mercedes emblem" and on the sides becomes 6 *loss lines* along the sextupole permanent magnets morphing into three near the ends. The terminus of all can be seen as marks on the inside of the plasma chamber.

The relative weakness of ion radial confinement compared to that of electrons creates a situation where ions are confined primarily by a "core" of cold electrons and distribute themselves via a process called *ambipolar diffusion* (ref. https://en.wikipedia.org/wiki/Ambipolar_diffusion), among others. This tends to order the ion distribution by charge, with the higher charge states closer to the central axis.

Charged particle losses to the chamber walls are dominated by electrons on the ends and ions on the sides. This charged-particle flow to the inner surfaces of the plasma chamber is called wall or *Simon current*.

Whether considering a thermalized plasma or a non-thermalized one in a magnetic field, ions, because of their much larger mass, are more mobile, creating a transition region between a largely equal positive and negative charge mix well inside the plasma to an unequal region approaching outside ground reference. A cross-section looks something like this:

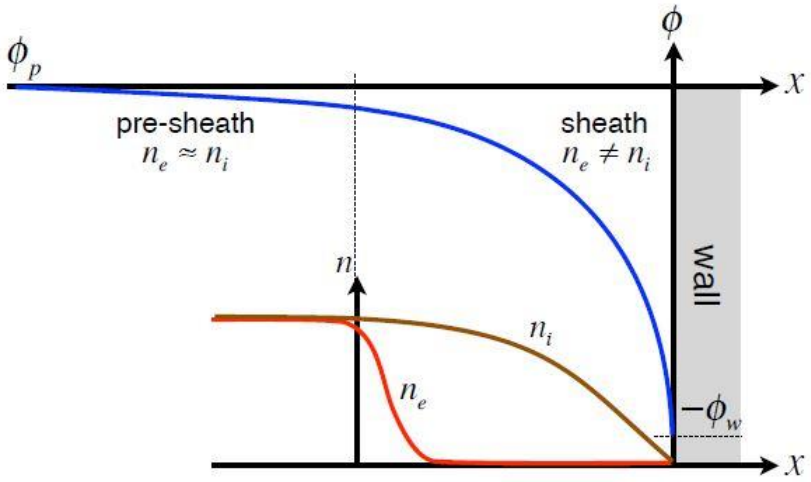

The *plasma potential*, show by the blue line and defined here by $\Phi_p = 0$ V becomes more negative with radius as the electron fraction of the plasma 'mix' gets left behind. Since we define the wall potential as ground, the resulting plasma potential is positive and in practice comes to about 10 – 20 V. This potential, while detrimental to radial ion confinement, when considered longitudinally provides some measure of "boost" of leaking ions toward the extraction aperture, but hinders the escape of ions from the plasma toward extraction if it is too high.

A general feature of all plasmas is that, taken as a whole, they are electrically conductive. Electromagnetic waves (microwaves) cannot penetrate a conducting surface! Considering the resonance shell, the distorted egg where the scalar magnetic field satisfies the resonant condition, as that electron density increases, that shell becomes ever more solid and efficient as a microwave shield. Upon initial source turn on at low power, microwave heating is unhindered; assuming roughly appropriate gas feed, more power gives more ionizations. At modest power settings however, one can easily see the self-shielding effect happening in real time. Each step up in power gives a corresponding jump in beam output, but one that decays from its local peak over the next couple of seconds as the increased electron density creates an ever-more effective conductive shell, reducing the efficiency of the incoming microwaves in heating cold electrons. Eventually, increased power becomes ineffective in increasing beam output.

Important in concept is to know that the plasma and ion distribution within the plasma chamber is very inhomogeneous yet, at half-centimeter scales or so, highly structured. A first cause of structure comes from the ECR resonance shell depicted in the Mechanical and Magnetic section previously. Hot free electrons are preferentially created on this shell, hence ions are preferentially ionized near this shell. Additionally, being stuck for the most part to magnetic field lines, hot electrons created here simply can't easily get to all regions of the chamber. "Creation structure" can be seen in these rather dramatic end-view X-ray images of argon plasma taken in 2004 by S. Biri.

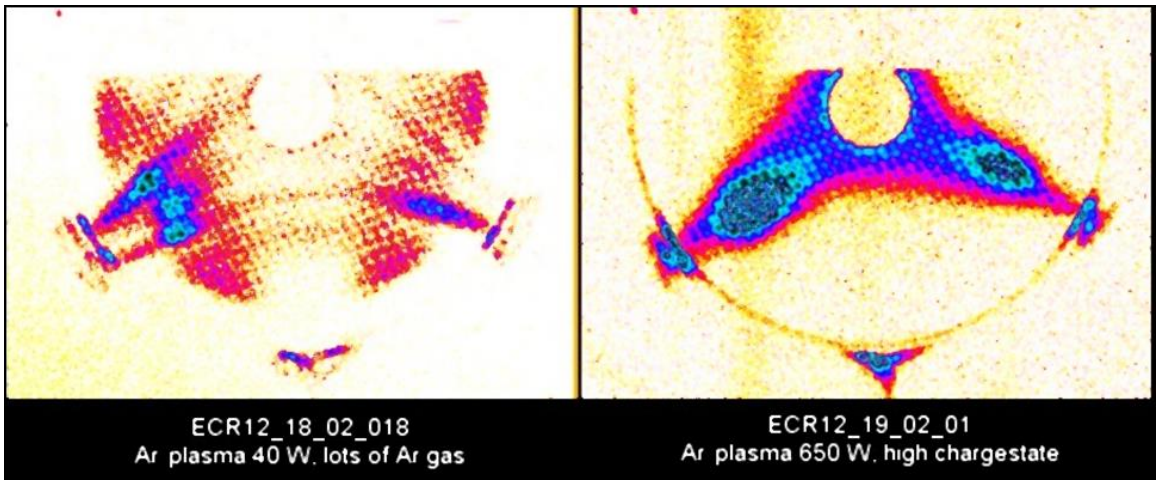

[*Strongly related, but outside the plasma chamber, let's follow some of the effects of this initial radial field sextupole asymmetry downsteam.*]

Thus forearmed, the following image of a $^{16}O^{3+}$ beam, extracted and analyzed, from the NSCL Artemis-A makes considerably more sense than otherwise:

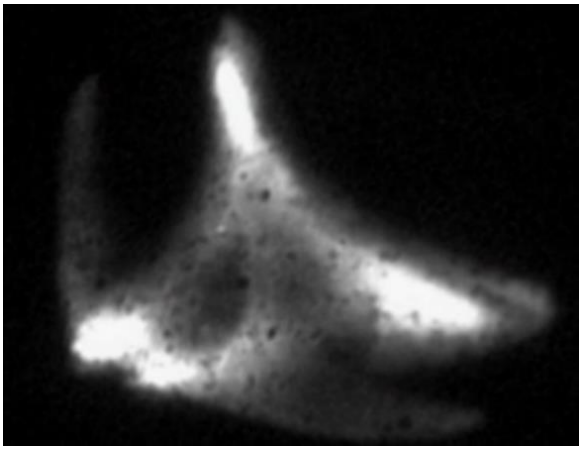

Star and triangle shapes, usually hollow, will often "appear" and "disappear" depending how the beam is focused in any particular location, but never really go away. They come from the complicated interaction of the hexapole field during the beam creation, and especially, extraction process, which serves to create a 2nd order x-y correlation, complicated by simple x-y correlations introduced by

downstream solenoids. While significant for very-low emittance injection into the K500, these distortions seem at present to be generally ignorable for linac operations, but worth noting to avoid confusion. For example, an ARTEMIS-B beam of $^{40}Ca^{9+}$ from December 2020 illustrates a typical-for-ECRIS "open triangle" on a LEBIT viewplate:

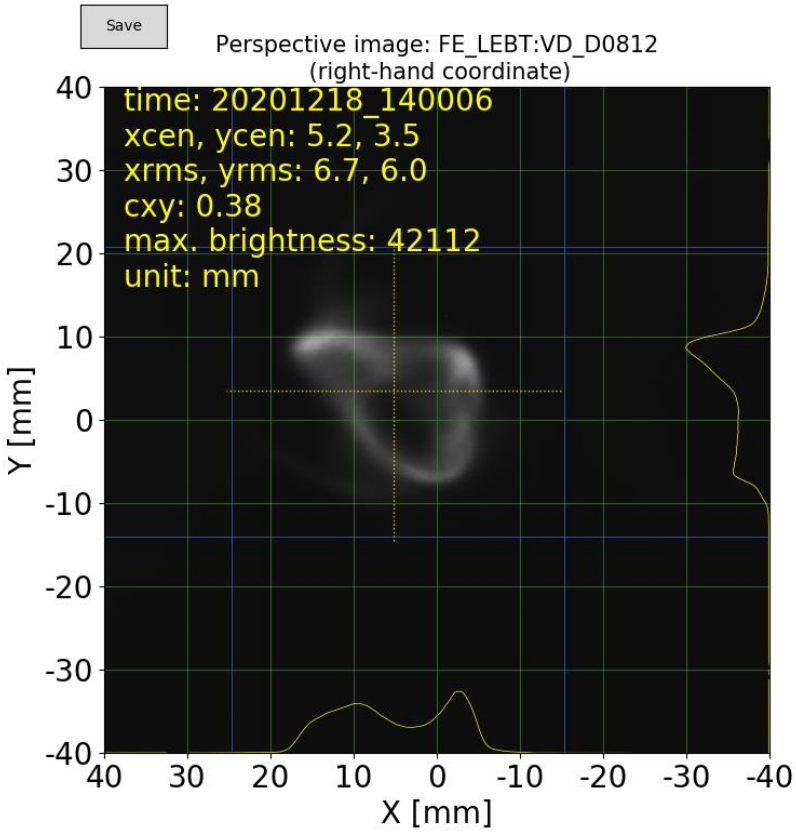

If one were to look at only the X and Y intensity profiles reconstructed along the sides of the plot, one might think that the beam consisted of two overlapping spots. However visually, the asymmetrical, non-uniform beam shape is clear, and one can immediately grasp the reason for the odd profile projections. Unfortunately, camera images are, with a couple of exceptions, available just in the early part of the LEBT beam line. Only 3-wire profile monitors are available downstream and understanding details of the beam shape without an accompanying camera picture is more difficult and ambiguous.

[*Now, a return to plasma chamber interior processes.*]

A key factor in source operation and beam output has proven to be the *bias disk*. All agree that it has clear benefits, but similarly agree that the mechanisms of how those benefits are achieved are less clear. Early ECR ion sources did not have them. In an attempt to increase the number of free electrons within the plasma chamber the rather obvious idea of providing extra electrons by shooting them into the plasma with an *electron gun* (https://en.wikipedia.org/wiki/Electron_gun) was experimentally tested with success. Overall beam output increased, but particularly so in the higher charge state region.

[Contextual Aside (already!): ECR ion source development thru the 1990's emphasized achieving very modest intensity of high charge state ions and particularly so when injecting into cyclotrons. The lower the M/Q, the higher the achievable energy for that ion from the accelerator at little additional cost. Direct beam experiments plus low-bandwidth detectors and analogue electronics often resulted in maximum useable beam currents well below 1 μA, so high source output in general was not a priority. Better data taking technology and the desire for producing secondary beams changed this considerably. Use of a stripping foil after acceleration to an intermediate energy allows the ECRIS to produce charge states near an intensity maximum in a middle range of what is possible, and stripped to higher CS later. Present ECRIS development at FRIB and similar facilities now emphasizes this intensity rather than the highest possible charge.]

However, electron guns supplying a sufficiently high current tend to be complicated, bulky, and take considerable power to run. (Getting AC line power to a source platform at high voltage is a non-trivial matter, involving at least an *isolation transformer* (https://en.wikipedia.org/wiki/Isolation_transformer) for modest potentials and a variety of complex schemes for voltages much above 100 kV.) Subsequently then, a much simpler on-axis hot filament, producing electrons by *thermionic emission* was installed on the injection end of the plasma chamber. This similarly improved source beam output, but also had considerable power requirements and filament lifetime was short. Finally, a negatively-biased disk in the same location was discovered to have the same benefits as the gun or filament while also being simple and cheap. Depending on the specific configuration, the power supply used provides -20 to -300V with a current draw of ~ 0.5 – 4 mA requiring a Watt or so of power. (This versus several kW for an electron gun).

The conventional wisdom was that the bias disk worked by "supplying electrons" to the plasma, but that has since proven to be too simplistic, if not entirely incorrect. Work on a full explanation is ongoing and it is far from settled. But at least a more-correct view is that the disk serves to reflect near-axis cold electrons that would otherwise be lost, back into the plasma. This effect is enhanced when combined with the extraction potential at the plasma aperture creating a column of trapped cold electrons near the axis. (Clear is that with a bias disk potential of, say -100V, approaching electrons of less than 100 eV will be repulsed in the opposite direction. Less clear is that the extraction aperture serves the same function. Consider an ECRIS which sits at ground potential. The bias disk remains negative in this view, but the extraction electrode and beam line must be negative in order to accelerate ions, so internal electrons reaching the edge of the extraction aperture see a repulsive force. Voltage can only be defined as a relation to something else.) Besides providing extra electrons for heating, the electrostatic forces from this column pull ions in closer to the center line, enhancing both their confinement time and the probability for extraction. Best evidence for this and of operational significance is that applying voltage to the disk results in an immediate change of the extracted beam size and shape. A negatively-biased disk may also serve to change the longitudinal charge distribution: ions leaking from the extraction aperture are then balanced by those being sucked onto the bias disk at the opposite end. Finally, by modifying the Simon currents, plasma potential is reduced, allowing ions to more easily escape.

Source performance has been shown to be relatively insensitive to the shape, size and material used for bias disks. (An approximate optimum seems to be a round disk of aluminum, with a diameter equal to the RF wavelength, that is, 21 mm for 14.5 GHz.) However, behavior is sensitive to its longitudinal location relative to the injection side field peak: outside, at, or inside (closer to the plasma)

of that peak. Bias disk position on FRIB sources seems to be generally in the latter condition (varies with which injection assembly is used) and is not remotely adjustable at the time of this writing.

Consider again the anisotropic flow of charges lost to the inner chamber surfaces. Even closely-adjacent surfaces may be exposed to considerably different charge flows and, were the surface dielectric, charge and hence voltage differences would accumulate. The conducting aluminum surface however prevents this and essentially "shorts out" the circuit. Introduction of a bias disk on its own, independently-drive circuit fundamentally changes the balance of Simon currents and thus fundamentally influences plasma behavior in a very direct manner. Likewise, the plasma affects the bias disk in a very direct manner.

Taken as an electrical circuit and a component of Simon current, the bias disk is directly connected to the plasma and becomes an extremely sensitive sensor of plasma conditions. That is, given that the voltage on it is held constant by its power supply, the current through that power supply is very sensitive to changes in source conditions. (With no plasma, the bias disk current is essentially zero. All current when running represents a flow through the plasma.) A small change in gas flow for example, may not change beam output or measured chamber pressure in any noticeable way, but almost always will change the bias disk current.

The bias disk can also be considered as a "single-type" Langmuir probe. These probes are a fundamental experimental tool for understanding plasmas. The Wikipedia article is well worth a read for a bit of knowledge about these probes but also for a compact introduction to several plasma physics concepts: https://en.wikipedia.org/wiki/Langmuir_probe. Mentioned briefly in this article is resistivity which leads to another important effect of the bias disk.

With voltage applied by the power supply and a resulting current, quite trivially one can calculate a V/I resistance (impedance). Since the plasma is at a net positive potential relative to the chamber surface, it must contain a net charge. It therefore has a capacitance, which is "seen" in the bias disk circuit (and the "Simon circuits") as an RC network https://en.wikipedia.org/wiki/RC_circuit. Bias disk voltage can thus serve to enhance or dampen oscillatory behavior. While this explanation is over-general, lacking and/or downright wrong, the conclusion is correct: the voltage applied to the bias disk has a huge effect on plasma and beam stability. This will be discussed further in the "Operation – Stability" section later.

There remains an entirely different way to use the bias disk: floating. That is, the power supply is disconnected entirely with the disk still insulated from the rest of the chamber. Oddly (amazingly?) this still enhances beam output considerably over the no-disk or grounded (to the plasma chamber)-disk case. Observed is that the disk will charge up to a few volts, say -5 to -20 V and vary with tuned conditions. The fact that the voltage seen is in the range of expected plasma potentials (but, as wired, with opposite sign) is probably not coincidental. It is also generally accepted that source stability is also enhanced.

Adequate Physics Model – Extraction from the Chamber

The general mechanical and electrical layout of the region of first beam acceleration is shown here:

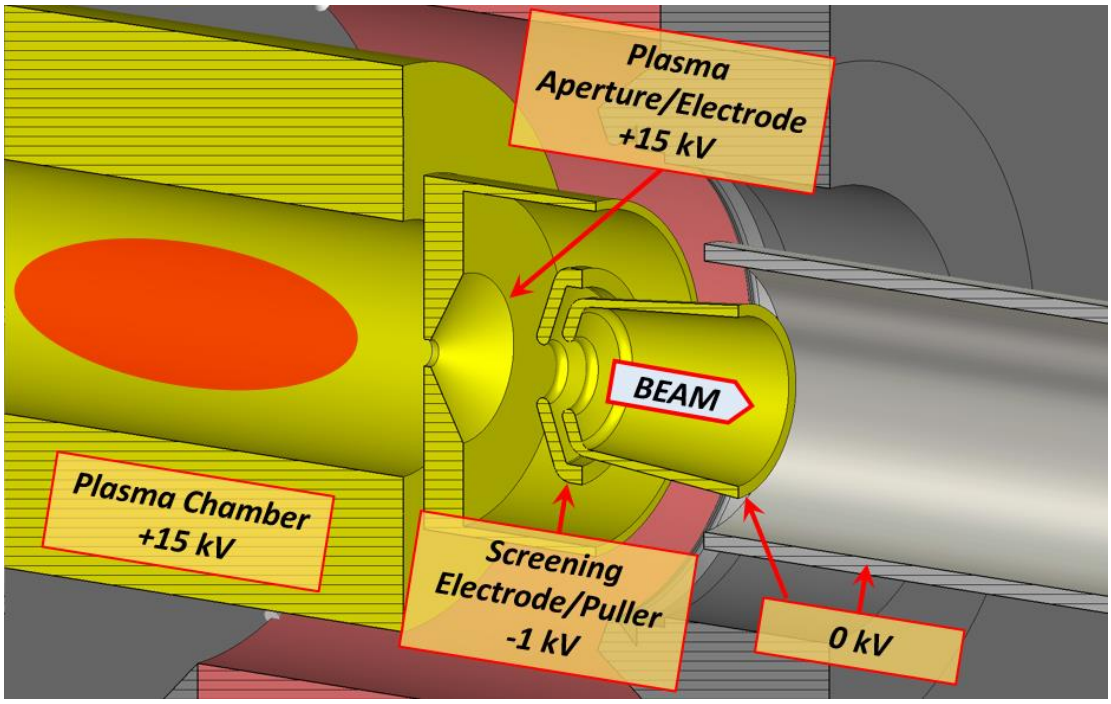

Beam exits the plasma chamber through a 10-12 mm hole. We want lots of it. That means in the region between the two electrodes it has a high charge density so transverse *space-charge* forces and low downstream velocity conspire to cause a large radial size expansion in beam size over even this short distance. (See https://en.wikipedia.org/wiki/Space_charge for further explanation.) The concave aperture and convex puller surfaces are intended to counteract this expansion. These shapes give the acceleration field a tilt toward the beamline axis, creating a focusing effect. A *Pierce geometry* is sketched below. While ARTEMIS uses this geometry, it is ideal only for a pure electron beam emitted from a flat surface. Ion beams require modifications to this shape, which also depends on the ion masses. A single ideal shape for all of our beams would be impossible to achieve.

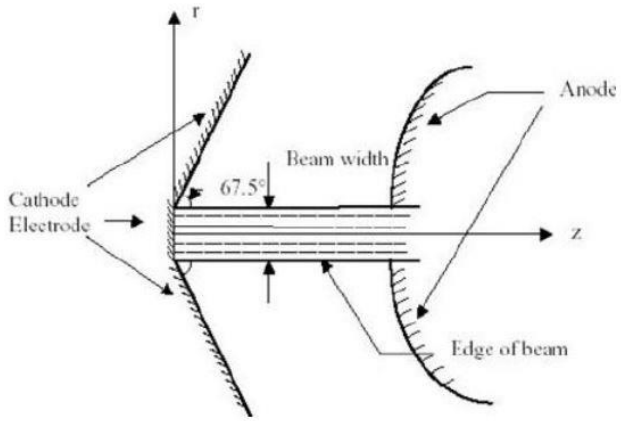

Simple considerations (*Child–Langmuir law*) would have us, for maximum intensity and with maximal space-charge forces, operating with the puller quite close to the plasma aperture, about 10 – 15mm where this focusing is strong. We generally do not approach the limits of this law and need less focusing, which can be accomplished by increasing this gap. Too, poor pumping speed in this region combined with outgassing from surfaces hit by beam conspire to create poor local vacuum. That, with a narrow gap creating a high field gradient encourages sparks across the gap which are to be avoided. So, we run with much larger gaps, presently 38 mm, between the aperture and extraction electrode (puller).

Details of what happens to beam in this region are complicated, but some practical things to list are:

1) Changing the puller position changes the initial beam focusing which, in turn, changes the beam shape at every location downstream.
2) The higher the source bias (extraction voltage), the smaller the effect of space charge and the higher the ultimate current limit.
3) Different extracted beam currents require different focusing strengths to achieve a consistent downstream beam condition. But we don't (yet) normally change that.
4) The puller position is not to be considered as a normal optimization parameter. Small changes make little difference in output beam current and larger ones have significant optical effects downstream.

The puller (extraction electrode) on ARTEMIS is normally negatively biased. Doing so increases the electric field strength in the gap, resulting in faster acceleration and hence, less time for space-charge forces to act within this critical region. While this might be considered as a first stage of an *accel-decel* beam extraction system, in practice here, that effect seems quite minor. Another intended effect is to stop electrons downstream from being attracted to the positively-biased source and *backstreaming* into the region, though it likely requires a smaller aperture in the puller and higher voltage than presently used. Some improvement to beam output has been noted by biasing the puller up to a voltage of about -700V, but, at present, none much beyond.

After the initial 15-20 kV acceleration, beam enters the *acceleration column* which is little more than a stack of metal rings, insulated from one another except through connecting ~ 80 MΩ resistors between each ring. (For examples, see https://www.pelletron.com/wp-content/uploads/2017/02/Accel-Tubes-v2.pdf.

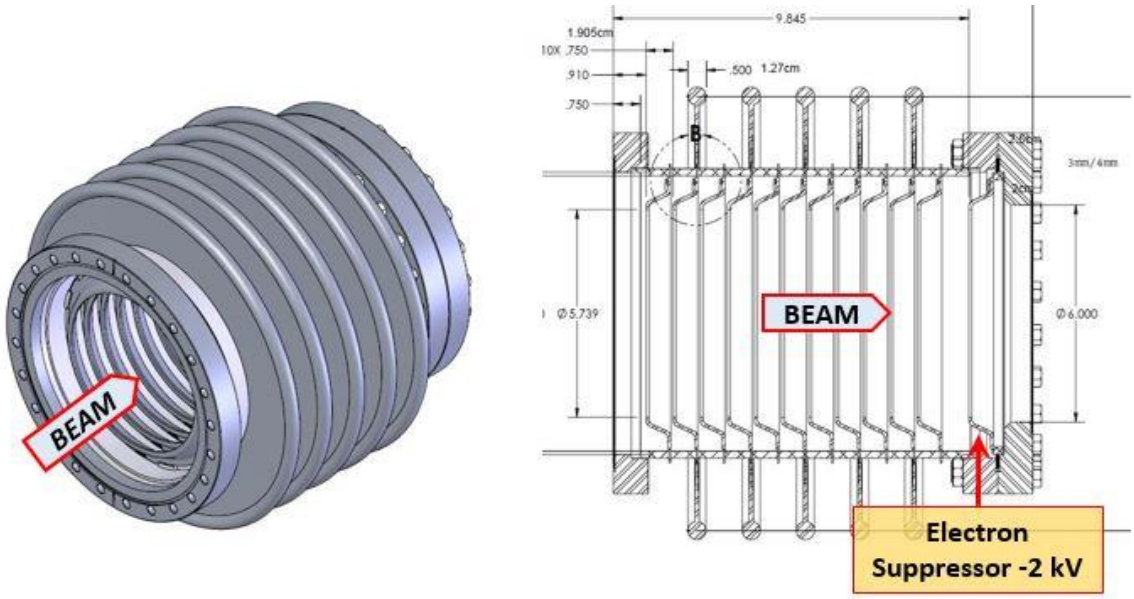

The ring furthest upstream is the first step down from platform potential and furthest downstream is, in principle, at ground. This stacking of rings is a standard practice. Rather than a “giant leap” across the full platform-to-ground potential, the potential is divided evenly into a series of gaps. Besides reducing chance of damage to the column by sparks, errant ions and back-streaming electrons, the voltage difference between successive rings not only accelerates the ion beam but provides a focusing effect similar to that of an *Einzel lens*. The arrangement also provides a measure of fault tolerance, if for example, beam induced sputtering creates an internal short between two adjacent rings. The last ring, which would otherwise be at ground potential, is instead biased to -2 kV. This is intended to stop any free electrons in the beam line (caused by beam impacting residual neutral gas) from accelerating up toward the positively-charged source.

[Aside: The path to failure. The FRIB ECRIS acceleration columns have a sequence of rings, forming 12 acceleration gaps. Suppose one starts with a platform potential of 96 kV, then the first ring is at 90 kV, the next, 84 kV, and so on in a sequence like this: 96 -> 90 -> 84 -> 78 -> 72 -> 66 etc. The voltage difference between rings is 6 kV. Now suppose that the insulation fails between the platform and the first ring. Then, rather than the voltage across each gap being $1/12^{th}$ of the starting voltage (6 kV), it becomes $1/11^{th}$ (8.73 kV) and the sequence of voltage drops is: 96 -> 96 -> 87.3 -> 78.5 -> 69.8 -> 61.1 etc. Beyond a slight change in the focusing strength, it’s likely that the failure would go unnoticed. But the remaining “good” insulators become more stressed by the higher gap potential, increasing the chances of another failure, followed by another increase in gap potential. When enough insulators have failed, sparks across the gaps become likely and the column would then have to be repaired or replaced. What here is supposed, has already happened on Artemis-B: the first ring is shorted to the platform, most likely due to a piece of construction debris left inside. No ill effects have yet been noted, but may appear over time.]

Ideally, all extracted ions would have motion entirely in the downstream direction and have zero motion in the transverse. This is not so, of course, and for essentially three reasons:

1) Ions within the source, while cold relative to the electrons, are hot by normal standards. The transverse thermal motion continues throughout the entire acceleration process and cannot be erased. The normalized, RMS emittance due to this contribution can be written [taken from Winklehner PhD thesis, 2013] for an idealized situation as follows, where r is the radius of the emitting disk in this case, the plasma chamber opening:

$$\varepsilon_T{}^{xx'-rms-norm} = 0.016\mathrm{r}\sqrt{\frac{kT_i}{M/Q}}$$

2) Ions are born in a very high magnetic field and are extracted to a field-free region. Just as Lenz's Law states, a changing magnetic field component orthogonal to a loop of wire induces a current in that wire, this change in field strength along the beam path induces a rotational impulse to ions during extraction. (A more-correct explanation is given by the Bush Theorem). Once in the field-free region, this impulse ceases, but the acquired linear transverse motion continues. Again, for an idealized situation, the contribution to normalized, RMS emittance from this effect can be written as:

$$\varepsilon_B{}^{xx'-rms-norm} = 0.032r^2B\frac{1}{M/Q}$$

Both effects are unavoidable. Because of the very high magnetic fields required for an ECRIS to produce high charge state ions, the latter, magnetic effect, dominates the thermal by at least a factor of 10. A reason to include these two lonely equations is to point out the latter's dependence on distance from the axis and inverse dependence on M/Q. For a given mass, a lower charge reduces emittance growth. However, a higher charge is better confined within the plasma chamber, hence exits closer to the axis. This effect appears to be stronger; experience shows that for a given mass, higher charge states have lower emittance ("beam's smaller").

3) The sextupole field imprints strong 2nd order aberrations on the beam, which often appear as a triangular or 3-pointed star shape. For this and spherical aberrations introduced by solenoidal focusing, the beam also often appears to be hollow.

Using these formulae and typical values, one can plot the effects of items #1 and #2 for various situations. A graph for the case of $^{82}Se^{14+}$ with an ion temperature of 4 eV and magnetic field of 0.85T is shown below (ARTEMIS aperture radius is 6 mm):

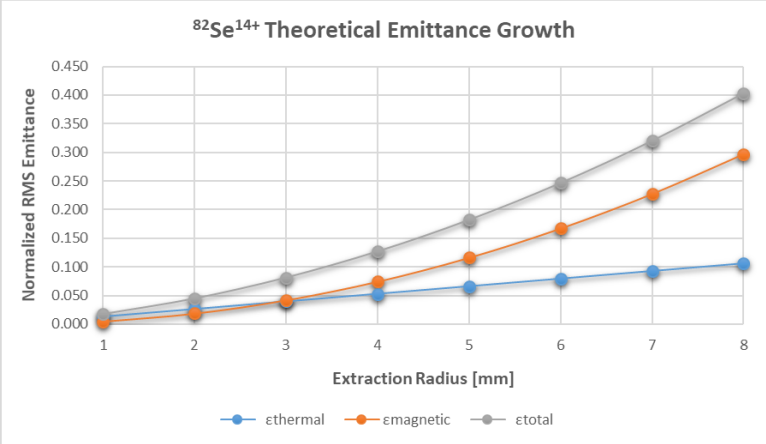

Important to know is that this is a very simplified ideal model, but it serves well as a basis for seeing the source from a larger overall perspective. For example one can see that using larger plasma aperture diameters in order to get comes with a rapidly increasing beam size and emittance. On the other hand, real-world measurements of higher charge state beams show that the emittance is less than what would be expected from this model, which implicitly assumes that ions are evenly distributed across the opening. This fact, plus many others, support the notion that the higher charge states are concentrated closer to the axis than those of lower states. Yet unavoidable is the conclusion that high-charge heavy ion sources requiring high magnetic fields produce beams of much lower quality than sources of electrons and protons without these high fields.

Adequate Physics Model – "Physics Calculator"

Many standard calculations, some, near-requirements for operation, some optional, can be done using the Physics Calculator.

Starting from the CS-Studio Launch Page:

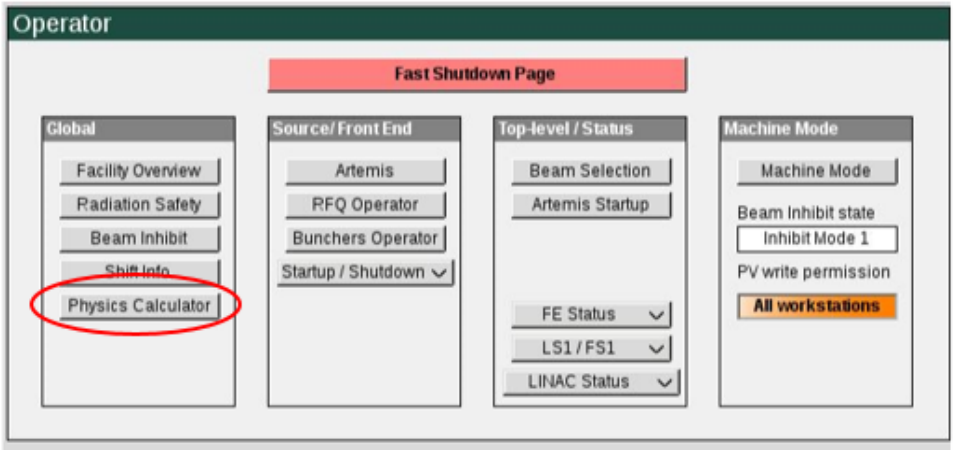

Leads one to this page:

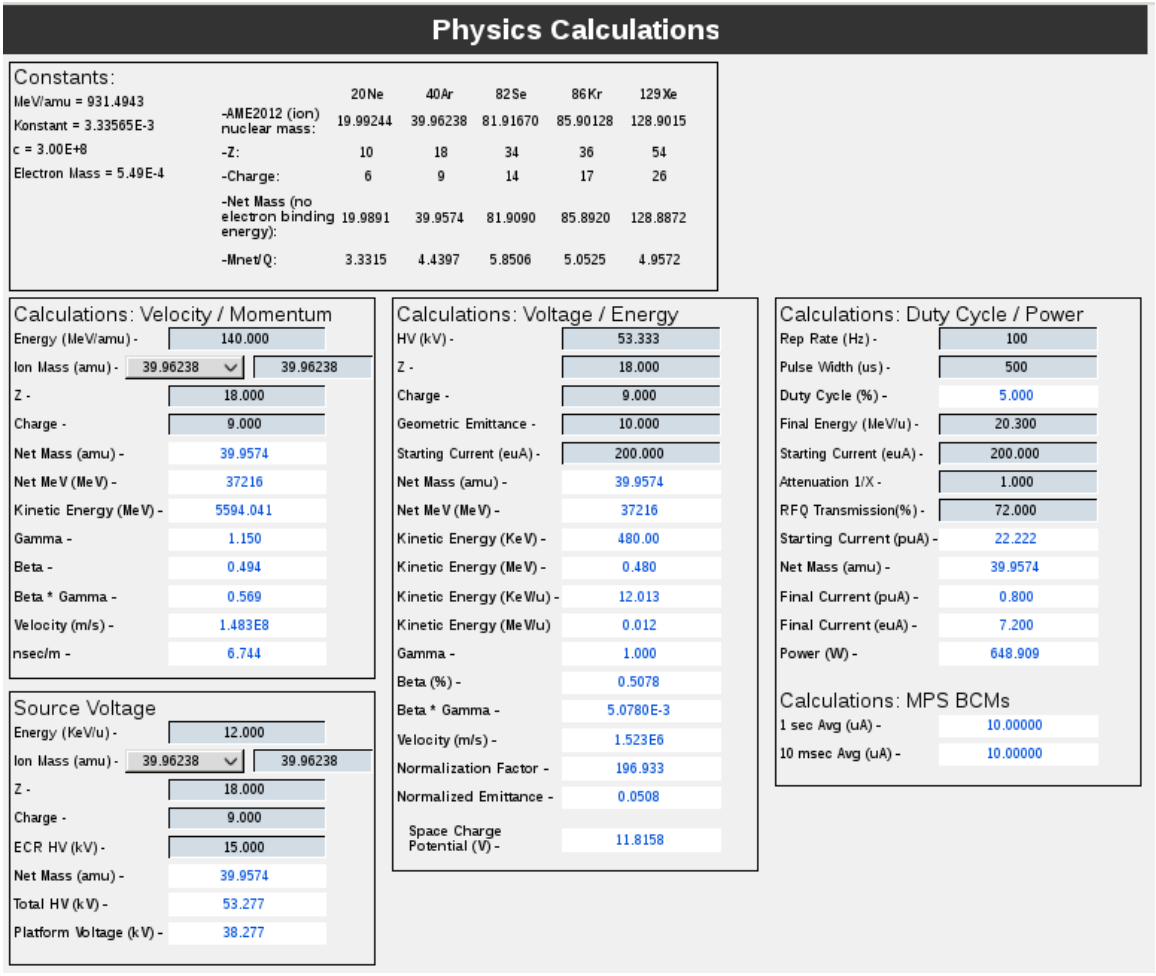

It's important to know that nothing one enters or calculates on this page changes any setting anywhere else, so one may safely experiment manipulate input parameters (in the grey boxes) and examine the results given as blue numbers. Specific uses include estimations of beam power for a given duty cycle (very important for remaining within the safety envelope) and determining the right source + platform potential to give the correct beam energy (12.013 keV/u) for injection into the RFQ. Over time, more features and functions will be added as operational needs evolve.

Operation – Knobs and Meters

An ECR ion source has a rather short list of devices under direct operator control. Specifically, ARTEMIS-B presently has the following "knobs", roughly in order of injection to extraction (`FE_ISRC1:` prefix assumed):

1) ECR (extraction/bias) voltage* `[HVP_D0679]`
2) Controls for heating solid-fed beam samples
3) Bias disk voltage `[PSB_D0679]`
4) Injection solenoid `[PSOL_D0682]`
5) Gas feed valves `[LKV_D0683]` and `[LKV_D0684]`
6) Microwave transmitter power `[RFA_D0679:POWR_CSET]`
7) Extraction solenoid `[PSOL_D0685]`
8) Puller voltage `[PSE_D0686]`
9) Puller position* `[DRV_D0686]`
10) Acceleration column electron screen* `[PSEL_D0698]`
11) Platform voltage* `[HVP_D0698]`

The * - marked parameters are part of the startup process, but not normally optimized in the tuning process. ECR voltage is presently being set to 15,000 V, though will likely be raised in the future as development progresses. The platform voltage is determined via the "Physics Calculator", desired ion and charge, and a final energy of 12.013 keV/u. The accelerator column electron screen seems to have little effect on the beam, so it's normally set to -2000V and left there. In contrast, the puller position has a large effect on the beam shape downstream, so it's only to be moved as part of beam development, not as a part of routine tuning.

The source is tuned by adjusting a "knob" from the above list and observing the result on some, or all, of the following directly-read parameters. Seldom will the change of one controllable device result in a change of only one of these read-back values! Again roughly in order of their physical location, these "meters" are:

1) ECR (extraction/bias) current, a.k.a. "drain current" `[HVP_D0679:I_RD]`
2) Voltage and power from solid-fed beam heating power supplies
3) Injection-side pressure `[IG_D0679:VP_RD]`
4) Bias disk current `[PSB_D0679:I_RD]`
5) Plasma chamber pressure `[IG_D0684:VP_RD] – not in service`
6) Microwave transmitter reflected power `[RFA_D0679:POWR_RD_RFL]`
7) Extraction-side pressure `[IG_D0686:VP_RD]`
8) Puller current** `[PSE_D0686:I_RD]`
9) Acceleration column electron screen current** `[PSEL_D0698:I_RD]`
10) Platform drain current `[HVP_D0698:I_RD]`
11) Beam current on the first Faraday cup `[FE_SCS1:FC_D0738:AVG_RD]`
12) "Scope" display of beam on FC_D0738

**- The acceleration column electron screen normally draws only a small current, insensitive to tuning and is therefore not useful to closely watch. Since the puller also has a large aperture relative to the

usual beam size, the measured puller current is also not a terribly useful tool. (This changes however as the puller position is moved closer to the plasma aperture and far away or when extraction voltage is unusually low).

Even from a simplified viewpoint, what these devices do and how they respond to changes requires considerable further explanation.

Operation – Vacuum – Hardware - Details

The pressure levels and responses to parameter changes play a vital role in source operation, hence it is vital to have a good "viewpoint" in mind when looking at the displayed read-back numbers. Shown below is an annotated snip of an *opi* for ARTEMIS vacuum systems. Note that the depicted beam direction is opposite that shown in previous sketches.

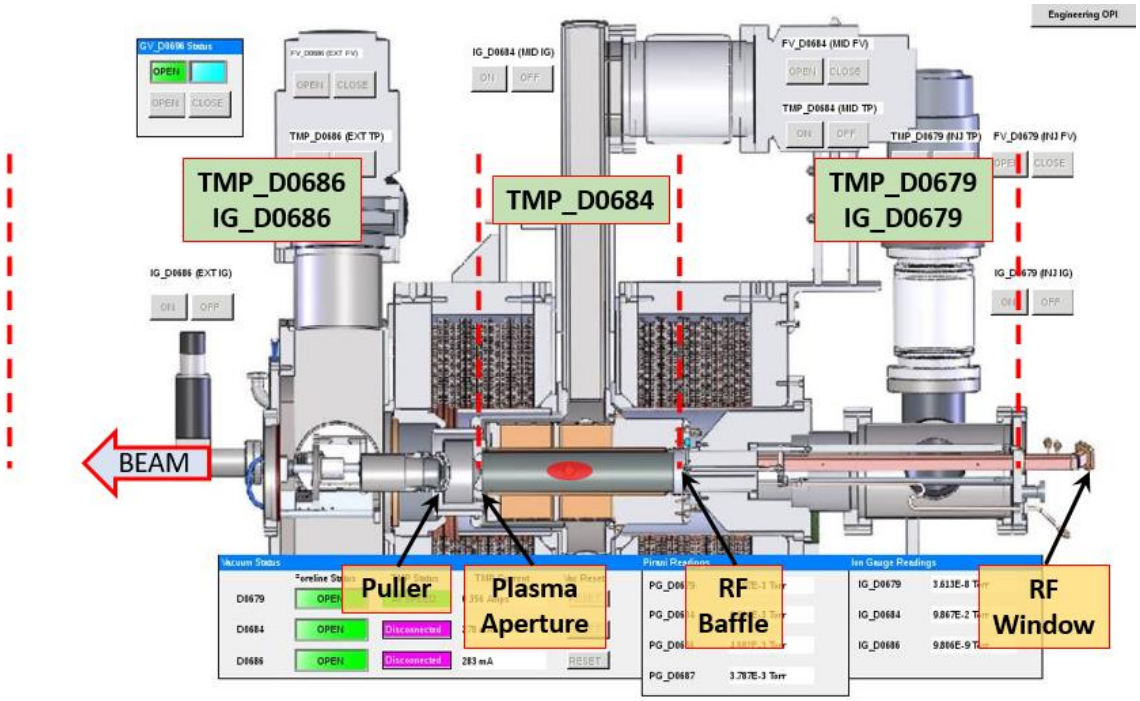

There are three distinct sections: injection, plasma chamber, and extraction, each with their own turbomolecular pump (TMP), but, very unfortunately, only two have associated ion gauges (IG). With no gate valves between them, at pressures over a few Torr when gas flow is laminar, these sections are quite 'connected' to each other. As the pressures become lower and flow becomes stochastic, conductance through the 12 mm diameter plasma aperture and especially past the RF baffle rapidly decreases, therefore the vacuum behavior of each section becomes more independent of the other.

The RF baffle and structure near it is a slip-fit within the surrounding cylinder, but, while designed to be RF-tight, is not classically vacuum tight. However, since there is no line-of-sight between one side and the other, it can result in a large pressure difference. The injection assembly 'stems' are far dirtier than the plasma chamber surface, plus have water lines and many opportunities for virtual leaks,

so, when the source is not running, pressure on the injection side [IG_D0679:VP_RD] will always be higher than in the plasma chamber [unmeasured] and extraction box [IG_D0686:VP_RD]. Another connection (usually) exists unfortunately, between the two sections in that TMP_D0679 and TMP_D0684 use the same scroll pump for backing vacuum. A pressure increase in one side of the plasma baffle will cause the backing pressure to rise, reducing pumping speed for the other.

An important note here relative to vacuum is that while the *outside* of the RF waveguide(s) belongs to the injection-side, the *inside* of the waveguide(s) belongs primarily to the plasma chamber. This is so because of the unobstructed opening of the waveguide end to the chamber.

A related very important note is that the vacuum inside the waveguide only extends to the RF window. Beyond that window is air. The window is a very thin sapphire piece of about 2 sq. cm. supported only on periphery, squeezed between two brass clamps and sealed with indium. It is fragile and when it fails, tends to fail catastrophically. Because it's a straight, unobstructed shot into the plasma chamber, complete venting happens fast, sending a pressure wave downstream into the LEBT, scattering broken window pieces along the way, until interrupted by automatic closing gate valves. This is bad. Be nice to the RF windows by:

1) Avoiding mechanical impact or bumps.
2) Avoiding sudden changes in pressure while intentionally letting the source up to air (nitrogen) or pumping it down again. Both should be done from as far downstream as possible since the plasma aperture and distance slows down pressure changes. Vent and rough pump from the extraction section, not from the injection side.
3) Not "fat-fingering" the microwave power to excessive values or trying to run at higher-than-previous levels on a whim. Some power is absorbed by window during operation and heat can break it. We don't yet know where the limits are, but haven't yet run over 900W for any length of time. Also, the standing waves within the guide change if the hookup is changed with different ends and fittings. That can also change the level of power lost in the window, so care is needed after changing to an untested arrangement.

The vacuum transition region between the plasma chamber and extraction box is also pressure-critical. More beam striking the puller surface can increase pressures via outgassing, sometimes to the point of causing voltage breakdown, especially when it is close to the extraction aperture. The voltage of the plasma aperture relative to platform voltage is 15 to 20 kV. For the puller, it's ~ -1 kV, thereby increasing the potential difference between the two surfaces. The distance between them is simply the puller position [DRV_D0686]. The electric field gradient of course increases directly with decreased distance.

Note also that this region is connected, when the D0696 gate valve is open, via the beam pipe to the downstream vacuum systems, so deviations there as seen on [FE_SCS1:CCG_D0707:VP_RD] - near the first dipole magnet will affect the extraction box pressure and vice versa. This pressure near the dipole will generally rise with beam intensity; rejected ions, that is, ones with an undesired M/Q, are lost in the dipole or nearby on the charge selection slits. It will also noticeably change when the emittance scanner or viewer plate is inserted. Outgassing of these surfaces causes this rise.

Operation – Vacuum - Heat and Environmental Effects

Internally, ECR ion sources producing multi-charged ions must operate in regions considered *ultra-high vacuum*. (From the physics point-of-view, it would make far more sense to call it "ultra-low pressure" rather than "ultra-high vacuum", but, engineers …). The phrase *base pressure* as we use it, refers to the pressure observed with no beam and all power supplies turned off. The lowest base pressure observed with the ion gauge installed on the ARTEMIS extraction box `[IG_D0686:VP_RD]` is about 1 x 10^{-8} Torr or 10 nTorr in engineering-style units.

[One thing to note is that the particular gauge installed while, in principle a sensitivity of 1 nTorr, the current output for this value is ~10 mA, which being read as it is on a 0 – 10V scale gets translated into a control system read-back in a rather coarse manner. A true UHV ion gauge would be a significant improvement, but can also generate X-rays under some conditions. Eventually, one hopes, CCG gauges will be installed to give more accurate values at these low pressures.]

The pressures observed with `[IG_D0686:VP_RD]` during more-or-less stable gas beam operation are generally in the range of 40-150 nTorr, with > ~300 nTorr (3 x 10^{-7} Torr) getting into a range of worrisome. Corresponding injection-side pressures, `[IG_D0679:VP_RD]`, will be about 100 nTorr, but will drift, usually higher, over a number of running hours. Solid feed beams, tend to run with higher pressures.

Since even the highest pressures are ultra-high vacuum, or close to it, outgassing from surfaces within the source vacuum system play a huge role in both the ion gauge readings and source performance. In turn, heat, both localized and in general, plays the major role in that rate of outgassing present. A perfectly clean aluminum surface will outgas very little. In the sorcerer's vernacular, this clean condition is called *conditioned*. However, undesirable contaminants will stick to, mix with, coat, or react with the otherwise pure surface and be released when heated, with resulting increased measured pressures.

a) Stick to = Adsorption https://en.wikipedia.org/wiki/Adsorption
b) Mix with = Absorption https://en.wikipedia.org/wiki/Absorption_(chemistry)
c) Coat = solids heated by the oven, subsequently condensed on the inside walls.
d) React with = chemistry specific to the materials involved

"Heating" as used here can be simply a temperature increase, more-or-less evaporating surface contaminants to be pumped away by the turbopumps. In vacuum technology, improving the base pressure by temporarily raising the temperature to a level considerably higher than that during its normal running state is called *baking-out* a system. (Components exposed to vacuum in the FRIB accelerator chain including all beam line components are baked-out before installation.)

In an ECRIS, the most dominant outgassing, however, occurs as a result of the plasma inside of it. While also an effect of increasing temperatures, the conditioning process is considerably more complicated and is very similar to a widely-used industrial technique called *(oxygen) plasma cleaning*. https://en.wikipedia.org/wiki/Plasma_cleaning

Environmental temperatures of both water and surrounding air probably also affect internal source pressure though the effects are not without some ambiguity. Base pressures usually rise when

the water skid (circulation) is first turned on, but this is usually simultaneous with turning on the solenoids as well. Three effects confound immediate conclusions however: (1) There may be slight water leaks from copper-tube joints of the water leads cooling the area behind the plasma baffle. (2) Magnetic fields will alter the electron flow within the ion gauges and thus affect the pressures read. (3) The solenoidal field wants to "spit out" the sextupole which is rigidly attached to the plasma chamber assembly. It may shift slightly and affect the various gasket and O-ring seals to outside air.

Changes in the room air and cooling water temperature have a major effect on source vacuum readings, beam output and other source parameters, as charted in the following example.

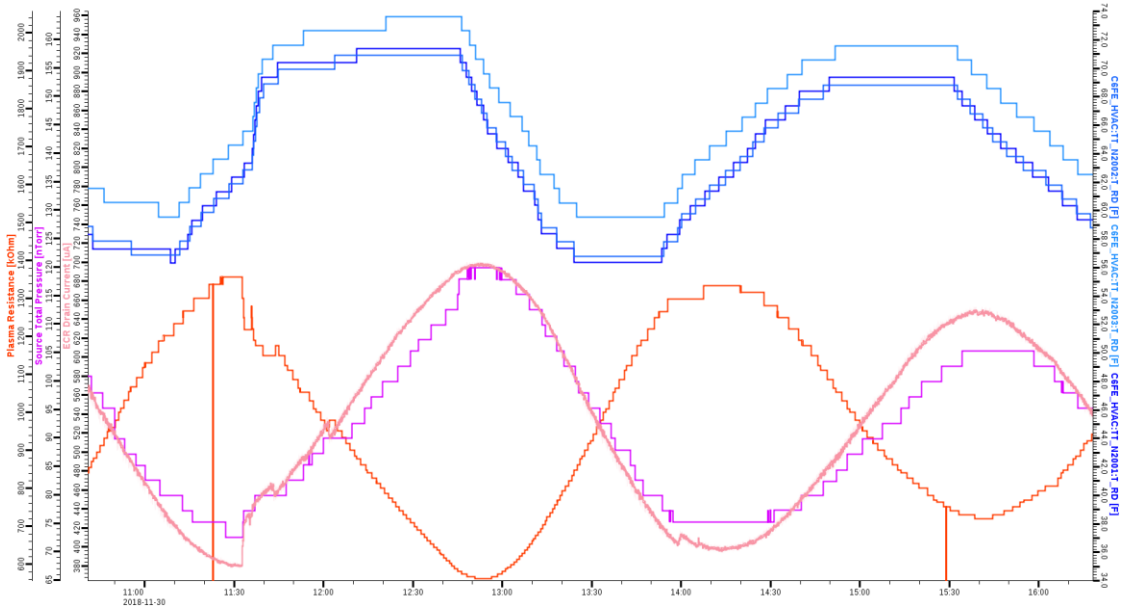

The top 3 blue lines are air temperature sensors of a misbehaving HVAC system in the front end, varying +/- 6 degrees F with a 2 ½ to 3 hours period. The purple line shows the sum of injection and extraction pressures, `[IG_D0679:VP_RD] + [IG_D0686:VP_RD]`, varying between 75 and 120 nTorr. The same periodicity also being shown by the orange line, ECR drain current `[HVP_D0679:I_RD]` ranging from 350 to 650 µA which implies a similar variation in output beam current (not shown here). The primary driver of these variations is thought to be the change of air temperature changing the pressure of the feed gas going into the leak valve thus changing the mass flow going into the plasma chamber. Similar source changes occur with variations in cooling water temperature.

Operation – Water Cooling "Skid" System

Chilled cooling water is supplied from outside of the Faraday cage to everything on the HV platform and therefore must cross multiple boundaries between large differences in voltage. To accomplish this without draining charge across these potentials, the resistivity of the water must be high. The normal value during operation is considered to be 6 MΩ/cm. However, as soon as the skid

system is turned off and flow through the resin bed (which serves to de-ionize the water) stops, the resistivity begins to fall. Stasis lasting a couple of weeks or so will drop it to that of something between distilled and rain water, ~0.1 - 0.2 MΩ/cm. Turning the skid back on recovers the desired resistivity, but may take hours or days to do, depending how low it has fallen and the time-in-service of the resin. Something around 2 - 3 MΩ/cm is considered "good enough" if time is a factor and one is far from the voltage maximum setting. However, there is always some charge escaping the source and platform which has to be compensated by current from the power supplies if a constant voltage is to be maintained. During operation, the accelerated beam itself is responsible for most of the drain current seen on `[HVP_D0679:I_RD]` and a substantial part of current seen on `[HVP_D0698:I_RD]`. The remaining amount is called *dark current*, which refers the beam-off leakage loss while holding a constant high voltage. Normally a fairly small and negligible value, it increases with poor water quality (or failure of other insulators).

This chart shows an example platform dark current vs. voltage relationship over a limited range:

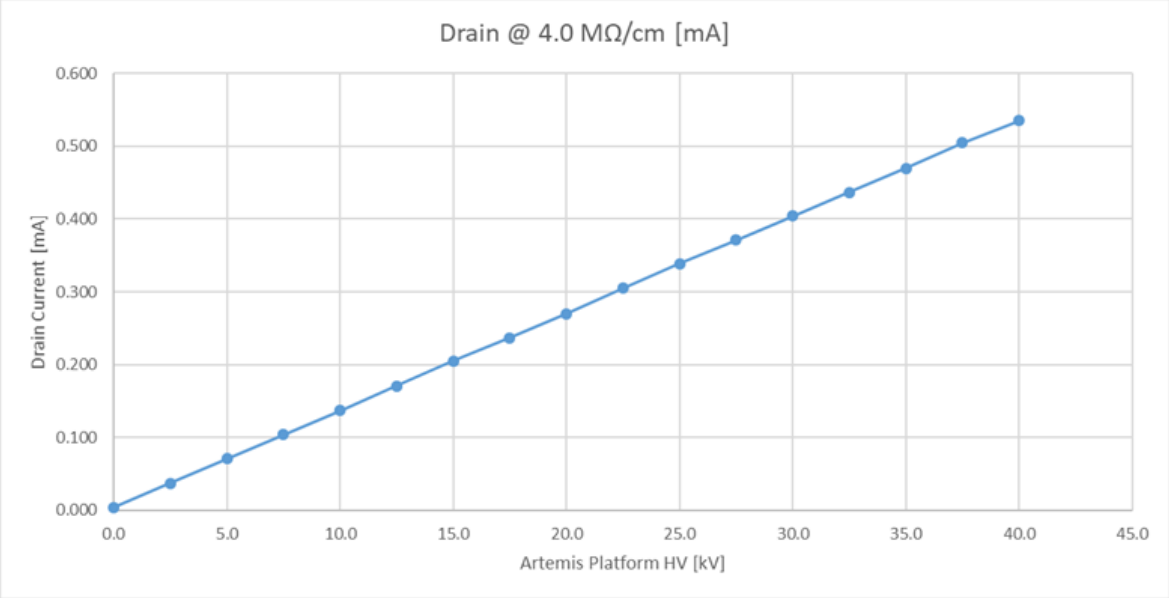

Notice the extreme linearity of this plot. Generally, as the voltage increase approaches a point of breakdown (spark), this line will hook upwards and thus give a warning that something bad is about to happen. During the water *polishing* process, the rise in resistivity will result in a decrease of dark current, and hence drain current for any given voltage or condition. The system is supposed to regulate to a value of 6 MΩ/cm, but is sometimes unreliable. Without control feedback, it will reach 10 MΩ/cm or so after some period of time. While this is relatively harmless, the resin bed lifetime is reduced somewhat.

[Many other factors are also important to cooling water suitability for accelerators. An older but concise and clear summary is available here: https://ieeexplore.ieee.org/abstract/document/1288995 (R. Dortwegt, "Low-conductivity Water Systems for Accelerators," Proceedings of the 2003 Particle Accelerator Conference, Portland, OR, 2003, pp. 630-634 Vol.1, doi: 10.1109/PAC.2003.1288995.)]

A serious issue we presently have is that the skid temperature controller does not hold a sufficiently-stable output water temperature. Changes of cooling water temperature affect beam output when the source is running, though it's not clear through what mechanisms.

The work-around solution at present is to disable temperature regulation by the controller local to the skid and to set the *mixing valve* locally. The majority of water volume is pumped "unaltered" in a closed loop to and from the HV cage devices. To control water chemistry, a portion is diverted by a bypass valve through the resin bed as mentioned above. Another portion is diverted by the mixing valve through a heat exchanger and is cooled by ~46 F water from the building chiller. The heat load on this water skid presented by cooled devices, primarily from the source solenoids `[PSOL_D0682]` & `[PSOL_D0685]` and first beam line solenoid `[PSOL_D0704]`, is rather low compared to the cooling power available. Hence, with these magnets on (source on), proper water temperature is achieved with the mixing valve set to 21% and, with these magnets off (source off), the valve set to 17%.

Very important to note is that one should never run the skid with the solenoids off and this valve set to 21%. The circulating water temperature will get too cold which can chill the plasma chamber assembly to the point that water can condense on it from surrounding air, compromising insulation and creating sparking paths to ground. Likewise, with magnets energized and a valve setting of 17%, overheating will occur within ~10 minutes. The only way to know what position the valve is in, is to look at the readout on the local controller – there is no PV for this purpose.

Operation – Platform & HV Cage Checks

Access to the High Voltage cage is granted to those who have completed the appropriate training, with additional training and authorization required to access the shielded area surrounding ARTEMIS. These requirements are not to be ignored.

The two source gas inlet valves, `[LKV_D0683]` and `[LKV_D0684]`, are fed via Poly-Flow plastic lines from three different manifold systems. In order to work properly and consistently, the gas must feed at a constant pressure of ~10 psi. (The input pressure specified by the valve documentation is 15 psi. If an expensive enriched gas is being used from a small "lecture" bottle feeding the auxiliary manifold mounted close to the source, the pressure may be set lower.) Verify that this is the case from the pressure gauges, trace each line back to its bottle of origin, verifying that each manual valve in the path is open. Leaving a valve closed during operation will slowly deplete the gas in the line, resulting in a tune drift and a requirement to continually open the main leak valve. If allowed to continue, the pressure will drop below atmosphere, encouraging contamination from outside. When the pressure drops low enough, discharges of source high voltage will occur, first resulting in a neon-tube lighting effect, followed by full sparks to ground and possible damage. This effect is generally described by *Paschen's law* (https://en.wikipedia.org/wiki/Paschen%27s_law) and can be visualized for some common gasses as shown below:

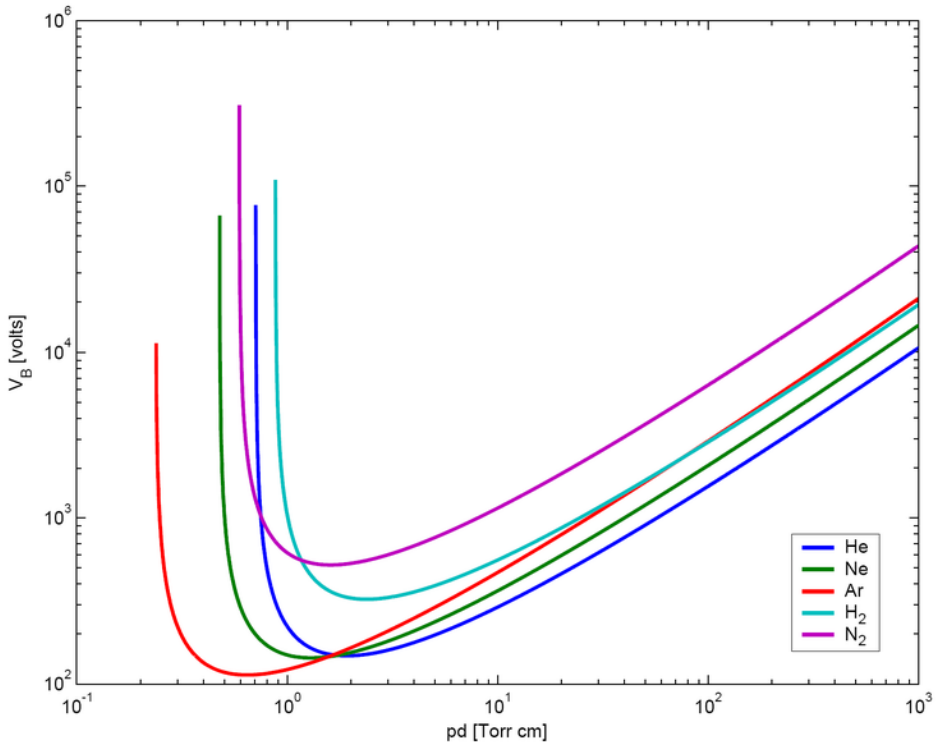

[During this inspection, note which gas is connected to which feed valve. Usually, the main (beam) gas is attached to `[LKV_D0683]` and supporting gas, mostly oxygen or helium, to `[LKV_D0684]`, but this is not always the case. Enter the appropriate information on the "Artemis Startup" opi.]

In the same vein, and especially after shutdowns or source configuration changes, look for loose and misrouted wires and cables, which can also become a cause of HV discharge. The plasma chamber assembly and everything electrically attached to it, including the power supplies of and wires from the floating rack, will be at 15 – 20 kV relative to platform ground. The insulation of wires will not protect against this kind of difference in potential. This is also true of "BNC-like" HV cabling which has sufficient insulation to hold HV between the inner and outer conductor, but not between the outer conductor and a high voltage (relative to it) exterior. HV sparks into or from cabling can destroy the power supplies or control circuits to which they are attached, as well as preventing `[HVP_D0679]` from holding a stable voltage.

Unstable HV or high currents from HV power supplies can also result from dirt left to accumulate on insulators. Over time, "tracking" patterns across insulator surfaces can provide a visual clue that current is leaking to ground. Hence, any dirt accumulation found on plastic or ceramic insulators should be removed before running the source.

Finally, *any* water leak, puddle or even dampness should be reported for immediate attention and action. The Poly-Flow lines and fittings on the end of the injection assembly are especially vulnerable to leaks and the region between the outer circumference of the plasma chamber and solenoid bore is especially sensitive to moisture and, once wet, very difficult to dry.

Operation – Charge State Distribution

(This section is unavoidably tedious and detailed. There are many intersecting factors and layers of information required for understanding and they defy neat organization. Minor apologies in advance).

The *Charge State Distribution* (CSD) is a key tool for ion source operations. Upon extraction from the source, the beam is a mixture of both ion species and charges. The energy gained by dropping from high potential to ground potential is simply the sum of the platform and ECRIS voltages multiplied by the charge of the ion, expressed in eV (electron volts) or keV. Ions with the same mass but different charges therefore will have different velocities and hence, momentums. Since the force exerted by a magnetic field is velocity (and charge) dependent, their paths diverge in the 90 degree *analysis magnet* `[FE_SCS1:PSD_D0717]`. In other words, they bend along two different radii as shown in this sketch from the classic description of the TRANSPORT matrix-based charged-particle optics code by K.L. Brown (CERN80-4).

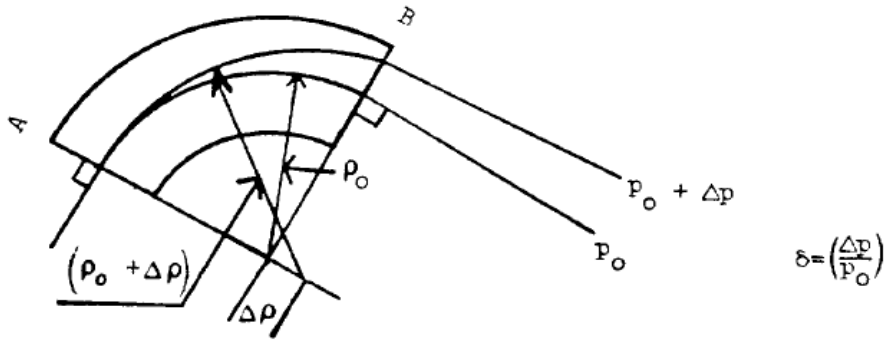

The field of the analysis magnet can be set to pass ions of the desired charge state along the *central ray*, centered on the downstream beam line. Off-axis ions can be stopped by slits blocking their deviant paths. Important to note is that because the initial acceleration is electrostatic, all ions and charges are focused/bent identically by electric fields in electrostatic quadrupoles (and bender/dipoles); path variations only appear in magnetic fields.

Ion mass of course comes into this, but the principle is the same: turn momentum difference into a path difference. Since the initial ion beam is also a mixture of species, a useful convention would be to divide the force equations by the mass (in *amu*) and consider q/m rather than mass and charge independently. Because however, source people are weird and it directly scales with *magnetic rigidity*, use of the inverse m/q is common. [Magnetic rigidity, which can be expressed in units of "Tesla meter," is a directly accessible measure of how difficult it is to bend a particle of a certain mass, charge and energy in a magnetic field. A 1 T*m ion in a 1 T field will bend with a radius of 1 meter. For scale, a $^{36}Ar^{10+}$ beam of 12.0 keV/u has a magnetic rigidity of 0.0567 T*m. At the full energy of 257 MeV/u, the value becomes 4.918 T*m for $^{36}Ar^{18+}$. {Since the practical limit of normal electromagnets is about 1 Tesla, the need for superconducting magnets for high-energy heavy ions, even one as "light" as ^{36}Ar, with M/Q = 3.6 stripped to M/Q = 2.0, is clear.} Protons, with M/Q = 1 have half the rigidity of these $^{36}Ar^{18+}$ ions at whatever equivalent energies per amu. Scaling is simple, in other words. Knowing the rigidity and

designed bend radius of the magnet also makes the conversion to field and reverse similarly easy. [FE_SCS1:PSD_D0717] has a bend radius, ρ ~ 25" = 0.637 m. Hence for the example $^{36}Ar^{10+}$ beam of 0.0567 T*m, the field in this magnet would be 0.089 T or 890 Gauss.]

So, ECRIS beams get spread in position after bending through the analysis magnet ordered via rigidity by M/Q, higher M/Q bent less, lower, more. Except for a few specific cases, no attempt is made to measure these positions directly. Rather, the horizontal (bend plane, a.k.a. *dispersion* plane) slits at [D0736], are closed symmetrically about the beamline axis and resulting current measured on the D0738 Faraday cup [FE_SCS1:FC_D0738:AVG_RD] behind them while the strength of the analysis magnet is varied and data recorded.

A calibrated and labeled result of such a scan of beam current versus M/Q for a ^{40}Ar beam is shown below. The X-axis derives from the power supply current which is basically linear with field, which is linear with rigidity, which in turn, is linear with M/Q. One could do the math, but more practical and sensible is to examine such a distribution, identify a clear known peak, assign that peak the correct M/Q value, which then allows a constant conversion value to be used, converting current into M/Q, to identify the others. Here is an example:

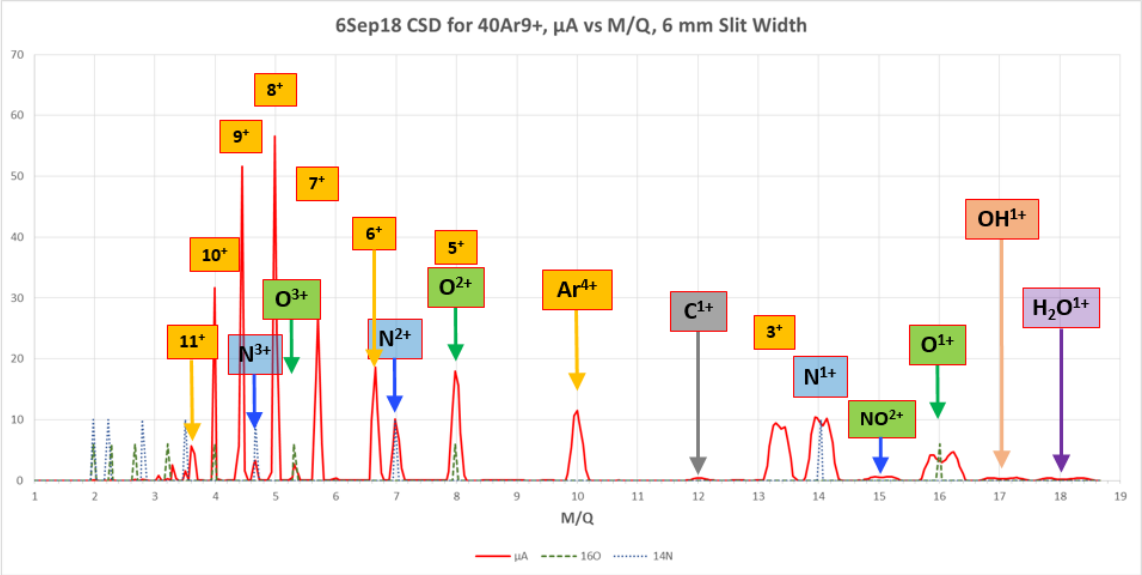

The charge states of argon are identified in yellow. Some oxygen and nitrogen peaks are labeled with text, and all with the artificial dotted line peaks inserted to guide the eye. Even though only argon gas is being added to the plasma, oxygen, nitrogen and carbon appear and will <u>always</u> appear at some level in <u>every</u> CSD scan. This particular CSD however, shows signature of a too-large air component. Possible reasons for that include: 1) a pretty substantial air leak into the plasma chamber, 2) it was taken not long after the source was vented, or 3) the feed gas line was not fully purged after a change. Carbon often appears due to hydrocarbons within the source chamber, in which case it also would show a peak at M/Q = 1 should the scan range be set to reach down that far. Hydrocarbons are also common contaminants in solid-feed material samples.

This permanent signature of carbon, nitrogen and oxygen can, however, be very useful when presented with complicated CSDs where exact calibration of M/Q is vital. It's also useful when dealing

with the primitive display made by the CS-Studio scan tool available for this purpose. (Hopefully a better dedicated CSD app will eventually be available for operator use).

An example scan tool CSD output is shown here:

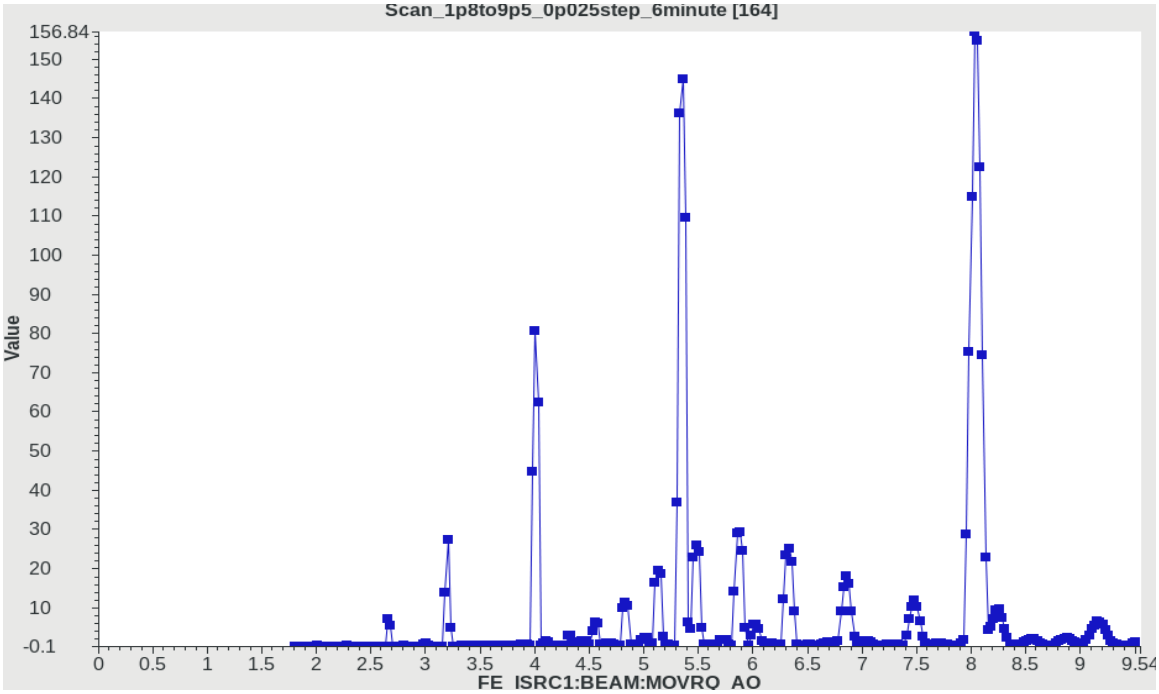

This CSD shows a "solid-feed" beam, selenium-82, which cannot alone maintain a stable plasma. As such, the main plasma is one of oxygen, serving as the *support gas*.

The first thing to see here is that each ion species displays itself in an "arc" of currents. The oxygen here *peaks* in the 2+ charge state. The selenium peaks at 14+, which is the desired CS, indicating that the source has been properly tuned to give a current maximum in that state. There is a 3rd such arc (of 40-argon) in this CSD that is buried at the low end of the current scale in this view. When the CSD is complicated, it's very important to select a peak based on its read-off M/Q value, but to identify the arc to which it belongs and verify that those other peaks fall at the expected M/Q! Delivery of the wrong ion can happen, has happened, and is considered to be very impolite.

[Notable Aside: By convention, for ease-of-communication and because it's easily measured, we usually refer to beam intensity in units of microamps, *μA*, which is then itself often typed as *uA*. The proper unit of measure would be *pps*, particles per second, but that for our purposes would be awkward. (As would measuring beam energy in SI units of Joules). However, since Amps are in units of Coulombs per second, which are an Avogadro's number of unit charges per second, they become essentially equivalent to pps if the charge of each particle is 1. We can then create a new unit, particle microamp, *pμA*, by dividing the current by the charge of the ions being considered. When confusion is possible, we call a normal μA, an electrical microamp, *eμA*. One can do the math, but an easy to remember approximate conversion is 1 pnA (1 x 10^{-9} pA) ~ 10^{9} pps. Relevant to CSDs, the use of eμA rather than pμA, exaggerates the peaks of high charge state relative to lower charge state in terms of their true intensity.]

Next, note that the $^{16}O^{4+}$ peak lines up nicely with this crude-form plot label at M/Q = 4, but the $^{16}O^{2+}$ peak misses M/Q = 8 by a significant amount. Such misalignments are typical and vary with over what range and how fast the analysis magnet is ramped as well as where the slit centerline is during the scan. When picking out the right isotope and charge for acceleration, it is important to cross-check the scaling with a nearby known peak and to not rely entirely on what is displayed in crude form by the scan.

Identification of ions and associated charge states from a CSD is greatly aided by making a table of ions, charges and displaying the math for easy, or at least, easier, comparison. One can make their own on-the-fly using Excel, or use a prepared version (CSD-AME2012) which can be provided. Here's a few select columns from it, with 3 M/Q boxes highlighted in yellow for emphsis:

Z =	6	7	8	8	18	18	34	34	34	34
Q	**12C**	**14N**	**16O**	**18O-0.20%**	**36Ar-0.33%**	**40Ar**	**78Se-23%**	**79Se-9%**	**80Se-50%**	**82Se-9%**
0	**12.00000**	**14.00307**	**16.99913**	**17.99916**	**35.96755**	**39.96238**	**77.91731**	**78.91850**	**79.91652**	**81.91670**
1	11.999	14.003	16.999	17.999	35.967	39.962	77.917	78.918	79.916	81.916
2	5.999	7.001	8.499	8.999	17.983	19.981	39.000	39.459	39.958	40.958
3	3.999	4.667	5.666	5.999	11.989	13.320	26.000	26.306	26.638	27.305
4	2.999	3.500	4.249	4.499	8.991	9.990	19.500	19.729	19.979	20.479
5	2.399	2.800	3.399	3.599	7.193	7.992	15.600	15.783	15.983	16.383
6	1.999	2.333	2.833	2.999	5.994	6.660	13.000	13.153	13.319	13.652
7		2.000	2.428	2.571	5.138	5.708	11.143	11.274	11.416	11.702
8			2.124	2.249	4.495	4.995	9.750	9.864	9.989	10.239
9					3.996	4.440	8.667	8.768	8.879	9.101
10					3.596	3.996	7.800	7.891	7.991	8.191
11					3.269	3.632	7.091	7.174	7.265	7.446
12					2.997	3.330	6.500	6.576	6.659	6.826
13					2.766	3.073	6.000	6.070	6.147	6.301
14					2.569	2.854	5.571	5.636	5.708	5.851
15					2.397	2.664	5.200	5.261	5.327	5.461
16					2.247	2.497	4.875	4.932	4.994	5.119
17					2.115	2.350	4.588	4.642	4.700	4.818
18					1.998	2.220	4.333	4.384	4.439	4.550
19							4.105	4.153	4.206	4.311
20							3.900	3.945	3.995	4.095

Developed over a number of years, the table contains several features that have proven useful. The red boxes highlight ions which have been accelerated by NSCL and have been or planned to be accelerated by FRIB. Ions that have been shown to be contaminants at various times are also included. Some of the stranger ones include chlorine (from scrubbing powder not washed out after plasma chamber cleaning), zinc (from exposed screws), and rhenium (from wire used to heat solid-feed crucibles). Oddly, since the plasma chamber is made from it, aluminum is seldom seem; the only occasions were when the chamber was in the process of being burned through. The percentage shown at the top with some of the names indicate the natural abundance of that isotope. Our bottles of oxygen and argon are very pure oxygen and argon, but contain the natural abundance of oxygen-18 and argon-36, so the resulting CSD will contain them. Fun factoid: We've run $^{40}Ar^{9+}$ through LS1 many times, with M/Q = 4.440 but, since it's so close that it can't be cut away with slits, $^{36}Ar^{8+}$ at M/Q = 4.495 gets a free ride thru. (Both shown in green boxes). Another feature of this particular table is that it uses the actual atomic masses instead of atomic numbers. While normally not necessary, such use makes the M/Q

values more correct. Here, $^{82}Se^{14+}$ is the desired combination with M/Q ~ 82/14 = 5.857. More exactly, using the real atomic mass of 81.91670 and subtracting the mass of missing electrons, M/Q = 5.851.

Consider now the following post-analyzed zoomed in version of the situation shown in the above raw scan:

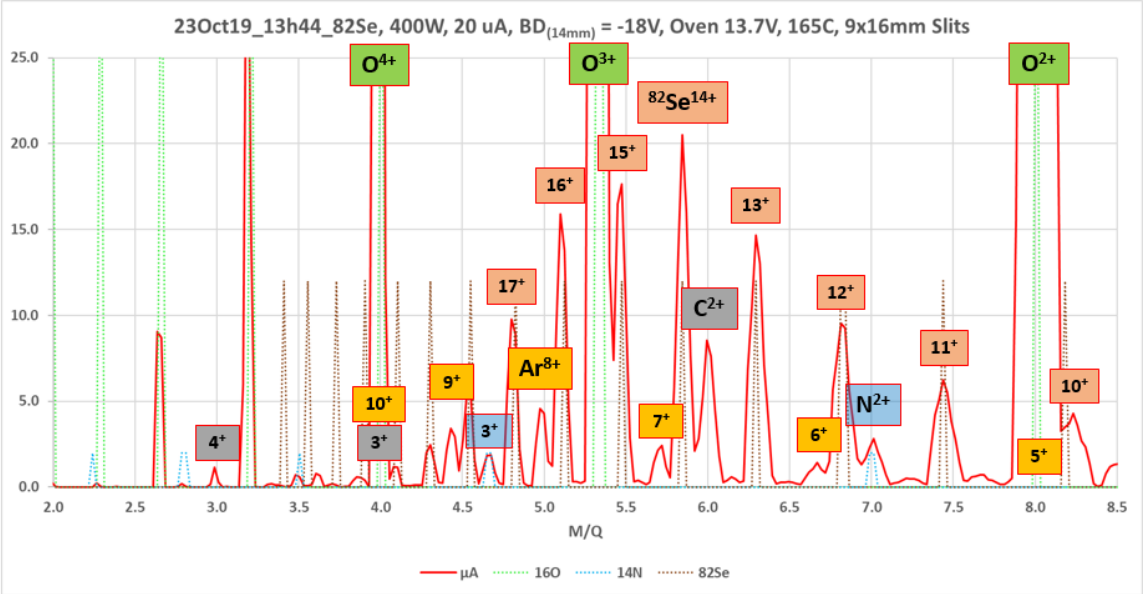

To the left of the M/Q = 5.85 $^{82}Se^{14+}$ is M/Q = 6.0 $^{12}C^{2+}$. (Since for Artemis, the analyzing magnet bends beam to the left and more rigid beam is bent less, the left/right sense of the CSD matches that of left/right w.r.t. the beam direction). To the right, is some "leftover" M/Q = 5.71 $^{40}Ar^{7+}$. (The source had been previously running a ^{40}Ar beam; it takes some time for the residual gas to dissipate completely.) The split on both sides of the $^{82}Se^{14+}$ is ~0.15 or ~2.6% in M/Q. With the dipole `[FE_SCS1:PSD_D0717]` set to pass this desired beam along the axis of the downstream beam line, undesired beams will not be aligned to this path.

$^{82}Se^{13+}$ and $^{82}Se^{15+}$ for example, are so far off-axis that, if not cut away by the analysis slits, they will be lost "somewhere" in the beam line before reaching the RFQ and potentially being accelerated. The "somewhere" loss of such beam can have two consequences. The immediate one is a possible worsening of beam line vacuum from beam striking components which may have not been previously exposed to beam before. This effect is usually minor and temporary. A more serious, but long term effect though is that of *sputtering* caused by the errant ions striking a surface: as an ion approaches, its electric field can cause atoms to be pulled from that surface. These lost atoms can migrate and accumulate on other surfaces, such as insulators and cavities, resulting in damage requiring repair. While it's unlikely that such damage will cause problems with a newly-built facility this month, this week or this year, early prevention saves later grief.

➔ Early prevention prevents later grief.

The 4-jaw slits located at `[D0736]`, the dispersive image, exist for eliminating undesirable ion species and tails. Use them.

Moving the [FE_SCS1:STPC01_D0736:SLR.VAL] slit in from the right (from a positive reading in mm toward negative) peaks with the higher M/Q values are progressively cut away. Doing the same with the [FE_SCS1:STPC01_D0736:SLL.VAL] slit cuts the lower M/Q peaks.

[The convenient way to both control and think of these slits is to use this display box from the "Artemis Setup" page:

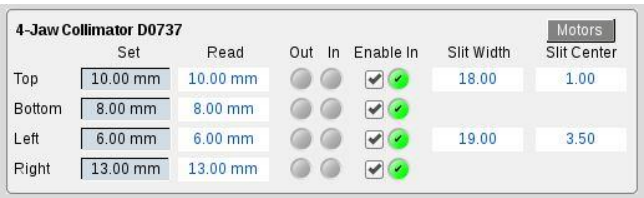

This 'box' example shows the left slit is 6 mm retracted from the defined geometric beam line center and the right 13 mm, so the width (gap) between them is 19 mm. However the center of that gap is offset to the right of geometric center by 3.5 mm, which is not atypical.]

But even with the relatively good resolution shown in this case, there is a limit to the ability of slits to cut away the undesired $^{82}Se^{14+}$ and $^{12}C^{2+}$ ions, even if one accepts a large intensity loss of desired beam from using a very narrow slit setting.

Turning for a moment to a much simpler use of CSDs, and starting with a following rule-of-thumb we have:

➔ Increasing microwave power and reducing feed gas generally move the peak of the distribution toward higher charge states. The converse conditions move it toward lower states.

Compare these two simple CSDs of ^{16}O, different only in the applied microwave power:

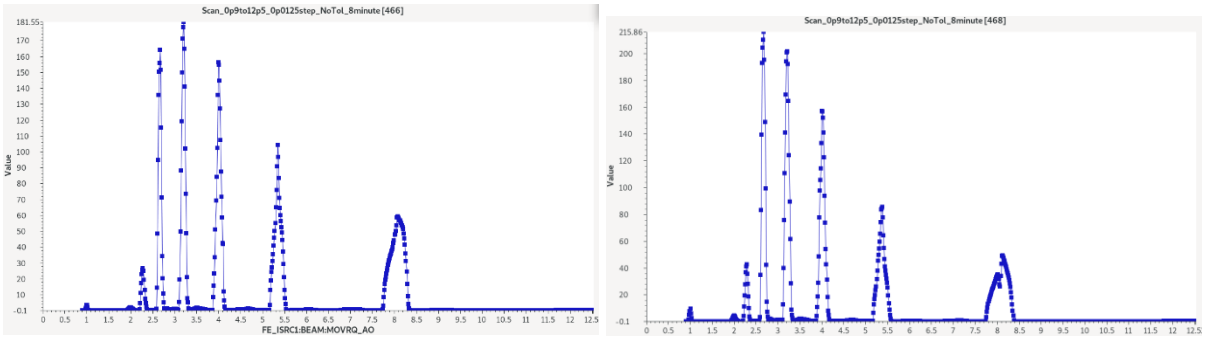

An increase of power moves the peak charge state from 5+ to 6+. The shell effect is clearly seen as the pair of inner K-shell electrons, 7+ and 8+, have much higher ionization potentials than that of the neighboring L-shell.

Operation – Charge State Distribution – (Resolution)

This first beam line segment, SCS1, is set up for use as a spectrometer. The ability of the system to physically separate ions is called the *resolving power* of the system. Given (and required) that the upstream optical elements (quadrupoles, solenoids and the focusing properties of the dipole) provide a true focus from the object (plasma aperture) to image (slit position) in the bend plane of the dipole, then the resolving power is proportional to the *dispersion* divided by the *magnification*. [Given a central ray particle of the nominal momentum, a different momentum will end up in a different place at the image location. That place depends on the dispersion of the system, given typically in cm/% momentum deviation. With a true optical focus, a 12 mm diameter beam coming to a 3 mm spot would have a magnification of 0.25]

Even when tuned perfectly, both dispersion and magnification are finite numbers, so there is a limit to how well the system can "see" closely-spaced peaks. This places a limit to how well closing slits can work in stopping unwanted beamlets. Consider the following sequence of slit-closings:

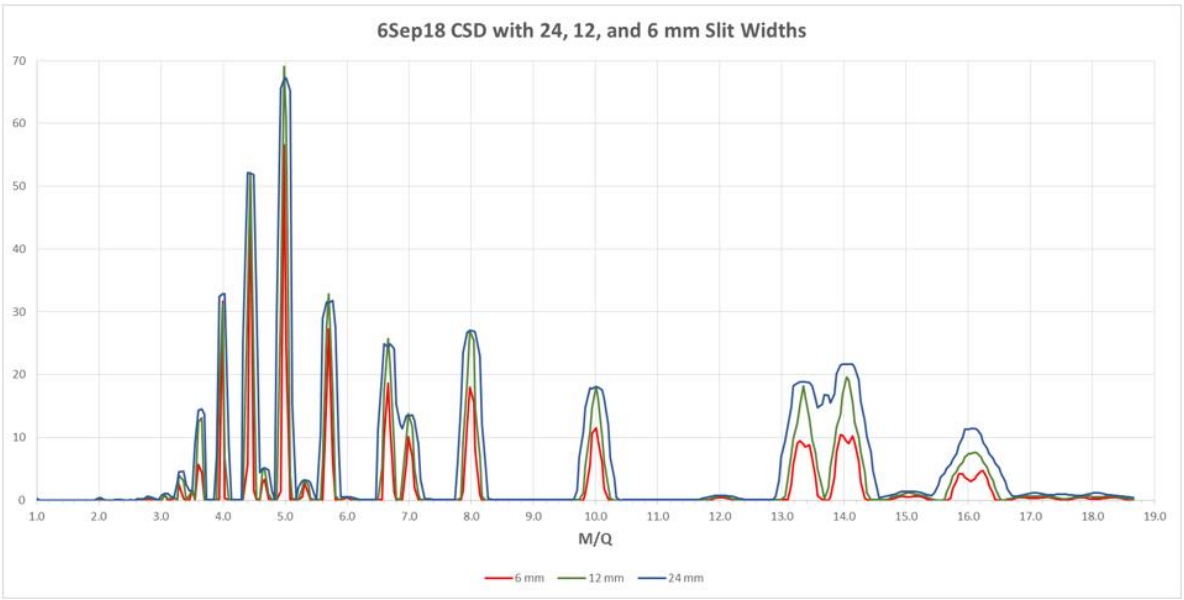

Note the magic appearance of distinct peaks from blobs, where $^{40}Ar^{6+}$ (M/Q = 6.66) splits from $^{14}N^{2+}$ (M/Q = 7.0) and $^{40}Ar^{3+}$ (M/Q = 13.3) splits from $^{14}N^{1+}$ (M/Q = 14.0) as the slits are progressively closed. What appears on the red trace to be double-peaking at M/Q = 13.3, 14.0 and 16.0 is not due to buried isotopes, but is an unfortunate space-charge effect caused by the first beam line solenoid `[PSOL_D0704]`. Tuned as it is to focus the desired M/Q properly, all beams with a lower M/Q are focused more strongly, causing them to pass inside of the column of desired beam. When sufficient total charge density can act over a long enough time, the desired beam is blown apart by repulsive forces into a donut, thereby reducing the resolution of the optical system.

Here is an example of the optical effect of adding more support gas (oxygen) to a ^{129}Xe beam:

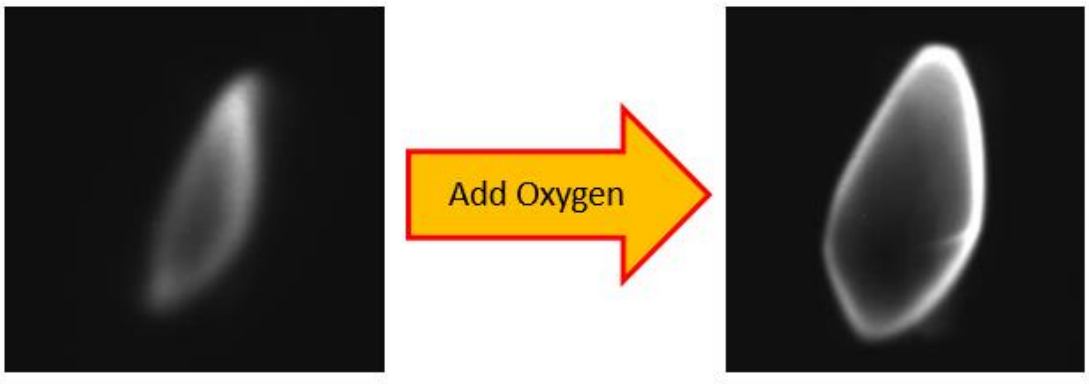

While less added oxygen gives a smaller detrimental effect, it also reduces the xenon beam intensity. Checkmate. An intensity scan with sufficient resolution therefore shows first one side of the oval, then the middle, then the other side resulting in a peak with a mid-region dip. [This "short-focus" problem is exactly why the focusing elements between ARTEMIS-A and SUSI and their respective analyzing dipole magnets are entirely electrostatic. Electrostatic elements focus all ion species in exactly the same way for a given source and platform potential.]

Another problem stemming from the present first beam line solenoid `[PSOL_D0690]` and the fact that present tunes do not use the next solenoid `[PSOL_D0704]` at all, is that as the solenoid is ramped up with the dipole `[FE_SCS1:PSD_D0717]` to pass the more rigid beams, it hits its current limit of 150 A near M/Q = 12 for an ECR extraction voltage of 15 kV, and lower, M/Q ~ 10, for a voltage of 20 kV. This can be seen in the Data Browser plot below, showing beam current (red), analyzing dipole current (green) and beam line solenoid current (yellow) taken during a CSD scan:

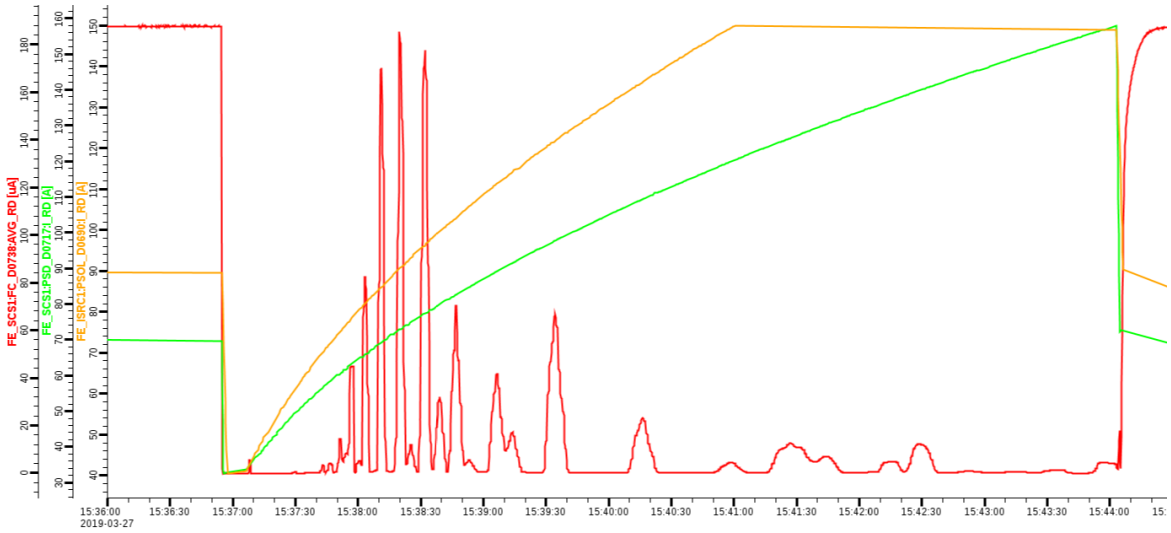

This causes a loss of transmission of the higher M/Q ion species through the dipole and larger-than-otherwise spot sizes at the slits and less beam intensity at `[FC_D0738]`. The "arc" of the CSD is thereby distorted, under-counting beam at the high end of the M/Q range. One might point out too, that, in principle the sum of all measured peaks in a wide-scan view should equal the ECR HV power

supply current, minus the dark current and assuming no transmission losses. Ideally, this give a solid number for transmission that can be used for goodness-of-tune purposes. Extra transmission losses in these high M/Q beams makes this difficult or impossible.

Closing slits however, will not improve the inherent state separation. That depends on both the initial design and how the beam line upstream of the analyzing slits is tuned. The designed beam spot at the slit position is a horizontal focus, but is not a vertical focus. It should, in other words, be small horizontally, so as to not overlap adjacent peaks, with no concern in this regard to its vertical size. It's useful to be able to see the spots directly using the camera and scintillation view plate at D0739, at the location of the slits. (However, maximum intensity is limited by fiat to 20 eμA)

The following is such an image, taken with a $^{129}Xe^{27+}$ beam (M/Q = 4.774):

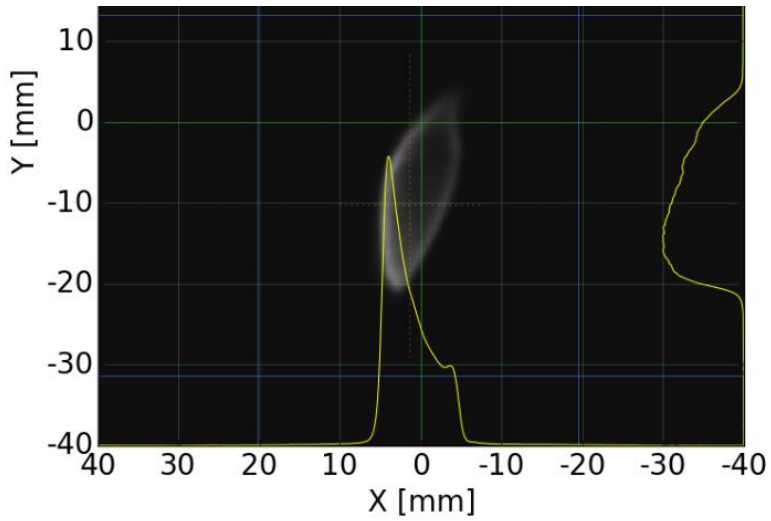

The space-charge effect of "short-focusing" clearly makes the spot far wider than it would otherwise be, hence making it more difficult for the slits to filter neighboring peaks.

This point becomes clear when the magnets are set for a $^{129}Xe^{26+}$ (M/Q = 4.957) beam:

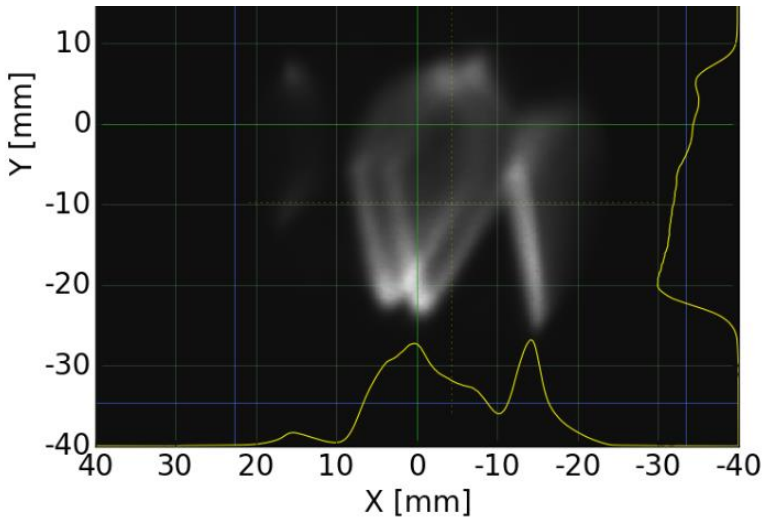

Closing slits will not block the contaminating ions, most likely $^{40}Ar^{8+}$ (M/Q = 4.995). (Use of beam with integer or half-integer values of M/Q is generally best avoided – the potential contaminants list often becomes quite lengthy for such values. And, again, values corresponding to charge states of oxygen, nitrogen and carbon should also be, if possible, avoided.)

Visible, but relatively mild in the preceding pair of pictures, is that these beam spots can be tilted. So far, the optics of SCS1, which were developed to give the correct orientation for about 50 eμA of $^{40}Ar^{9+}$, aren't changed for other beams and other intensities. As the vertical beam extent tilts into horizontal, resolution decreases. For this, as well as more complex reasons, getting the best resolution from a CSD for particle species identification and/or cleanup may also require closing of the vertical slits too, as shown in the these beam pictures of an argon beam:

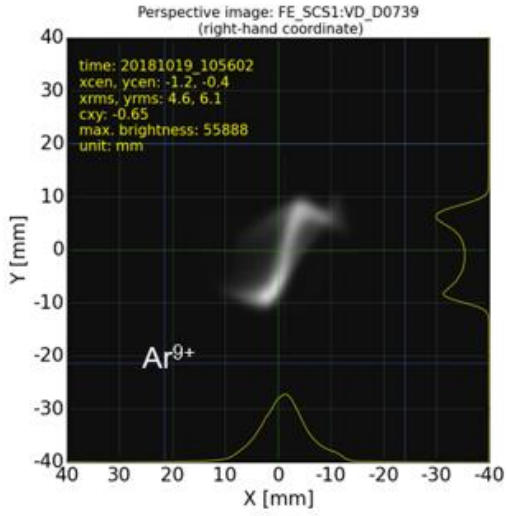

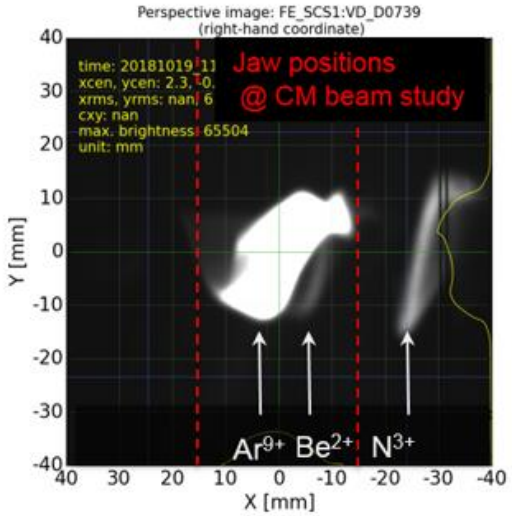

On the left, the $^{40}Ar^{9+}$ beam is shown, but contaminants are not visible. With the exposure set to show these contaminants on the right, the argon is too bright, but the upper 'hook' overlaps the $^{9}Be^{2+}$. Closure of the vertical slits would, in this case, help, and even more so if the spot were further tilted.

Further development is needed to improve this situation. Presently, slit setting around a typical well-behaved beam will be about 12 mm wide horizontally and about 30 mm vertically, passing 95% of the desired beam's maximum measured intensity. A tilted condition will force the horizontal to be wider and allow the vertical to be narrower for the same transmission. Slit settings for low-intensity and solid-feed beams can be narrower without significant losses of the desired ions, and generally will have better resolution. The optics will have to be somewhat different to achieve this however; a simple scaling between low- and high-intensity optimized settings will not give the best results.

Operation – Contaminants – (Acceleration)

As hopefully the previous section made clear, even when all is perfect, there are situations where beam selection by M/Q is insufficient to eliminate undesired ion species from being transmitted downstream. Depending on the degree of mismatch with the central tune, the undesirables stand a good chance of being lost somewhere before reaching full energy and striking the production target. The 'somewhere' is important. Losses at energies below the Coulomb barrier for nuclear reaction, about 1 MeV/u, are far less important than losing the same particles at energies higher than that. Heating effects from lost beam increase with power and, above the Coulomb barrier, immediate and induced radioactivity become a major factor. If high-energy losses cannot be avoided, it's best for those losses to happen in a controlled way, which requires knowing which, if any, contaminants will be accelerated along with the desired ion.

Cyclotrons have a great advantage over linear accelerators in this regard. First, the multiple orbits create essentially a very good spectrometer, separating ions by M/Q to a resolution of less than 1 part in 5,000. Second, as in the K500, loss of the undesirable beam injected from the ECRIS happens at low energy and is localized within the machine and its vault. A second filtering happens at an energy of about 10 MeV/u with a stripping injection into the K1200 to a different M/Q, followed by a second 1 in 5000 filtering during acceleration.

The question of whether a contaminant beam will accelerate in LS1 (to ~20 MeV/u) is both more difficult to answer and more important to answer. The importance is dictated by a long-term beam loss limit of 1W per meter. (The higher the energy of the particles, the fewer allowed to be lost). Eventually, there will be an app to correctly answer that question, but a partial examination of the $^{40}Ar^{7+}$/ $^{82}Se^{14+}$/ $^{12}C^{2+}$ case discussed above is illustrative and can be approximated using the "Physics Calculator" from the "FE Launch Page".

(This layout sketch may be useful for orientation …)

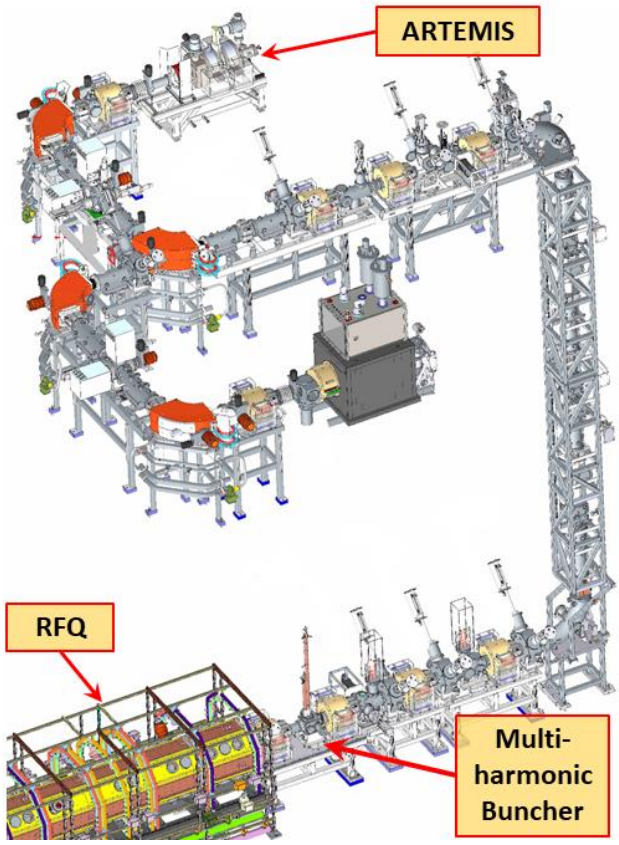

All three beams will reach the entrance of the RFQ, since the two contaminants are about 1.2% in momentum away from the main beam and the acceptance of the beam line is about 5%. The source + platform voltage is set to deliver 12.013 keV/u $^{82}Se^{14+}$, and the two contaminants each have a different energy and velocity as shown in this table:

Ion	M/Q	HV [kV]	Energy [keV/u]	Velocity^{-1} [nsec/m]	MHB to RFQ [nsec]	Δ RFQ Arrival Time	Δt [RF deg]
$^{40}Ar^{7+}$	5.708	70.284	12.312	649	779	-9	-260 (+100)
$^{82}Se^{14+}$	5.851	70.284	12.013	657	788	0	0
$^{12}C^{2+}$	5.999	70.284	11.715	665	798	10	+290 (-70)

Now examine this plot of phase and energy acceptance of the RFQ measured by A. Plastun:

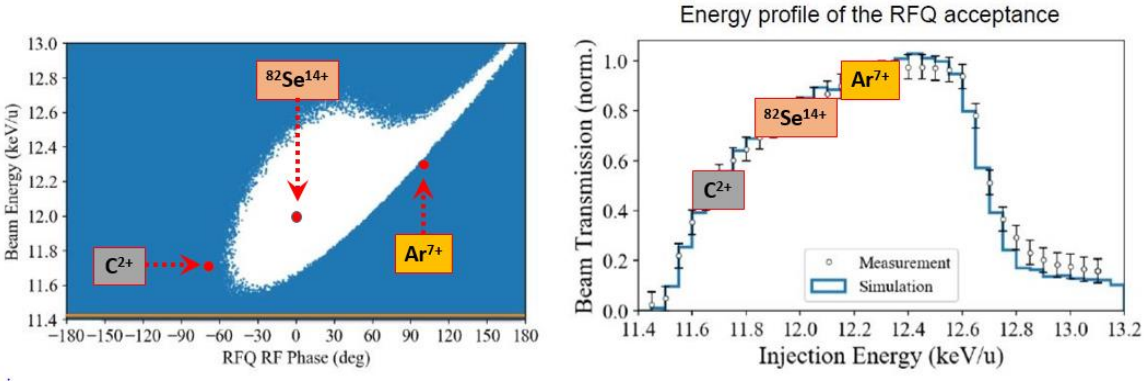

Looking only at energy and where each ion fits on the right-hand plot, one can see all 3 ions fitting into this energy acceptance window of the RFQ and being accelerated to 0.5 MeV/u, albeit with different rates of transmission. (Yes, considering only a DC, not-bunched beam, a wrong ion has a higher transmission percentage than the correct one!)

However, 1.2 meters before reaching the RFQ, the incoming DC beam is bunched by the *Multi-Harmonic Buncher* (MHB). The buncher phase is set to deliver the correct beam longitudinal compression into 40.25 MHz pulses at the point of entry to the RFQ. That phasing depends on the *time-of-flight* (TOF) of the particle group in reaching the RFQ. Because the contaminant ions have velocities different than the desired beam, they reach the RFQ at the wrong time and either may not accelerate or accelerate with reduced intensity. An illustrative first approximation can be done in the remaining table columns and is shown with the red dots on the left-hand plot, above. While it appears that the $^{40}Ar^{7+}$ and $^{12}C^{2+}$ both fail to fit into the RFQ phase vs. energy acceptance window, the dots only represent a centroid of a much larger pattern.

Based on simulation, this plot is a much better and more complete representation of the situation:

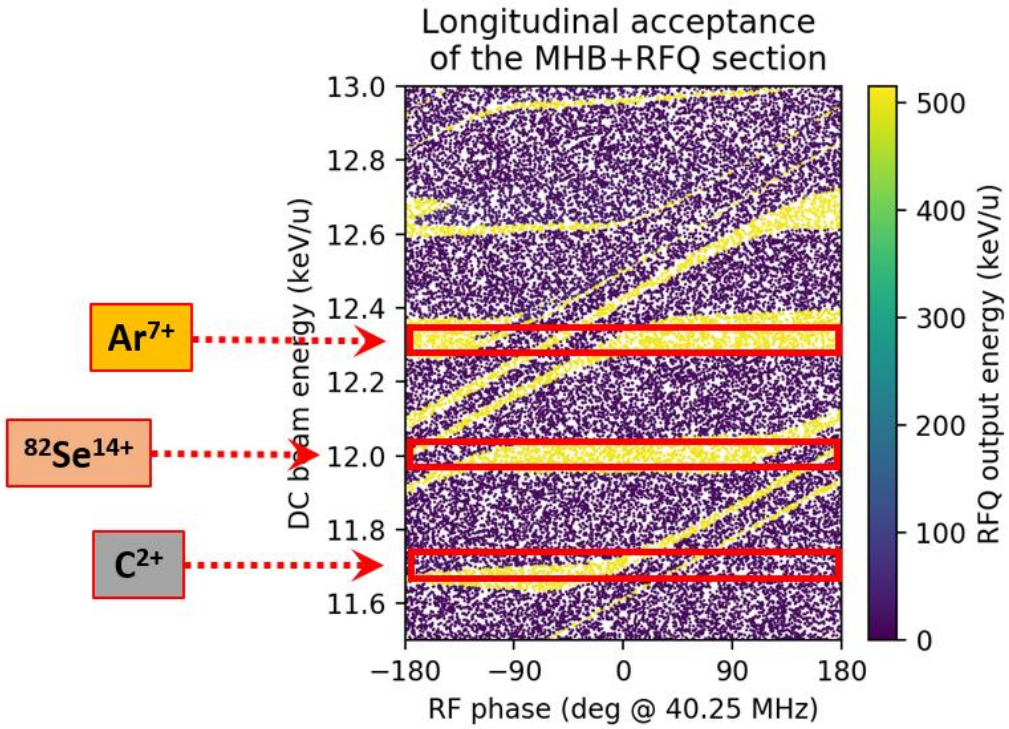

Using this plot, all one needs to know as "input" is the ion beam energy per nucleon. Moving horizontally across, the zones represented in yellow will accelerate through the combined MHB and RFQ system. Those in purple will not. The net transmission of each ion is then roughly the percentage of yellow in each 'box' compared to the total area of the box. (The energy spread shown here, i.e. the height of the box, is exaggerated for clarity). So, here, about ¾ of the incoming selenium, 2/3 of the argon and 1/6 of the carbon is accelerated by the RFQ.

An actionable summary of the discussion on contaminates is:

1) Set up SCS1 to give adequate separation between ion species at the slits.
2) Close the slits around the desired beam to cut away undesired species.
3) Use the CSD to identify contaminants that may pass the slit cut.
4) Notify the Accelerator Physics Group of these possible contaminants.

Operation – Beam Stability

Perhaps too obvious to state, but ion source plasmas are fundamentally unstable. Far from thermal equilibrium and consisting of charged particles wishing a quick return to neutrality, their continuance depends on externally provided microwave power and several different restoring mechanisms. The physics of plasma instability is both interesting and wildly complicated, the details of which are beyond the goals of this guide.

However, I think that it's possible to make useful categorizations of these instabilities (and what to do about them) based on their observed time characteristics.

Instability Generalized:

A) Long term (minutes)
B) Macro – Hz to kHz
C) Micro – MHz
D) Mode "switching" or "notching" – fast transition, long duration

What one observes as stable beam output depends on how the beam current is observed. Beam may appear to be stable in one view, but not so in another. An older, simple display method is pictured below.

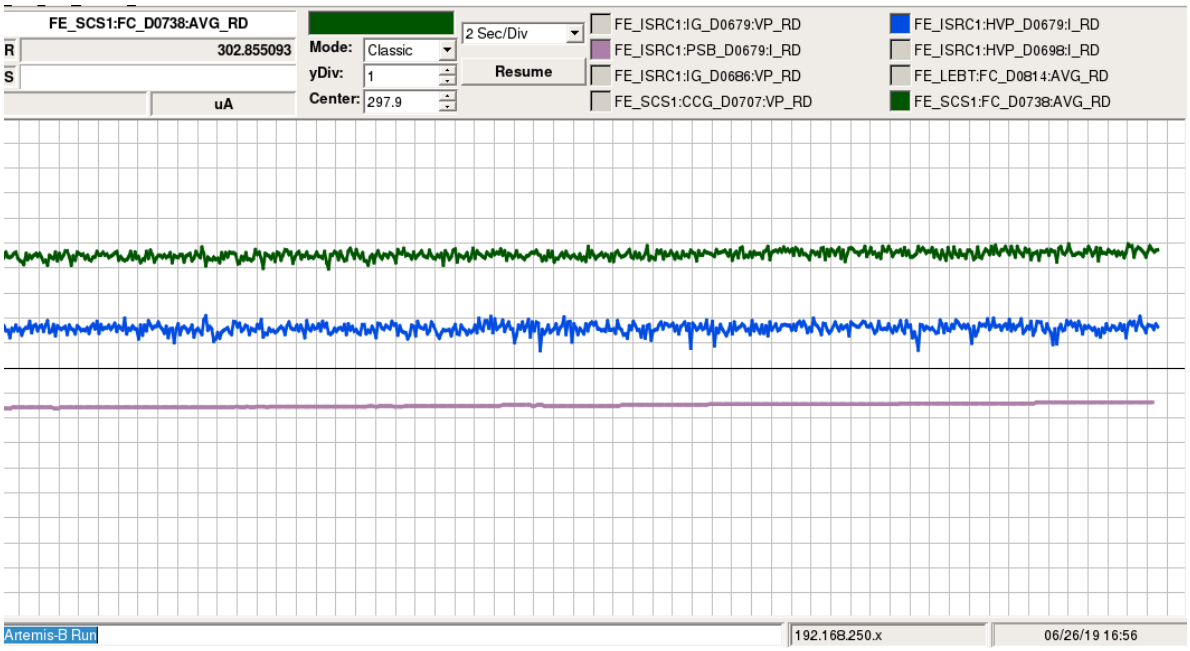

The green line represents an averaged D0738 beam current `[FE_SCS1:FC_D0738:AVG_RD]`, with a vertical scale of 1 µA/division, horizontal scale of 2 sec/division and 100 ms sampling. "Stable" yes, but much more information is available by using what amounts to a good oscilloscope built into the FRIB diagnostics system. From this overview page, one can select which Faraday cup to view.

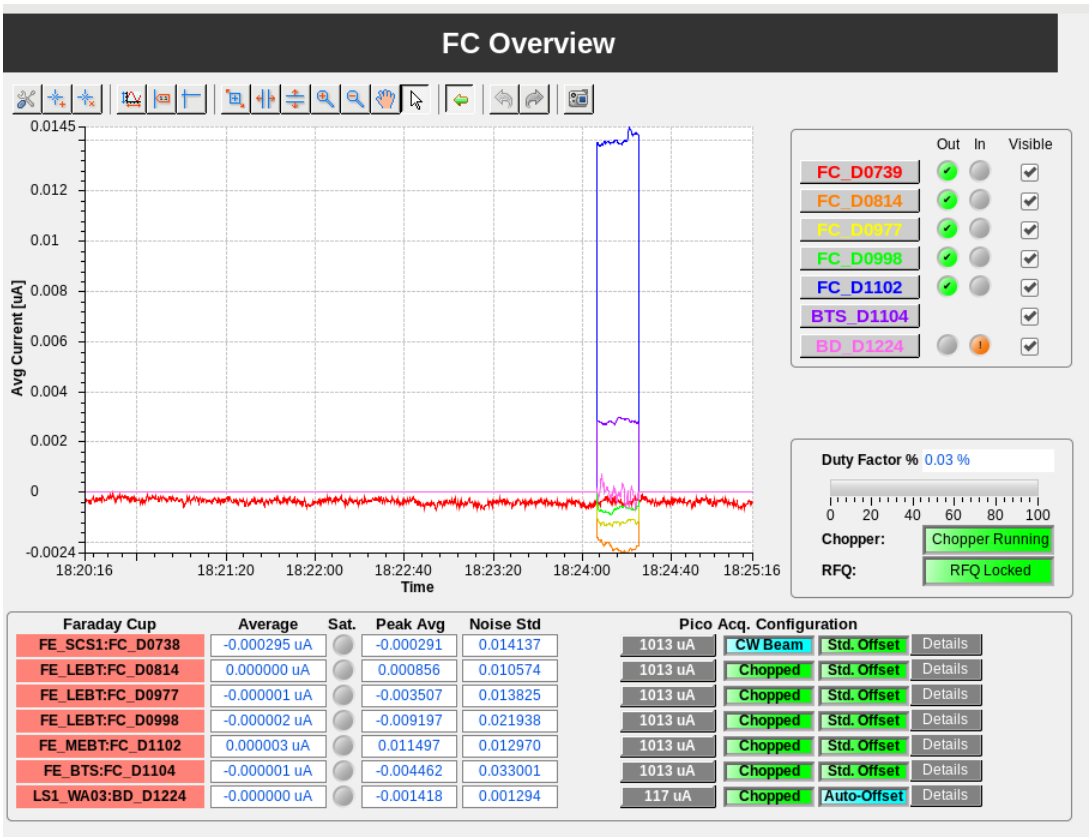

Selecting an individual Faraday cup here takes one to a “scope display”. The equivalent of the above simple chart trace is shown here:

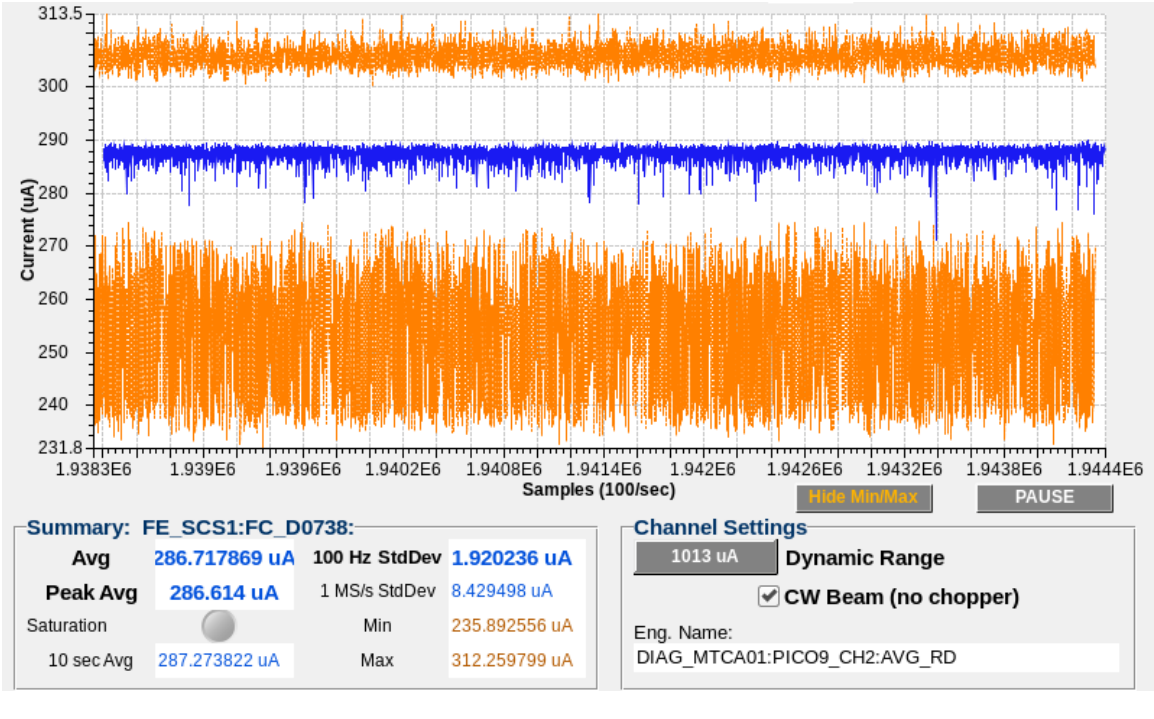

The blue line is the averaged beam current, similar to the previous green line, but with more visible detail. The orange lines are new, showing the maximum and minimum currents seen during the average period, sampled at ~1 MHz. These bands are typical, getting closer to the average at low microwave input power and spreading with increasing power, among other influences.

The condition shown above isn't terrible, but could, by some tweaking of source parameters, be improved upon as shown below:

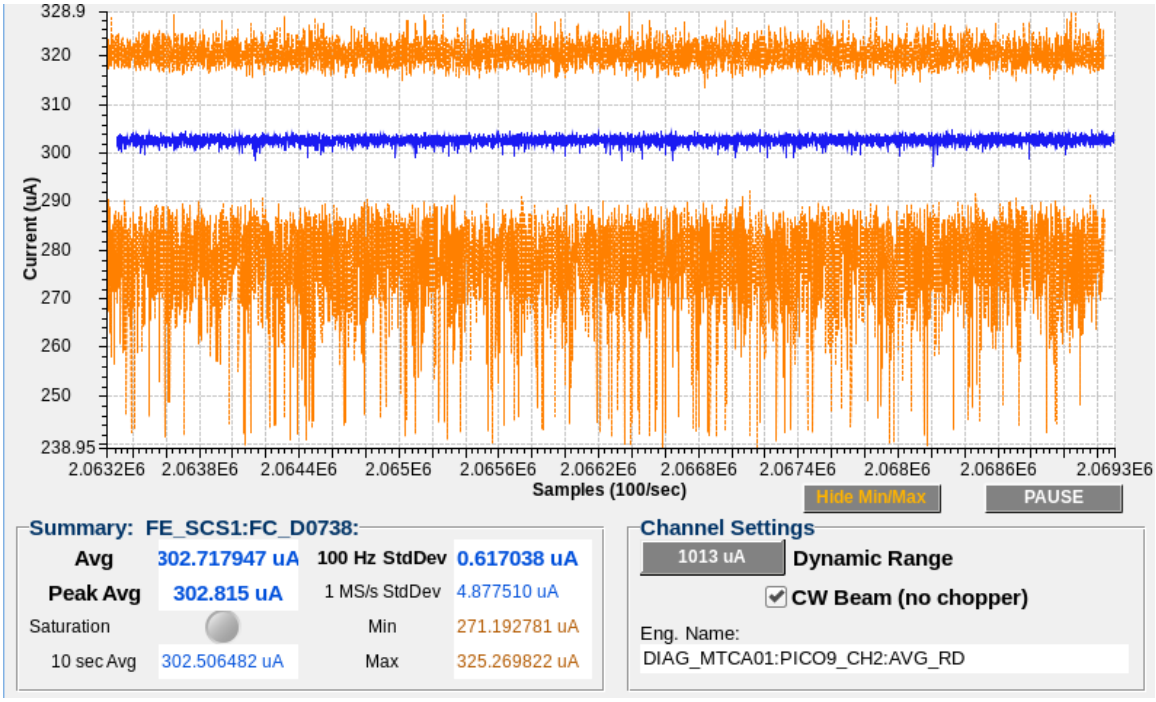

Auto-scaling can make interpretation a little tricky, but the number of fast current dropouts to a low value has clearly been greatly reduced as have the spikes appearing in the average current line (blue).

If one isolates a single dropout and examines it on a shorter time scale, one can see that it happens very fast. The “scope” diagnostics allow one to do this and take a snapshot of captured events. For example, this one:

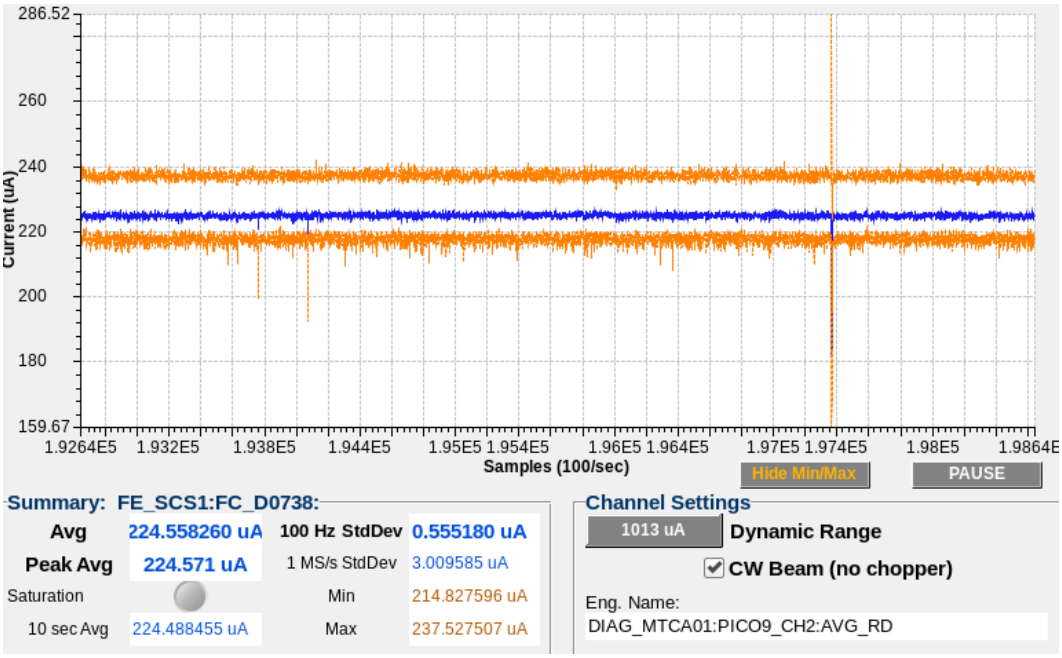

- becomes something like this on 100 µsec time scales:

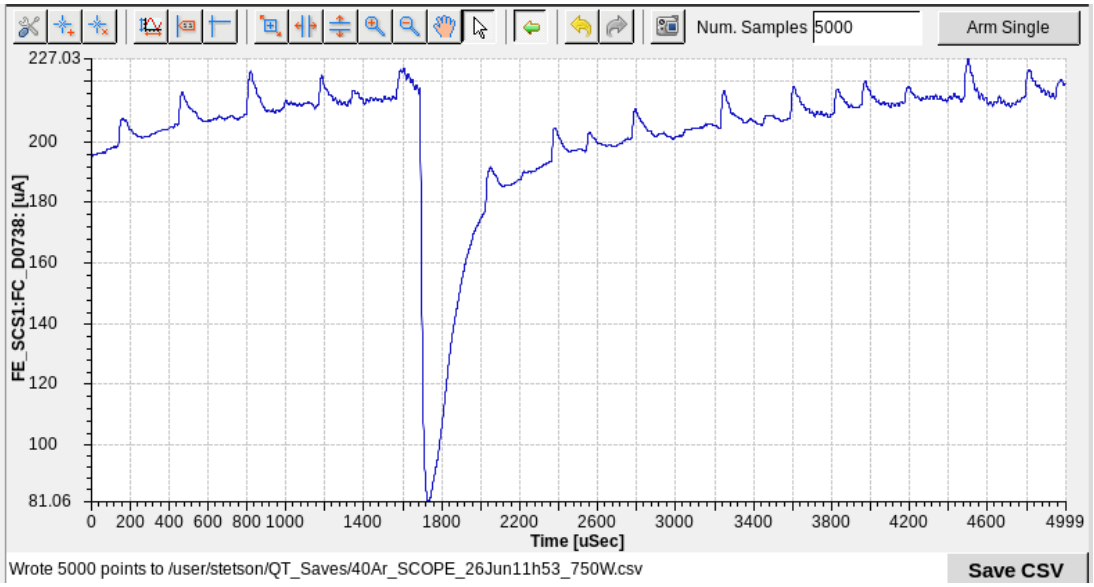

The drops occur within a few microseconds. However, because the standard Faraday cups are a hunk of metal, they have capacitance resulting in a cutoff frequency response of about 100 kHz. (Specifically designed low-capacitance devices are used at a few locations downstream in the linac to increase the limit, the so-called fast Faraday cups.) The dropout shown is thus, in reality, both deeper and faster than can be responded to by the measurement hardware.

But, besides aesthetics, why bother with such fine-tuning if the beam output is basically stable? The biggest reason is that these fast “insignificant” instabilities tend to increase as the tune shifts closer to significant mode shifts. Adjusting operating parameters to reduce instabilities usually leads one away from falling from a cliff from a high to low output mode.

In this next example, the condition in the middle is better than the condition to the sides, even though the averaged current appears to be about the same:

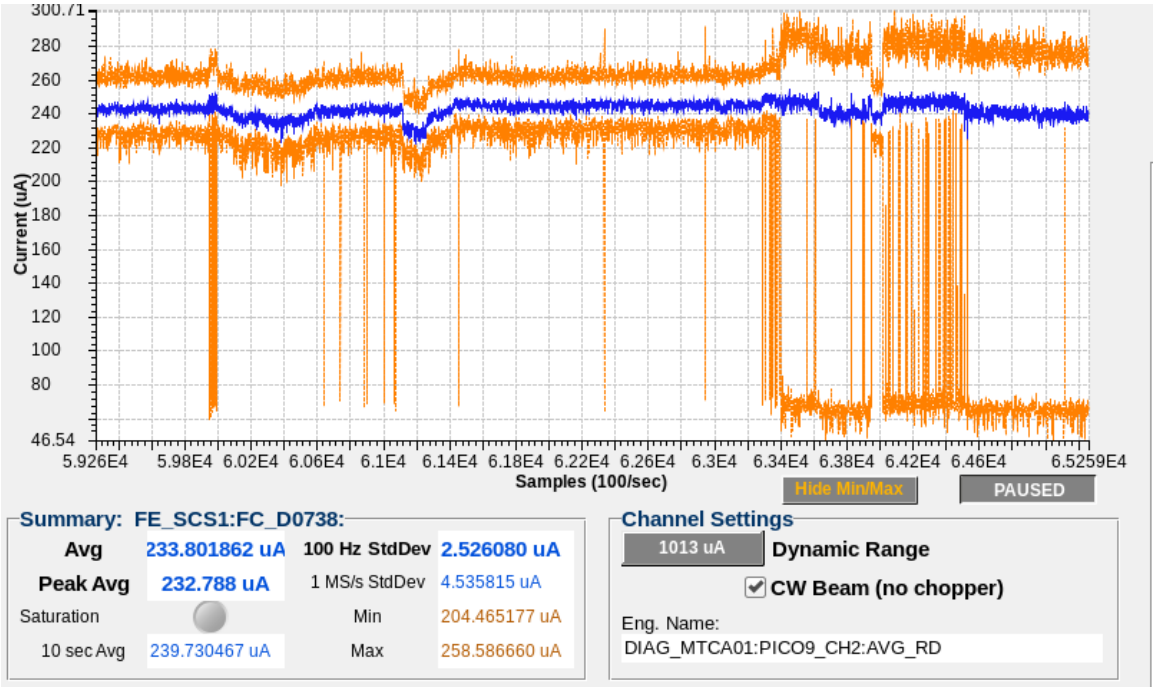

In contrast to these chaotic “spike” instabilities, the source can exhibit periodic collective oscillations. Periods of serious oscillations seen can range from roughly 50 Hz to 25 kHz. Consider this example showing a transition between two states:

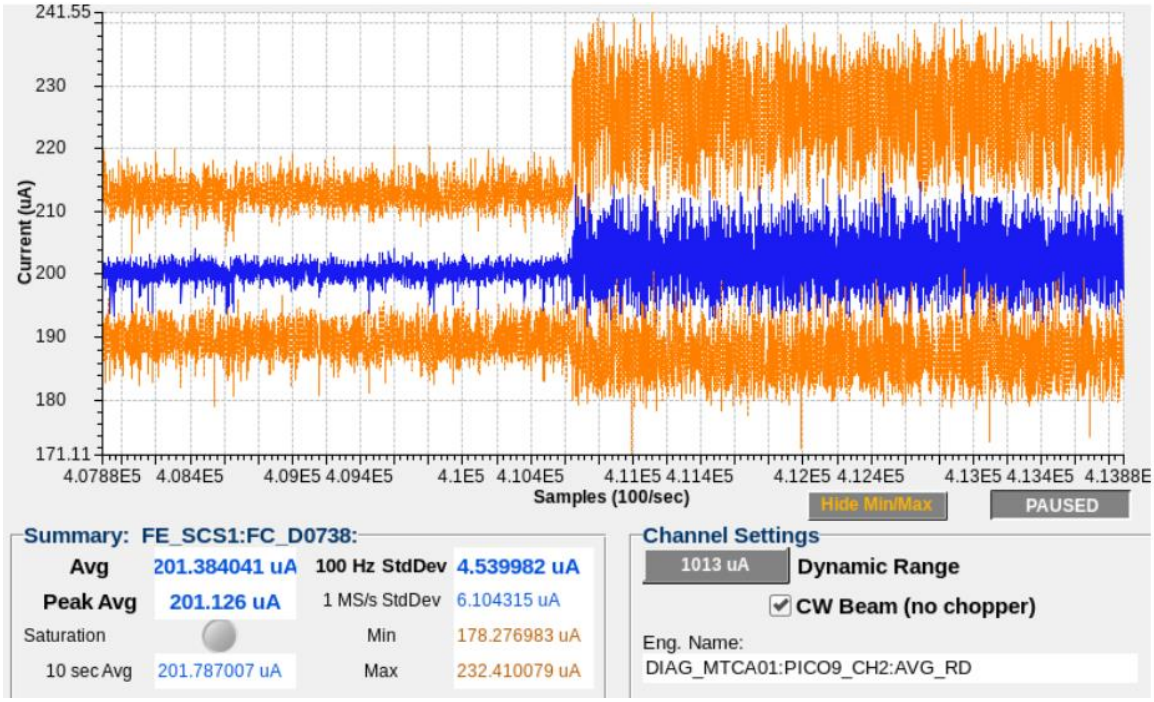

The better condition on the left above corresponded to a 160 msec sample shown here:

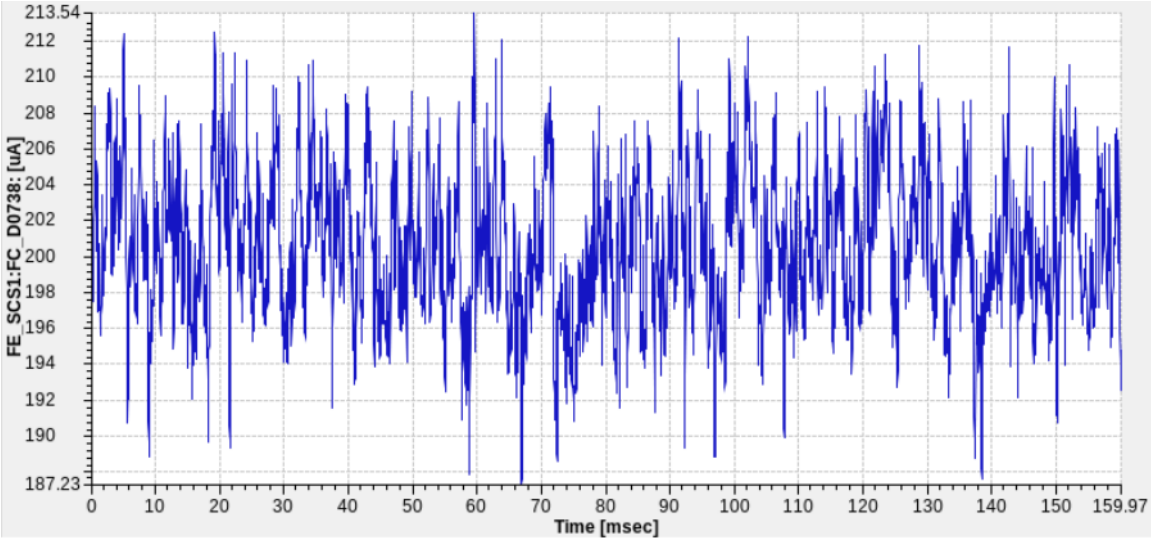

As always, the beam is not totally stable, but variations are small and stochastic. A change in bias disk voltage created a distinct transition into the second condition with this very different corresponding 160 ms sample:

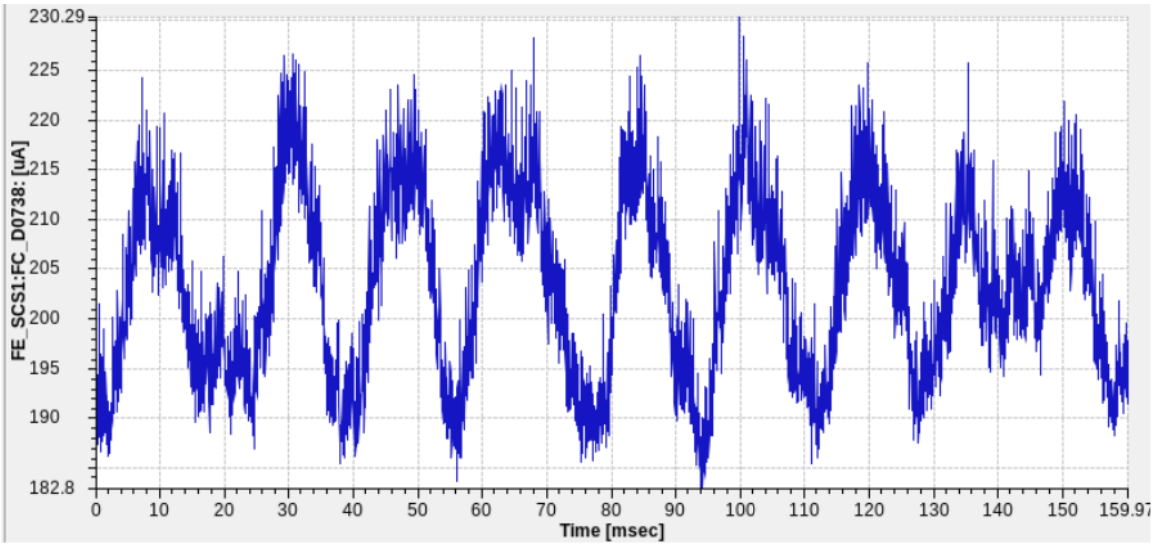

Not only is the range of variability somewhat larger, a ~55 Hz oscillation was clear and represents a situation best avoided. (While in the early commissioning period of FRIB beam power are low, that will change. At high power, beam loading on the SRF cavities is considerable and affects the cavity tune. Strong periodic source beam instability may therefore create issues in the linac downstream.)

The form of such periodic underlying oscillations varies greatly and defies simple explanation. This ^{40}Ca beam for example show a superposition of both a large oscillation around 130 Hz and a smaller one of about 1.2 kHz.

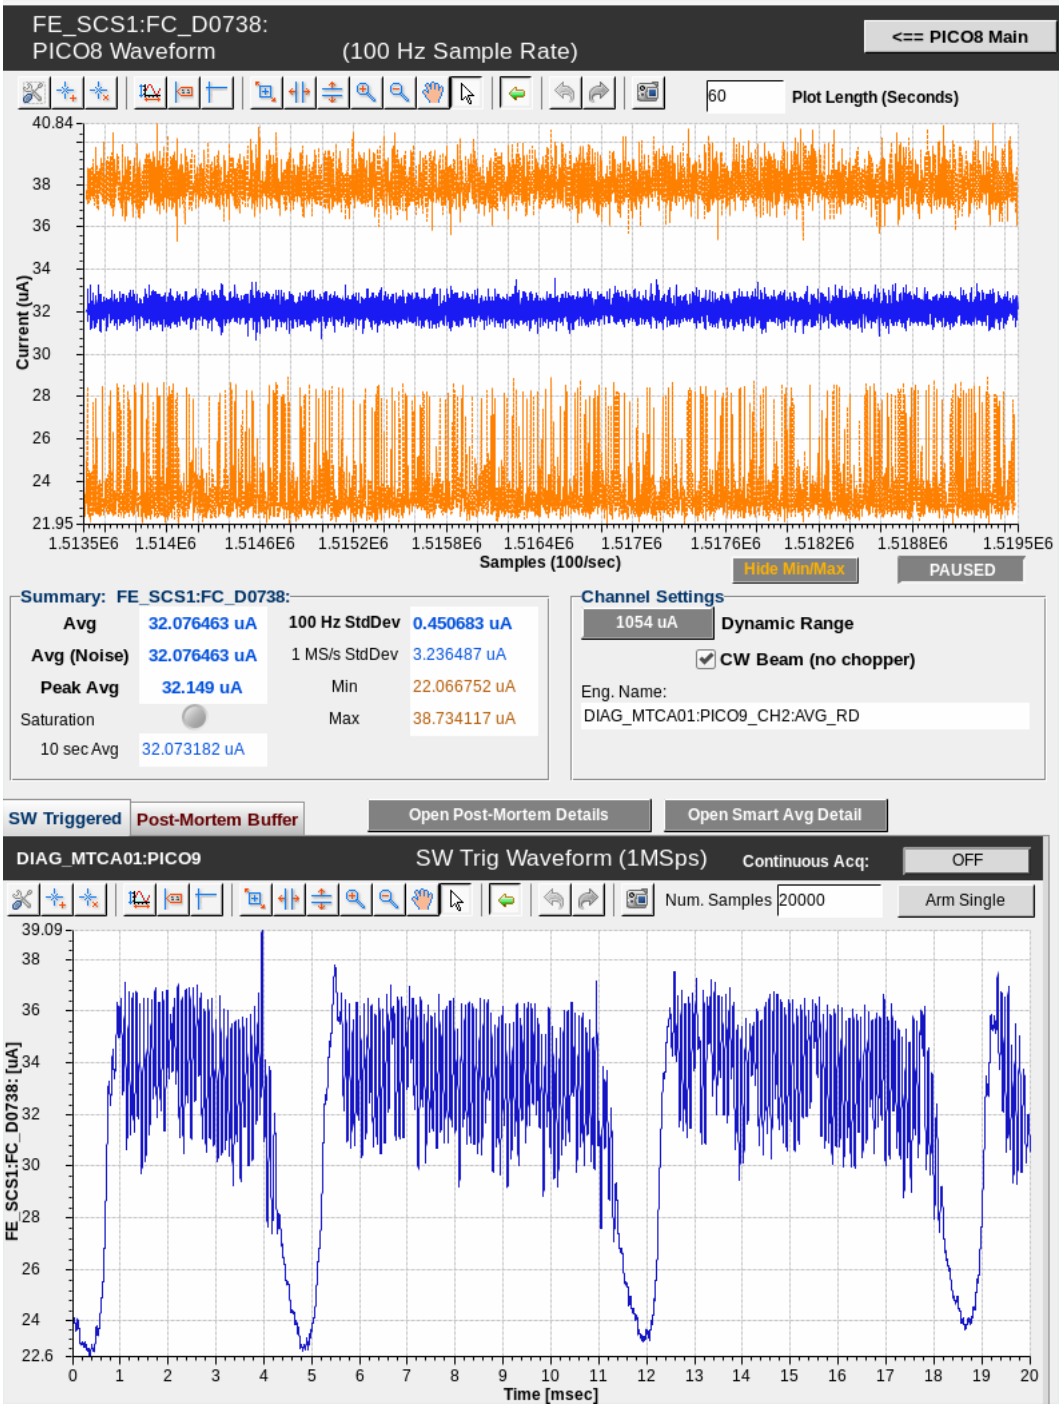

[Aside: Since many such oscillations fall within the range of hearing, this brings up a curiosity from the good old days: the use of a pickup antenna, amplifier and loudspeaker to literally tune sources by ear. Adjustment of source parameters until tones are minimized, ideally leaving only static, was a good/fast/cheap way to reduce or eliminate these periodic beam variations.]

It should be said that often one accepts somewhat reduced beam intensity in exchange for a more quiescent beam and plasma condition, as in the following example. In this sequence, the bias disk voltage is changed in increments from -23 V to -29 V.

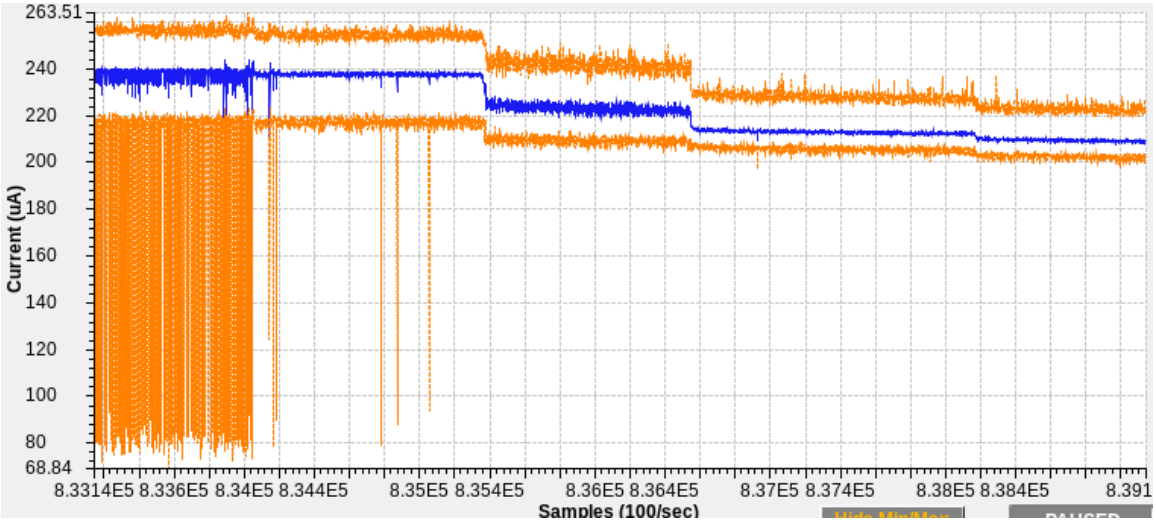

Why continue past the first curative step? Because at minimal cost of intensity, it takes one further away from the onset of that bad initial instability, allowing a greater tolerance for inevitable long term drifts in temperature, pressure and such. One wants a little extra distance while walking near the edge of a crevasse.

Another type of often-seen instability is a “notching” behavior of output beam intensity, repeating every few seconds, but with a somewhat ill-defined periodicity. While the beam current changes are often, at first, small, they usually indicate that the tune is approaching a condition where either there is a full mode shift to a much lower output or a degradation into more rapid, larger and severe instabilities. Although many factors could be involved, a first and easy reversible step is to iteratively adjust the confinement solenoids by a couple of amps and observe the effect. Shown below is an example of the effect on notching behavior in a ~18 µA beam of setting the injection-side solenoid current from 510 amps to 505 amps and returning to 510 amps:

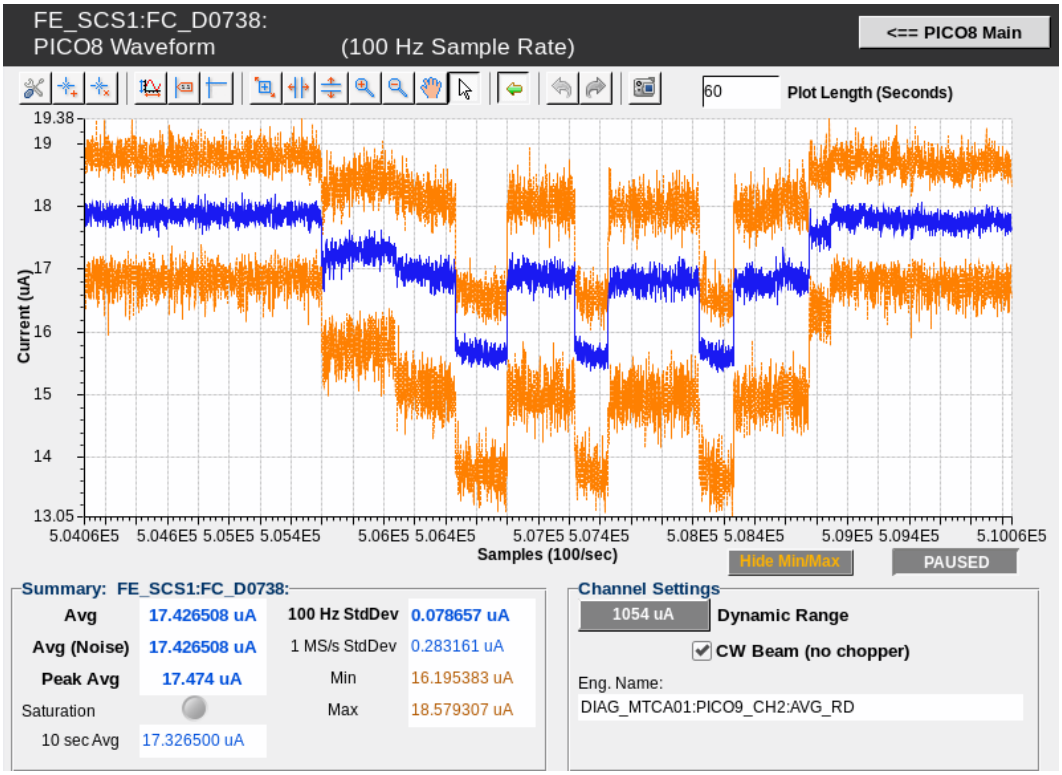

Instability or drifting over much longer times is also common and virtually certain over the first 4 – 8 hours after starting the ion source. An example over a ~2.2 hour duration is shown below for a $^{40}Ar^{9+}$ beam desired to be of 50 µA:

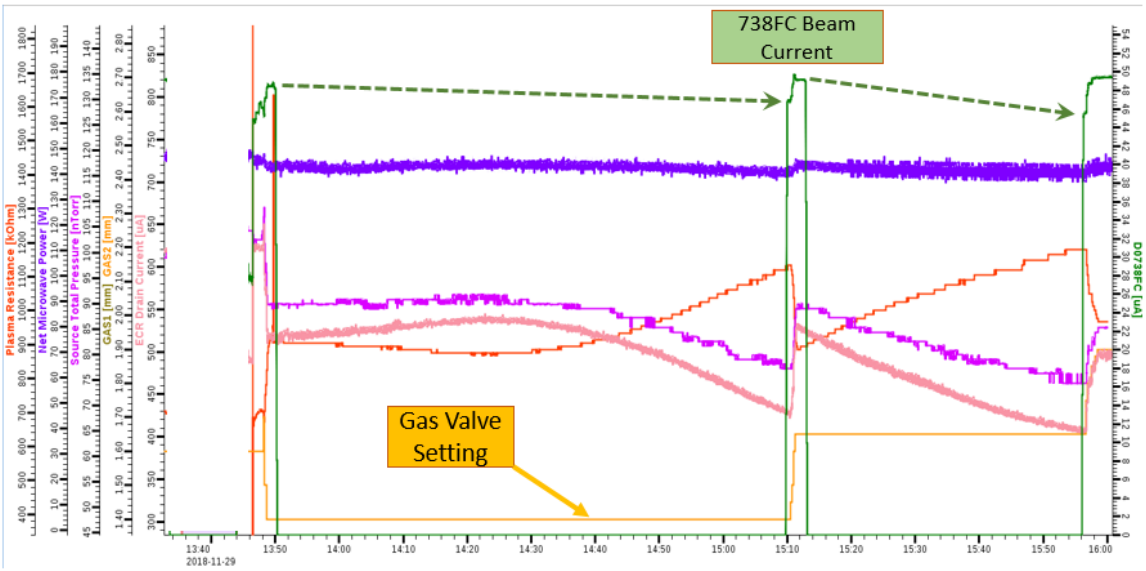

At 13:50 beam output has the desired value and the Faraday cup is removed to send beam downstream. However, read back source parameters, pressure (violet), bias disk current (red) and ECR drain current (pink) are seen to drift over the next 1.2 hours. The cup is inserted, gas valve opened by 0.2 to recover the original condition, then removed again. These drifts continued, and the process is then repeated. One can note here that smaller more frequent gas flow changes, keeping the drain current, bias disk current and pressures more constant likely (but not ALWAYS) would have kept the beam intensity more constant even without interruption of beam delivery by inserting the Faraday cup. At a minimum, drifts in any of those three readback values are a clue that beam output is also changing and/or the source tune may be approaching instability.

It's generally better to make several small corrections rather than waiting for these slow, long shifts to become substantial, possibly introducing other instabilities which may not be so easily remedied.

Closeout

Hopefully, this guide provides a useful beginning framework for those new to ECRIS operation. As a sub-field of both atomic and plasma physics, it's both an interesting and deep subject. Additional reading and exploration of the subject is recommended.

Though important, missing from this guide at present is a detailed discussion of the concepts of emittance within a context of beams that are both X-Y correlated and containing significant spherical and 2nd order aberrations. Perhaps an appendix on the subject can be added in the future.