Momentum spread and optics at the UITF

Preliminary results and mysteries

M. Bruker*

December 2021

This document summarizes why a two-week beam study took me two months and still raised more questions than it answered.

Many thanks to (at least): Matt, Joe, Tomasz, Rama, John M., Keith, Tom, Dennis, Jay, Marcy, Xi, Yan, Michael T., Yves, Sarah.

*bruker@jlab.org

Contents

1 Introduction

The UITF has always been set up somewhat empirically to suit the needs of the experiments. To a large extent, this is due to the available models being incapable of predicting the beam parameters for reasons that have not been understood. Unfortunately, none of the experiments so far have depended on a predictable and stable beam, which means we were never forced to put in the effort needed to get to the bottom of things. With the harmonic kicker test and potentially other applications on the horizon, we have been trying to gain at least some understanding of the straight-ahead MeV section of the machine, which proves to be complicated enough.

The original purpose of these studies was to measure the full phase space at the exit of the booster as a function of energy (save for the bunch length, which cannot be measured in the current setup) to optimize our model to the point where it would be able to predict the beam parameters at the experiment. While this idea is still worth pursuing, it has turned out to be less simple than we had imagined. What follows is a summary of what we pretend to understand and, more importantly, what we admit we (i.e., I) do not understand.

The beam line segment relevant to these studies is shown in Fig. 1; it is designed for beam energies up to about 8 MeV as provided by the booster cryomodule. Refer to the quick reference drawing [1] for more details.

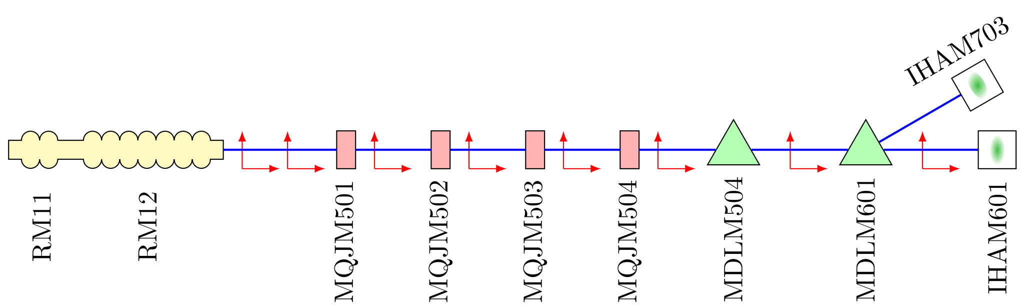

Figure 1: Relevant components of the MeV beam line.

2 Momentum measurements in the MeV spectrometer line (M700)

The M700 "spectrometer" beam line is used to measure the absolute central beam momentum by deflecting the beam with a dipole magnet; in addition, it also serves to measure the relative momentum distribution around this central value by transforming it into a distribution of transverse displacements in a predictable way. This distribution, however, is convolved with the intrinsic transverse shape of the beam as determined by the optical functions and the transverse emittance. Unless the intrinsic transverse width can be made negligibly small compared to any dispersive contribution, deconvolution will be required, which needs the intrinsic width as an input; however, it cannot be measured directly because momentum spread cannot be switched off in real life[1]. Accurate measurements of the momentum distribution are therefore contingent on an optical model of the beam line and a measurement of the transverse phase space before the dipole. With the quadrupoles MQJM701 and MQJM702 degaussed, the rest of the spectrometer beam line can be considered a drift; the dispersive and multipolar contributions from correctors and stray fields (e.g., ion pump and earth) are usually neglected in the hope that they are not the end of the world. Whether or not these assumptions are reasonable I cannot say with confidence.

Table 1 lists the distances between the relevant elements.

2.1 Dispersion and edge focusing

The DL dipole used for dispersive measurements bends by a nominal angle of 30°. Lately, there has been some confusion about what else this dipole does to the beam. The usual Elegant models do not seem to apply to it very well, seeing as it is not a sector dipole but a rectangular one that is mounted at zero angle with respect to the straight beam line. While not necessarily detrimental in any

[1] My intuition tells me one might be able to conceive some kind of parameter scan (e.g., quad scan) that, given η_x and ϵ_x (which are measurable), would yield α, β, and $\delta p/p$ in a single curve fit, eliminating the need to know β in advance. An optics expert can probably explain to me why this makes no sense.

Table 1: Longitudinal positions of the elements in the 700 line relative to the dipole, sorted by measured position. The measured positions refer to the centers of the elements as determined with a tape measure from the cabinet, which I cannot read very well because it is in prehistoric units. The "drawing" positions were extracted from [2]. Both correctors are at non-nominal positions.

Element	Type	Δs_{UED} (m)	$\Delta s_{\mathrm{drawing}}$ (m)	$\Delta s_{\mathrm{measured}}$ (m)
MDLM601	Dipole	0	0	0
IPMM701	BPM	0.4377	0.47	0.46
MLHM701	Corrector	0.6747	0.58	0.53
MQJM701	Quad	1.3953	1.30	1.34
MLHM702	Corrector	1.8394	1.87	1.50
IPMM702	BPM	1.7394	1.78	1.75
MQJM702	Quad	2.0367	2.09	1.98
IHAM703	Harp	2.2792	2.31	2.31
ITVM703	Viewer	2.3935	2.44	2.43

way, this arrangement is somewhat unusual and not covered in a lot of texts.

The optical properties of the dipole can be inferred from a GPT simulation. To my knowledge, no measured field map exists, much less for the particular copy of the magnet we have at the UITF. The field map I am using was simulated by Jay and is found at

`/a/opsuser/benesch/magnets/inj_maps/MDL_30deg.table ;`

it is depicted in Fig. 2 at $y = 0$. (x, z) is the bending plane.

First, to simulate the focusing action of the dipole, we send a bunch of finite transverse size and zero angular spread through the magnet and observe the final angle of each particle as a function of its initial displacement. Figure 3 shows the $x'(x)$ and $y'(y)$ dependencies; inter-plane coupling is negligible.

The physical reasons for the focusing action in x and y, respectively, are different: In sector magnets or rotated rectangular magnets, the focusing in x (the bending plane) is caused by a displacement-dependent difference in path length,

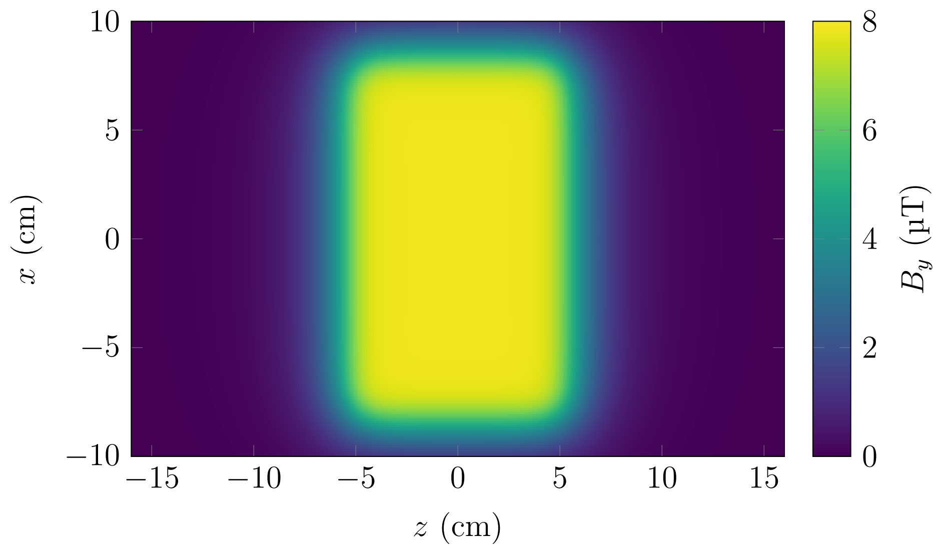

Figure 2: Field map of the unrotated DL dipole at $y = 0$ normalized to $\int B_y \, \mathrm{d}z = 1\,\mu\mathrm{T\,m}$ on axis.

altering $\int B_y \, \mathrm{d}z$. In the case of this arrangement, the path length would ideally be independent of the displacement, but the field is not, resulting in a tiny focusing effect with a focal length of about 14 m. Conversely, the focusing action in y (perpendicular to the bending plane) is not caused by subtle changes in the field integral but by fringe fields. This effect is much stronger and results in a focal length of 0.5 m, which, considering the much larger distance to the downstream harp (2.31 m), inevitably leads to a large spot size in y. While annoying, this is usually considered irrelevant.

The effects of a nonzero angle of incidence on the beam shape have not been studied.

Figure 4 shows the dispersion at IHAM703 in the bending plane:

$$\eta_{\mathrm{GPT}} = 1.3217(2)\,\mathrm{m}. \tag{1}$$

There is no dispersion in the other plane. The dispersion barely depends on the vertical displacement of the beam with a relative change of about 2×10^{-4}/mm.

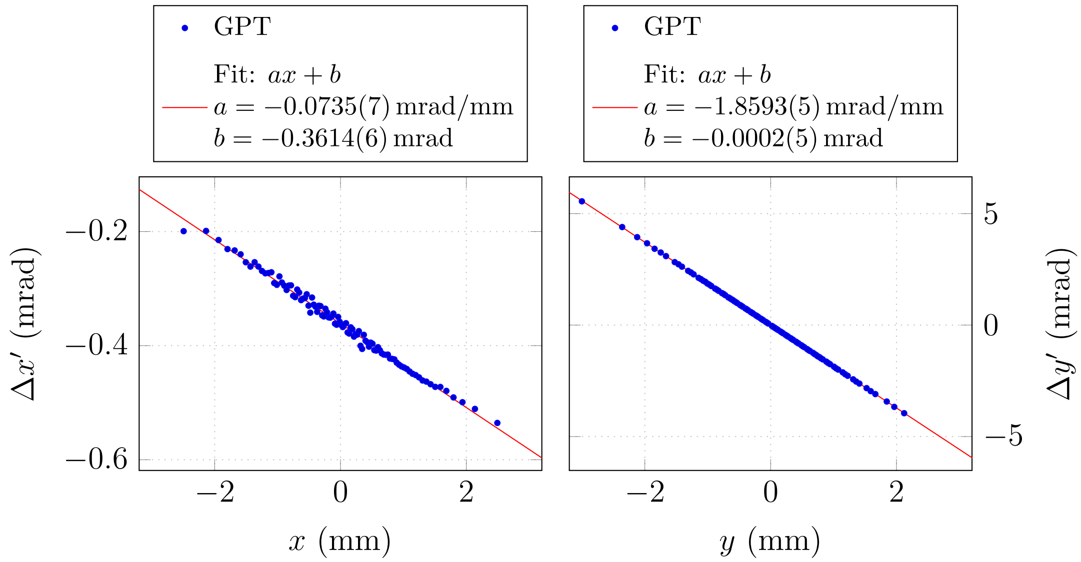

Figure 3: Transverse kick angle as a function of initial displacement as simulated with GPT. $p_0 = 8\,\mathrm{MeV}/c$, $\int B_y\,\mathrm{d}z = 13\,384\,\mu\mathrm{T\,m}$. In the absence of an initial angle distribution, the inverse of the slope of these curves is the focal length in the respective axis. The nonzero ordinate intercept in x is caused by a slight misalignment of the rotated coordinate system with respect to the deflection angle. As both analyses were computed from the same initial particle distribution, we also see a tiny $(y,\ x')$ coupling effect in the left plot.

All these results are independent of the central momentum (not shown). It is also somewhat less than inspiring to find that the value is close to what it would be for an ideal, zero-length, edgeless dipole:

$$\eta_{\mathrm{ideal}} = 2\Delta s \tan\frac{\theta}{2} = 1.24\,\mathrm{m} \tag{2}$$

assuming $\theta = 30°$ and the distance between the centers of the elements $\Delta s = 2.31\,\mathrm{m}$.

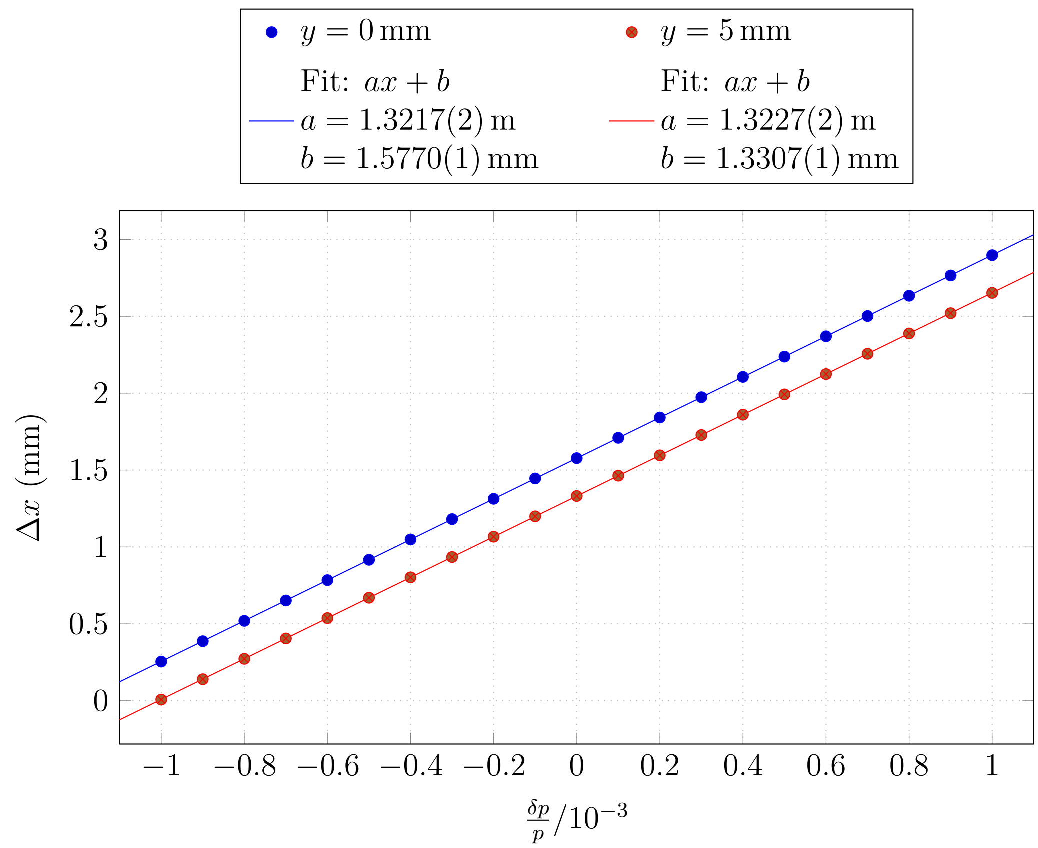

Figure 4: Transverse displacement at IHAM703 as a function of relative momentum deviation as simulated with GPT. $p_0 = 8\,\mathrm{MeV}/c$, $\int B_y\,\mathrm{d}z = 13\,384\,\mu\mathrm{T\,m}$. The nonzero ordinate intercept is caused by a slight misalignment of the rotated coordinate system with respect to the dipole in terms of both angle and position of its exit axis. We see that any vertical displacement of the beam is mostly irrelevant.

2.2 Beam-based dispersion measurement

2.2.1 Digression on viewer calibration

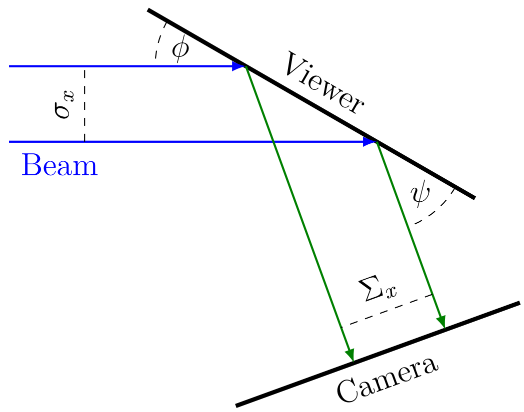

Figure 5: Beam/viewer/camera relationship with arbitrary angles in x. $\Sigma_x = \sigma_x \frac{\sin\psi}{\sin\phi}$. Nominally, $\psi = \phi = 45°$; then, $\Sigma_x = \sigma_x$.

The viewers at the UITF are mounted at a nominal angle of $\phi = 45°$ in x with respect to the beam line. This non-zero angle distorts the aspect ratio of the beam image on the viewer when observed in the viewer plane; the x axis is stretched by a factor of $\sin\phi$ ($\sqrt{2}$ for 45°), whereas the y axis is untransformed. However, in theory, if a camera looks at a viewer from an angle ψ in x with respect to the viewer plane (see Fig. 5), this transformation can be compensated for: The total beam image distortion in the camera plane is

$$\frac{\Sigma_x}{\sigma_x} = \frac{\sin\psi}{\sin\phi}. \tag{3}$$

In the case of $\psi = \phi$, i.e., the normal of the camera plane is perpendicular to the beam pipe, the aspect ratio of the beam is untransformed even though that of the viewer geometry itself appears distorted in the picture.

In reality, we generally have to assume both angles to be different from 45° and from each other[2]. ψ can be calculated from the image of the viewer itself (which is known to have equal x and y dimensions) projected onto the camera plane. We can either cross our fingers hoping for ϕ to be close to the nominal value or perform a beam-based measurement of it by applying a known displacement to the beam and using the resulting displacement of the image to solve Eq. 3 for ϕ.

The 703 viewer constitutes a convenient case in that it has fiducial marks on it that allow for easy calibration; see Fig. 6. The aspect ratio is $x/y = 0.62(2)$; therefore, $\psi = 38(1)°$.

Table 2: Results of the linear fits $a_i(x - \overline{x}) + b_i$ to the data shown in Fig. 8. The "corrected" rows include a factor of $1/\sqrt{2}$ to account for the 45° rotation of the harp in the (x, y) plane.

i			a_i		b_i/px
1	701H	Viewer	0.815(12)	px/(µT m)	94.8(6)
	701H	Harp (raw)	0.097(3)	mm/(µT m)	56.3(1)
2	701H	Harp (corrected)	0.069(2)	mm/(µT m)	
3	701V	Viewer	−1.072(4)	px/(µT m)	219.2(2)
	701V	Harp (raw)	0.087(1)	mm/(µT m)	25.0(1)
4	701V	Harp (corrected)	0.061(1)	mm/(µT m)	

Displacing the beam by a known amount is difficult in that it involves trust in either another piece of diagnostic hardware or a theoretical model. We are not yet at a point where we have much trust in either, especially seeing as the BPMs were not considered usable at the time these measurements were performed, but we can start by using the harp as a reference and checking the results for consistency. To this end, the beam was steered with the 701 corrector in both axes and the displacement observed on both the harp and the viewer. The results are shown in Fig. 8 and Table 2.

[2] At this point, I feel like more time has been spent talking about what the viewer cameras really show than it would have taken to design a viewer/camera assembly with predictable geometry.

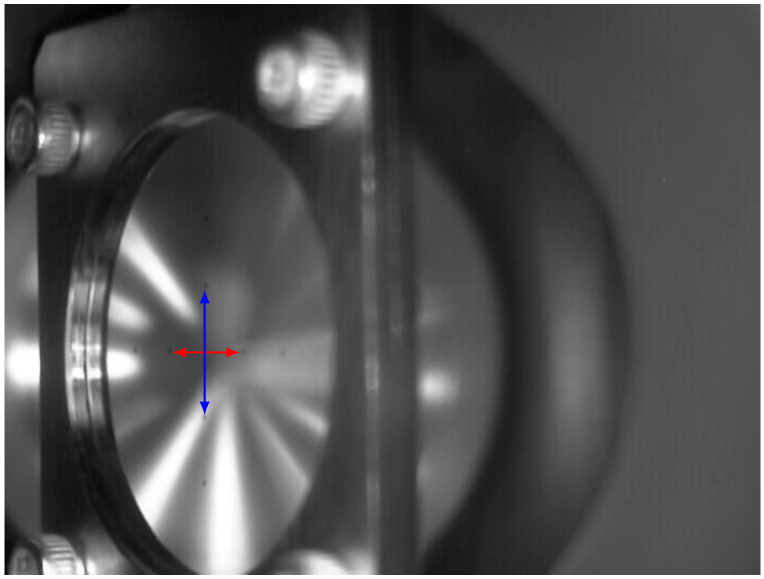

Figure 6: Picture of ITVM703 as seen by the respective camera (averaged over 200 frames to denoise it). The spacing of the dots indicated by the arrows is $\frac{10}{\sqrt{2}}$ mm in both x and y. This length corresponds to 69(2) px in x and 112(2) px in y (conservative error estimates). We shall pretend this coordinate system is unrotated.

Translating the y beam position on the viewer into physical units should be reasonably free of unpredictable errors; I would like the same to be true for both axes of the harp. From Table 2, we can determine the vertical scale of the viewer in a beam-based way, which is $a_3/a_4 = 17.5(1)$ px/mm. Accounting for the drift length between harp and viewer, which increases the displacement on the viewer by 7 % compared to the harp when steering with this corrector (see Table 1), the corrected value is 16.3(1) px/mm. According to Fig. 6, the true (geometric) vertical scale of the viewer is 15.8 px/mm; an unexplained discrepancy of 3 % is acceptable here.

Based on the harp results alone, it is interesting to see that the horizontal corrector is stronger than the vertical one by 11 % for the same alleged integrated

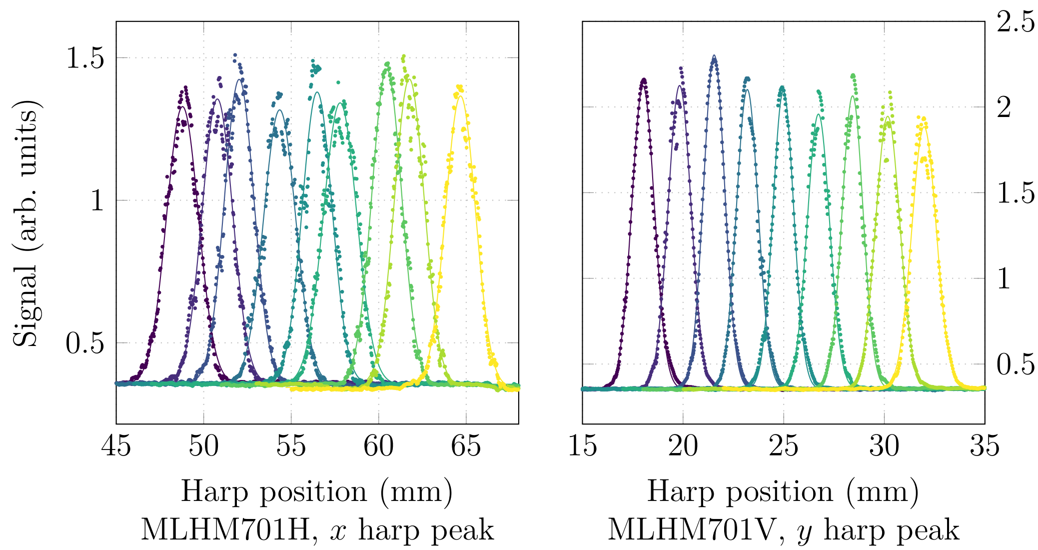

Figure 7: Using the 703 harp as a BPM to measure the beam displacement when steering with MLHM701H/V. Note that the 703 harp has an unintuitive orientation with respect to the beam line: x and y are swapped (here, the naming corresponds to the actual axes, not the HarpFitter convention, i.e., the x peak is at the right-hand side of the scan), and y is inverted. The colors correspond to the corrector settings shown in Fig. 8.

field. This observation has been made at the CEBAF injector before[3], though with smaller differences. On average, the correctors are about as strong as advertised: in theory, correctors should kick by 0.037 mrad/(µT m) at 8.0 MeV/c, which gives 0.065 mm/(µT m) after 1.75 m of drift. There is some uncertainty due to the non-zero length of the correctors changing the effective drift length in a way that is calculable but not too relevant now.

Assuming the calibrations of the x and y axes of the harp are in better agreement with each other than those of the corrector, the aspect ratio of the beam on the viewer is

$$\frac{a_1/a_3}{a_2/a_4} = 0.68(2), \tag{4}$$

giving an angle between viewer and beam line of

$$\phi = \arcsin\left(\frac{\sin\psi}{\frac{a_1/a_3}{a_2/a_4}}\right) = 65(6)^\circ. \tag{5}$$

Barring some gross oversight during assembly, this value seems excessive; Marcy says anything between 40 and 50° is to be expected as the assembly is not designed for precise alignment, but 65° should even be obvious to the naked eye. However, a large number of measured quantities with unexplored systematic uncertainties contribute to this number, and the condition of Eq. 5 is also somewhat unfavorable. I am willing to leave this discrepancy a mystery for the time being; the actual angle can be surveyed in due time.

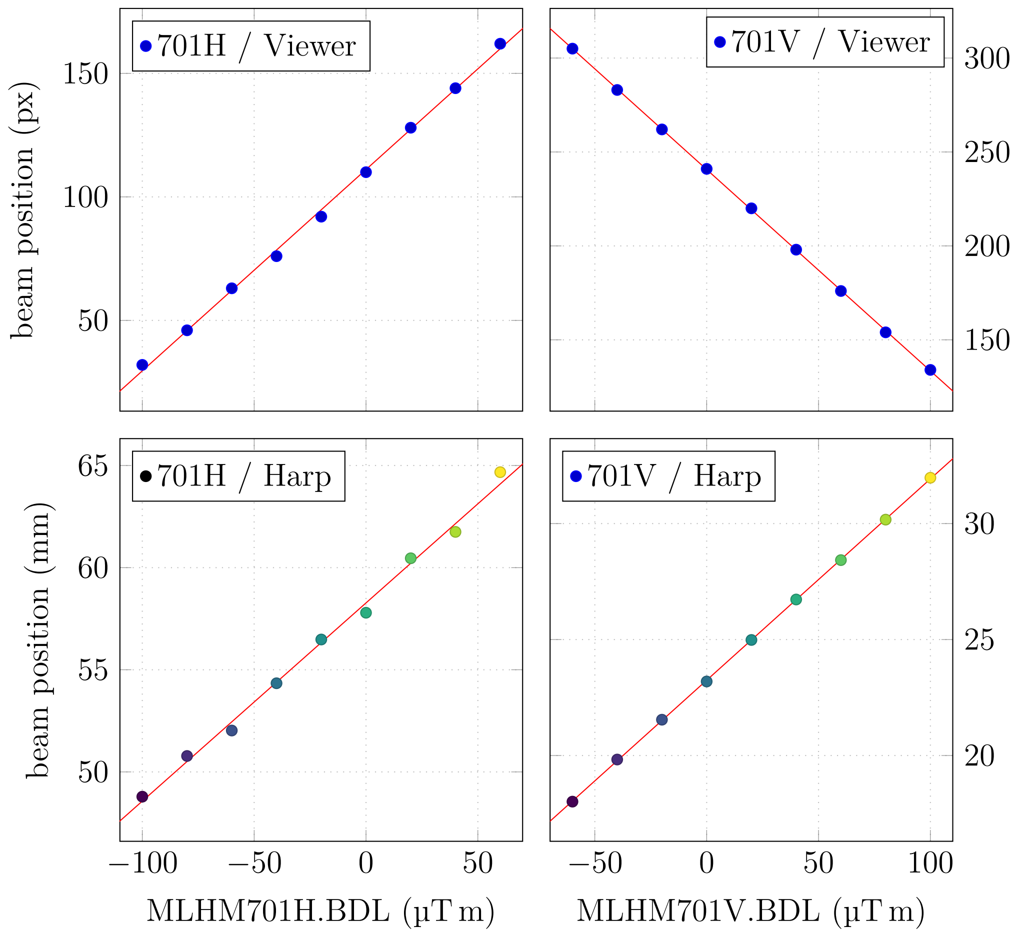

Figure 8: Beam displacement on viewer ITVM703 (top) and IHAM703 (bottom) in response to an angle change at MLHM701H (left) and V (right). The slopes of the lines are listed in Table 2; the colors of the harp points correspond to the data shown in Fig. 7. While the y axis of the harp is inverted, so is the 701V corrector (the latter was fixed after the measurement), which explains the opposite slopes of the viewer and harp lines in the vertical case.

2.2.2 Dispersion measurement

The dispersion can be measured directly by varying the beam momentum by a known relative amount and observing the beam displacement. To determine what change in cavity field will cause a certain relative amount of momentum change, we set the cavity field to a range of values around the nominal point and adjust the dipole to give the same deflection angle. This strategy avoids altering the trajectory through the dipole field and therefore only relies on the dipole calibration being correct for one particular deflection angle, and only in a relative sense. The result is shown in Fig. 9.

This measurement relies on the assumption that varying the beam momentum in this range leaves the entire orbit unperturbed, in particular the launch into the dipole. No effort was made to check this; the study should be repeated in a more methodically sound fashion by using BPMs.

The dispersion measurement is shown in Fig. 10 using both the viewer and the harp in the spectrometer line. Conveniently, the monitors measuring the beam displacement are the same ones used to measure the momentum spread, so the dispersion values so obtained do not need to be propagated.

It has been brought up that measuring the dispersion in this way involves a circular argument in that if one believed the dipole calibration to be accurate (which, for lack of a momentum reference in real life[3], can only be determined from simulation), the dispersion obtained from this same simulation would need to be accurate as well. It is all the more interesting to find that the measured dispersion (1.20 m, see Fig. 10) is 9 % smaller than the value computed from the field map (1.32 m, see Fig. 4). This measurement, however, includes all dispersive contributions between the booster and the spectrometer harp, not just the dipole. While one would assume the others to be negligible in comparison, this will need to be checked. It is also possible for parasitic focusing fields to exist between the dipole and the harp; for example, at the time of this measurement, the M701 and M702 quads were cycled at zero nominal field but not degaussed. The potential

[3] While it has been shown that gun-energy beam can be transported into the spectrometer line [4], allowing for an exact calibration in theory, systematic errors caused by the orbit and, in particular, the earth's field are hard to quantify and put an upper limit in the order of a few percent on the quality of such a calibration.

impact of these differences is yet to be evaluated (could be measured, simulated, or both).

What is even more interesting is that the viewer seems to have a different horizontal scale in this measurement than it did in the corrector-based calibration study described in Section 2.2.1. Assuming a vertical scale of 15.8 px/mm and $\psi = 38°$ (both from the fiducial marks as explained in Sec. 2.2.1), accounting for the drift between viewer and harp magnifying the displacement from the dipole kick by 5 %, the horizontal calibration obtained from this measurement is 13.2 px/mm, giving $\phi = 48°$. While this number seems much more reasonable than the one measured earlier, the reason for the discrepancy is unclear.

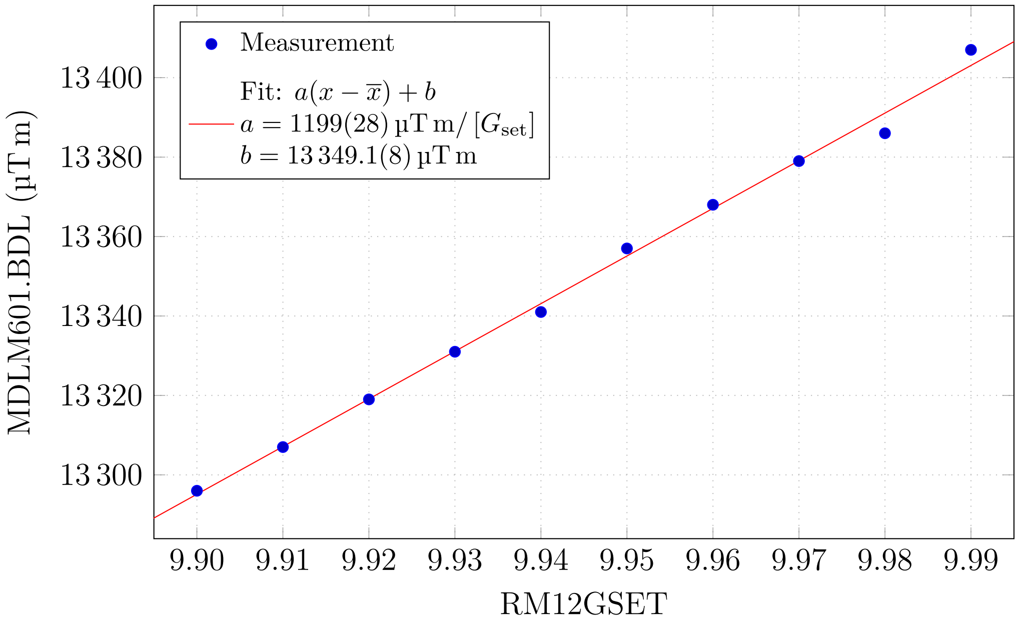

Figure 9: Dipole strength (assumed proportional to beam momentum) as a function of G_set around the reference point of 9.95, $p_0 = 8.0\,\mathrm{MeV}/c$. The relative momentum change per unit of G_set is $\xi = a/b = 9.0(2)\,\%/\,[G_\mathrm{set}]$. *(Note: From a purist's point of view, this unit is difficult to handle in that people often confuse the meaning of G_set and G_mes with the energy gain per unit length, while they are actually proportional to the field amplitude. As these quantities are not proportional for non-relativistic particles, I am hesitant to use the unit* $\mathrm{MV/m}$ *and will use* $[G_\mathrm{set}]$ *instead for lack of a less misleading unit.)* The numbers were obtained by adjusting the magnet to center the spot on the viewer every time (using the centroid value from the profile fit as a reference). The magnet was cycled at every point and iteratively readjusted if needed. Because this measurement takes a few minutes, some energy fluctuation is unavoidable, so reproducibility is not stellar.

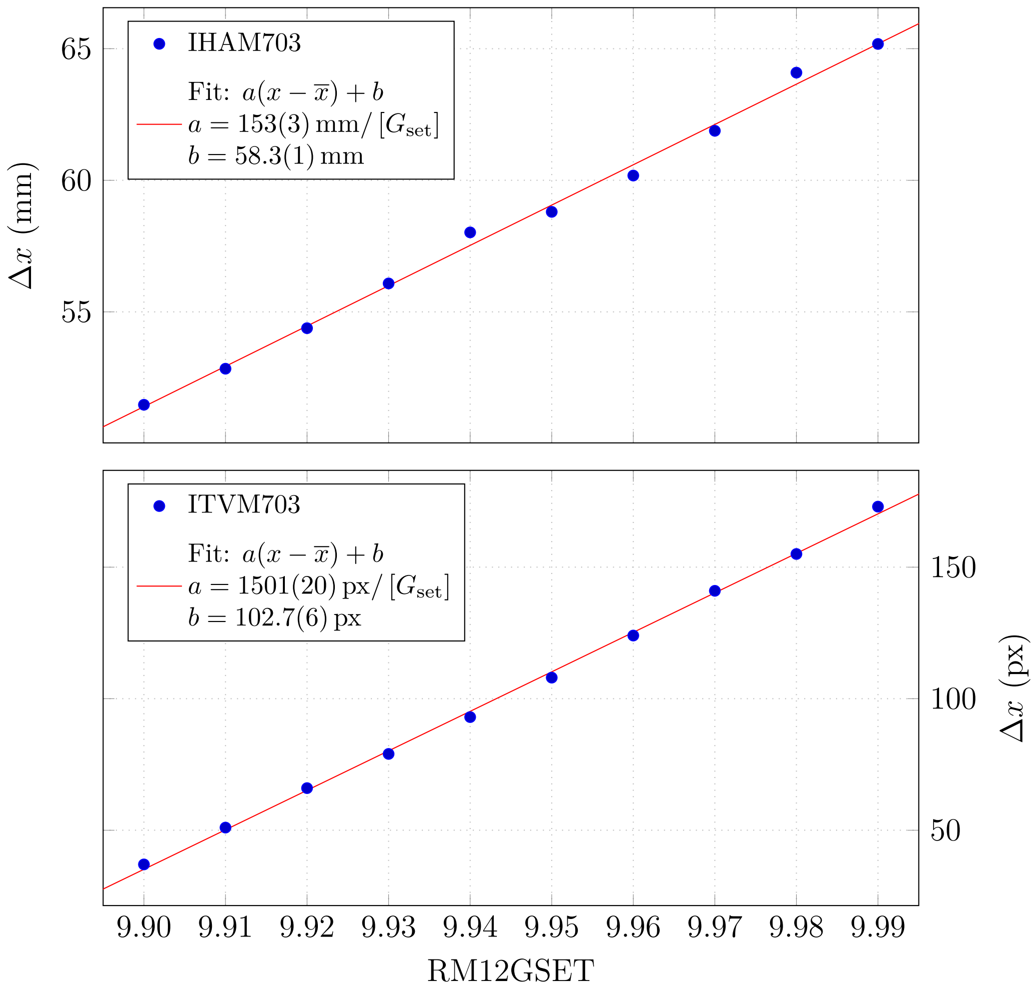

Figure 10: Dispersion at IHAM703 (top) and ITVM703 (bottom) in their respective units of length. In the case of the harp, the slope needs to be scaled down by a factor of $\sqrt{2}$ to account for the harp rotation. With ξ from Fig. 9, the dispersion is $\eta = a/\xi$. This gives a dispersion of 1.20(3) m for the harp and 16 708(447) px for the viewer.

2.3 Harp-based momentum spread measurement

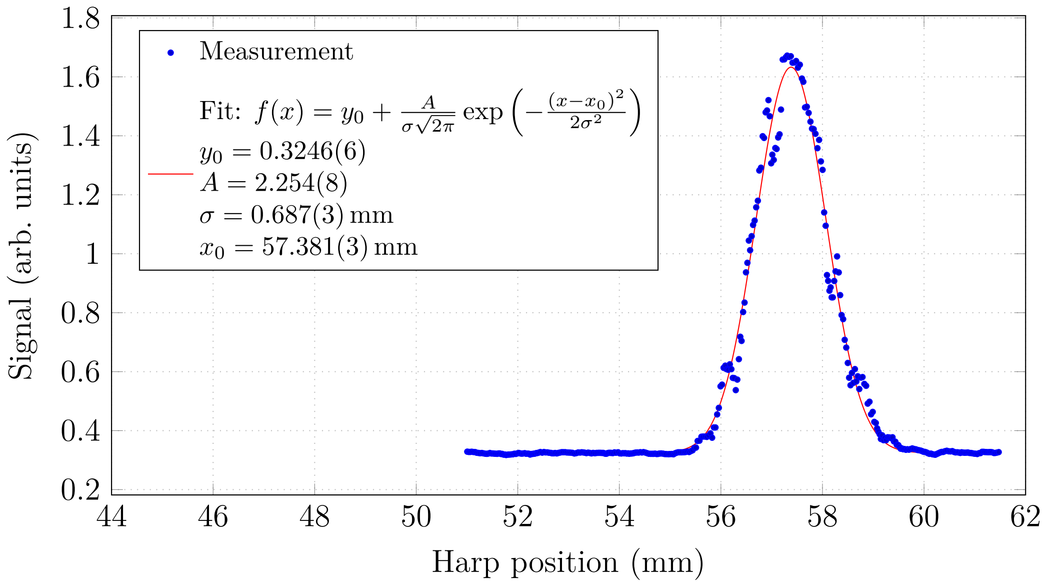

Figure 11: Beam profile measurement at IHAM703 with all cavities set up for minimum momentum spread and the quadrupoles set up empirically for minimum transverse beam size on the 703 viewer. $p_0 = 8.0\,\mathrm{MeV}/c$.

Figure 11 shows a profile measurement at IHAM703 under optimum conditions. The shape of the x peak is a convolution of betatron-optical and dispersive displacement of particles within the beam on the one hand and unwanted AC steering of the whole beam on the other hand, the spectrum of which is—a priori—unknown but assumed to contain frequencies well above the time scale of the harp, though the tune-mode signal being mains-synchronous may suppress the 60 Hz content. Though not necessarily constant over time, the AC steering part can effectively be treated as an increase in emittance for now.

Assuming both the momentum spread and the intrinsic transverse phase space are Gaussian, the width is

$$\sigma_x = \sqrt{\beta_x \epsilon_x + \left(\eta_x \frac{\delta p}{p_0}\right)^2}. \tag{6}$$

Because β_x cannot be measured directly, we shall ignore its contribution for the time being; with $\eta_x = 1.20\,\mathrm{m}$ as shown in Fig. 10, the measured σ_x (scaled down by a factor of $\sqrt{2}$ to account for the harp rotation) gives an upper bound of the momentum spread:

$$\frac{\delta p}{p_0} \leq 4 \times 10^{-4}. \tag{7}$$

Since it takes about 2 s for the harp to scan over the peak at 2 mm/s, this number includes contributions from subsonic and audio-band noise that might be present on the beam on top of the "true" momentum spread (in an intra-bunch sense), but it ignores any slower drift. The stability is explored on different time scales in Sections 2.4 and 2.5.

The primary reason why the intrinsic momentum spread cannot be made zero is the variation of the RF field strength around the nominal arrival phase in any cavity. The impact of this variation in the booster is mitigated by the buncher, which, somewhat paradoxically, is based on the same principle as the problem it is trying to solve, adding an energy spread proportional to the upstream bunch length (assuming the bunch is short compared to the RF period), except it does so in a non-relativistic context such that the imprinted velocity profile results in the bunch being focused longitudinally. Apart from being defined by the laser, the bunch length at the entrance of the buncher can be increased by space charge, making it current-dependent, and some initial energy spread from the gun (which is otherwise negligible in terms of its direct contribution to the total energy spread).

The energy spread added by the buncher inevitably stays with the bunch for the rest of its life; the upside of this compromise is that one can hope to make the bunch negligibly short at the position where it is made relativistic so that both the bunch length and the absolute momentum spread stay constant from that point onward. In practice, the acceleration to relativistic velocity happens over some non-zero length, which includes not only the 2-cell but also part of the 7-cell cavity. The notion of having a longitudinally frozen beam after the 2-cell is a misconception: with our current parameter set, β at the exit of the 2-cell is only 0.90 (660 keV kinetic energy per Yan's measurement). As a result, the optimal longitudinal focal length of the buncher is a (weak) function of the booster parameters.

The parameter space, having many dimensions, is cumbersome to explore experimentally. Apart from the necessity of resteering the beam line in response to big momentum changes, the MeV bunch length cannot be measured in the current setup; only the momentum spread can. Some simple model-based studies of the longitudinal phase space were described earlier this year to explain the measured energy calibration of the booster [5, 6, 7, 8].

Among the simplest experiments one can conceive to see if the results make sense is a variation of the buncher amplitude around its empirically determined optimum. The momentum spread is then determined in the same way as before; the result of this measurement is shown in Fig. 12 (individual harp data not shown). The simulated lines come from the tracking model mentioned earlier with the parameters of the booster cavities chosen to give on-crest acceleration to the momenta measured in the real machine (1.05 MeV/c at the 2-cell exit and 8.0 MeV/c at the 7-cell exit, respectively). The lack of agreement is indicative of an incorrect model (both in terms of simulation and interpretation of the measurement), including but not limited to the contribution of the intrinsic beam width to the width of the measured harp peak. It may be interesting to repeat this comparison without the 7-cell to reduce the number of uncertain parameters; however, at such a low momentum, the geometric emittance may be too high to make meaningful momentum spread measurements. One could also try to measure the momentum spread after the buncher directly to see if it agrees with the model based on the measured bunch length. It is known from a recent direct measurement that the RMS bunch length at the chopper (almost directly upstream of the buncher) is about 20 ps, though the shape is non-Gaussian [9]. The simulation would become more accurate if the true shape of the distribution were taken into account.

If one wanted to be careful, it would be a good idea to check the centering in the booster and the downstream magnets after changing the buncher amplitude to avoid systematic effects due to the buncher steering the beam; it is known that due to the relative arrangement of correctors and apertures, the buncher cannot be centered in and will always steer by a tiny amount.

To avoid repeating the surprise we had with the G_{set} calibration of the booster [7], Fig. 13 shows the buncher power as a function of amplitude setpoint. In the setpoint range of interest, there is no deviation from quadratic behavior, so the relative scaling of the amplitude should be well-behaved.

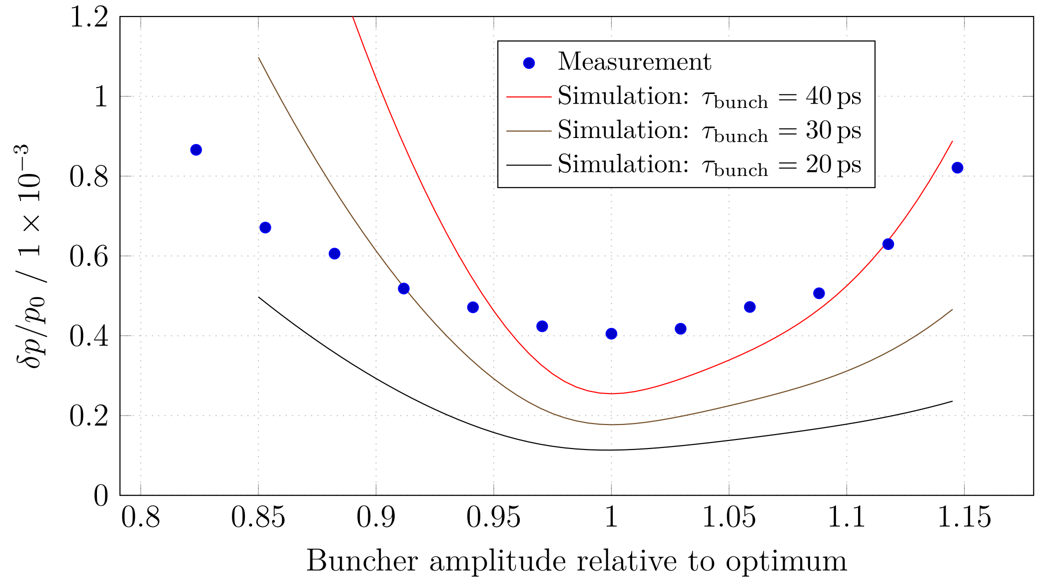

Figure 12: Buncher amplitude scan for comparison with model-based predictions. In the measurement, $G_{\mathrm{set}} = 34$ is considered optimal. The simulated bunch lengths τ_{bunch} refer to the RMS width of a Gaussian distribution. Even assuming that the measured values are overestimates because of Eq. 6, the agreement is unconvincing at best.

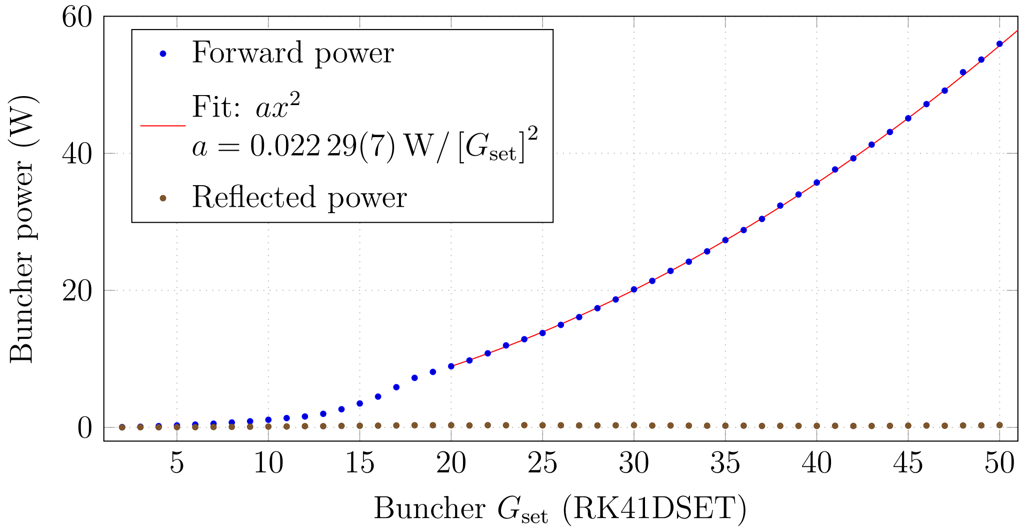

Figure 13: Buncher power as a function of nominal field. Above a G_{set} value of 20, the square law is close to perfect for the forward power, while no power is reflected, so the field can be assumed to be proportional to G_{set}.

2.4 Viewer-based momentum stability study

While working with the dispersive MeV beam lines 700 and 800, we have always seen the beam move inexplicably on various time scales between seconds and hours. As this problem does not appear as significantly in the non-dispersive parts of the machine, we feel compelled to look for the reason in the RF system. Results from an earlier test without beam [10] suggest a relative field amplitude stability of 2×10^{-4} and 5×10^{-4} (RMS) for the 2-cell and 7-cell cavity, respectively; these results are attributed to the noisy Test Lab area and unstable helium pressure. For serious beam studies at the UITF, this would be an issue in and of itself, but seeing as the data sample that the numbers come from only covers 200 ms, the total stability over a longer time window may be even worse; an attempt to assess this will be made in Section 2.5.3. Throughout the months, there has been a frequent need for setpoint adjustments to obtain the nominal beam position in the dispersive line as specified by the users [11]. The peak-to-peak magnitude of the changes we have seen over the months is as big as 2 %, which is not close to acceptable, though to what extent it is actually the momentum that changes has not been investigated enough.

We started by recording a one-hour-long movie of the 703 viewer being illuminated with a CW beam, thereby avoiding modulation or masking that may occur with a macropulsed beam in case the beam motion includes components that happen to be synchronous to the macropulse structure. This measurement was facilitated by an extra optical attenuator inserted into the laser path to decrease the beam current to about 10 pA for humane treatment of the viewer.

Figure 14 shows an example of how each camera frame is reduced to the statistical quantities of interest, i.e., horizontal beam position and width. The beam displacement is assumed to be proportional to the change in beam momentum.

Assuming the frame rate of the camera is sufficiently constant, we can divide the time-domain data of the central beam position into slices and Fourier-transform them to get a sense of the frequency content of the beam momentum. This spectrogram is shown in Fig. 16.

The frequencies have to be interpreted with caution as we have to assume some of them are a result of aliasing; at a sampling rate of only 30 Hz, we are certainly violating the sampling theorem. It is, however, very interesting to see that this crude method works at all, at least for frequencies not too far above the sampling rate; high frequencies will be averaged away due to the photoresponse of the viewer and the integration time of the camera; they will only contribute to the effective width of the observed beam spot. The degree to which this happens cannot be discerned with this method because there is no easy way to change the bandwidth.

What we learn from this exercise is that the beam motion we have always seen on the viewer predominantly consists of sharp spectral peaks, most notably one at 3.5 Hz, with barely any contribution from random noise. The 3.5 Hz peak may be an aliasing artifact of 26.5 Hz, which commonly shows up in the LLRF spectra of the booster, see below. Some of the spectral content does not appear until 10 min into the measurement; attempts to find a correlation with other measurements recorded at the same time, e.g., microphonics, were unsuccessful, so I cannot offer an explanation.

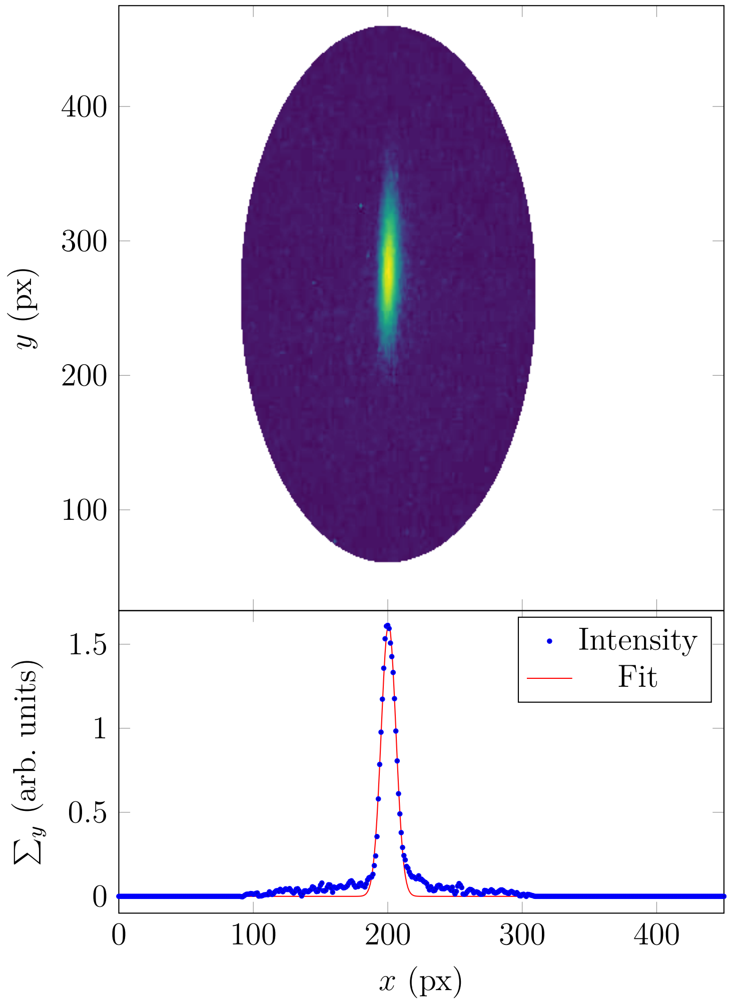

Figure 14: Example of a camera frame of the M703 viewer. The data has an elliptical exclusion region applied to it to remove background unrelated to the viewer. An arbitrarily defined elliptical ring within the viewer area serves as a reference for background correction. The vertical axis is then summed away and a Gaussian fit applied to obtain an estimate of the central value and width of the beam spot. The method may not be 100 % clean, but it gave the results I wanted without days of coding.

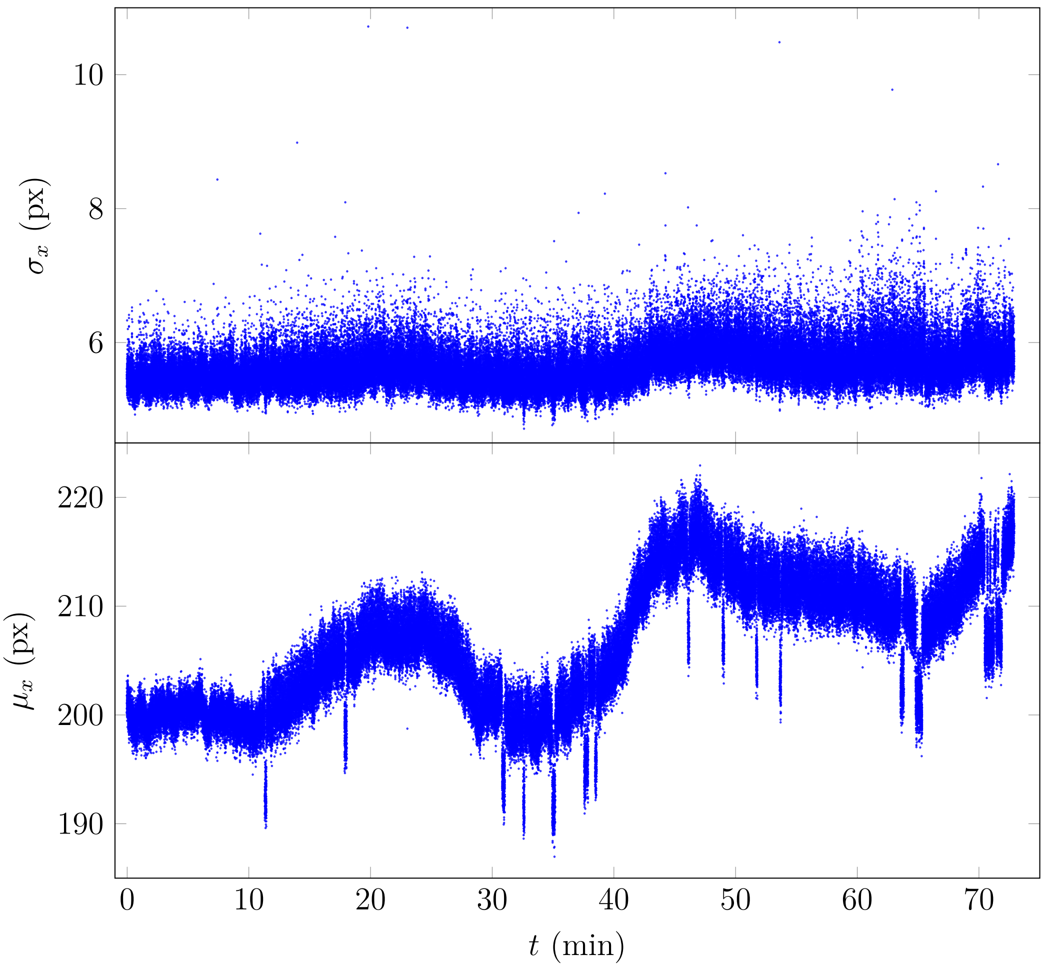

Figure 15: Evaluation of the one-hour-long movie of the M703 viewer (131 149 images in total at 30 Hz frame rate). The central position (bottom plot) shows three features: long-term drift, short-term excursions (about 10 s each), and rapid fluctuations that look like noise but are actually concentrated at sharply peaked frequencies, see Fig. 16. These data were recorded on 11/10/2021 starting at 8 p.m.

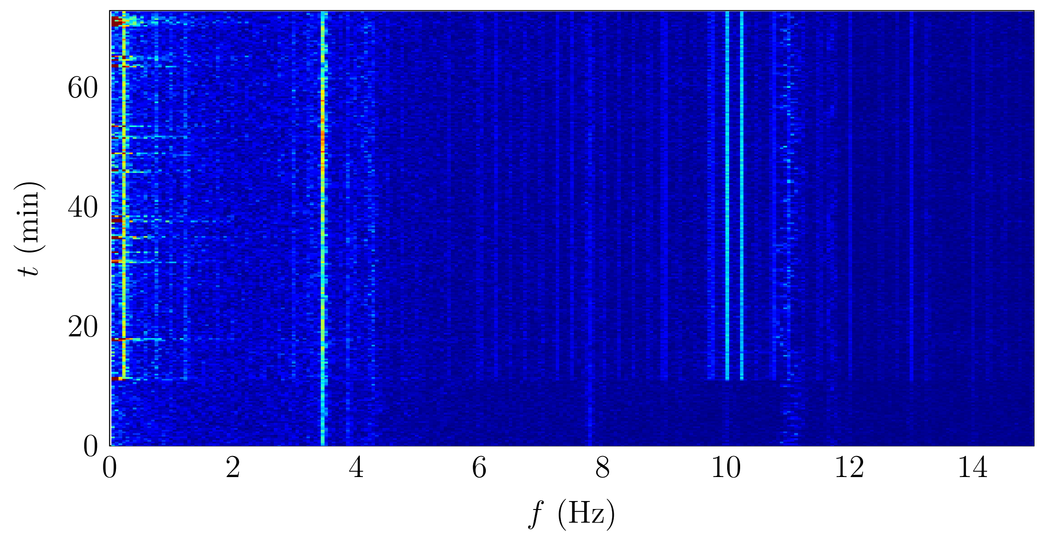

Figure 16: Spectrogram of the beam position on the M703 viewer based on the data shown in Fig. 15. Color scale in arbitrary units.

2.5 BPM-based study of momentum and orbit stability

2.5.1 Spectral features with CW beam

BPMs allow convenient access to the spectrum of the beam position over a much wider range of frequencies than what is available from viewers. However, seeing this spectrum in all its beauty requires a CW beam, which makes such a study problematic at the UITF because of the radiologically limited average beam current (< 150 nA). While the BPMs themselves are designed to work at such a low current, there are unresolved prohibitive issues with power RF from the booster cryomodule crosstalking into the receivers. The current setup therefore needs a current in excess of 1 µA in order for the SNR to be acceptable in CW mode. We obtained temporary beam authorization for 10 µA to perform a CW measurement with continuous RadCon monitoring, but this option is not routinely available, preventing us from making a day-long measurement.

Hoping to see any frequency correlations or other interesting features, we recorded the demodulated spectra from both booster cavities and all relevant

BPMs while running 8 µA of CW beam into the M703 dump for an hour. The sampling time is a compromise between aliasing-free bandwidth and low-end frequency resolution; it was chosen as 40 µs for the BPMs and 100 µs for the SRF cavities. First, to get a sense of the overall spectral content, the spectrum of all BPMs averaged over the whole run is shown in Fig. 17. The spectra in the non-dispersive beam line are dominated by mains hum, which seems hard to avoid entirely; some of it will be elucidated later. There are also other peaks, e.g., at 27 Hz and 103 Hz, which may be mechanical vibrations in the machine or something else. Unsurprisingly, the dispersive beam line hears a cacophony of other notes in addition, which are presumably added by the booster.

Figure 18 is an attempt at correlating the time-averaged spectrum of the dispersive BPM with the detuning and field probe spectra of the booster cavities. Some correlations with microphonics are rather obvious (e.g., 11 Hz as seen before in the viewer measurement, see Fig. 16), others less so. Most prominently, the beam spectrum contains loud notes at 120 Hz and 692 Hz, which appear in the field of the 7-cell cavity but are not caused by detuning, at least not in any obvious way. To explain the 120 Hz peak, considering the lack of similarly high 60 Hz content, I would not look for a ground loop but rather for supply ripple in the RF system, e.g., the HVPS of the HPA; of course, it might be something else entirely, and the peak at 692 Hz is not clear at all. It is interesting that the latter and its harmonics, while not dominating the cavity spectrum, seem to account for the dominant part of the beam momentum spectrum. The 27 Hz peak is visible in all non-dispersive BPMs as well as both SRF detuning spectra, so it may be mechanical motion including but not limited to the booster. How much of it is in the beam momentum vs. the orbit displacement is unclear.

Whoever enjoys looking at spectra in a time-windowed way will find such data in Figs. 19 and 20 (same data, different frequency scales). These representations show some extra features that might otherwise not be obvious:

1. The transverse position of the beam sees intermittent “honking” at the mains frequency with a respective duration of multiple minutes and a period of about 20 minutes, which is highly suspicious. Section 2.5.3 offers a possible explanation.

2. Every 5 min, a burp goes through the RF and lights up the whole spectrum

for a single acquisition frame, so this may be some transient mechanical excitation. While this problem in itself does not lead to FSD trips—at least not during these particular measurements, but things might compound—it is very visible in the beam. Whatever it is that periodically bonks the machine, I do not like it.

3. While some time-dependent components are visible in the field spectrum, their relative amplitude is too small to contribute to the beam significantly. In particular, noise is not a relevant problem.

Lastly, the second-time-scale beam position distribution at the M702 BPM is shown in Fig. 21 (bottom) along with the evolution of its RMS width over time (top). Not accounting for the outliers, the numbers may serve as an estimate of how the audio-band perturbations contribute to the effective momentum spread in applications of such time scale (e.g., harps, though those will most likely not see mains harmonics when operated in tune mode). We see that the contribution is about 2×10^{-4} (at $\eta_x = 1\,\mathrm{m}$), which is not overwhelming but quite significant. As the bottom figure shows, the components comprising the fluctuation are random enough to make the distribution somewhat Gaussian-like, so while I am not fond of the inflationary use of RMS values of non-normal-distributed quantities, one may assume that the convolution between the true momentum spread and the time-dependent changes is close to a quadratic addition of the RMS values.

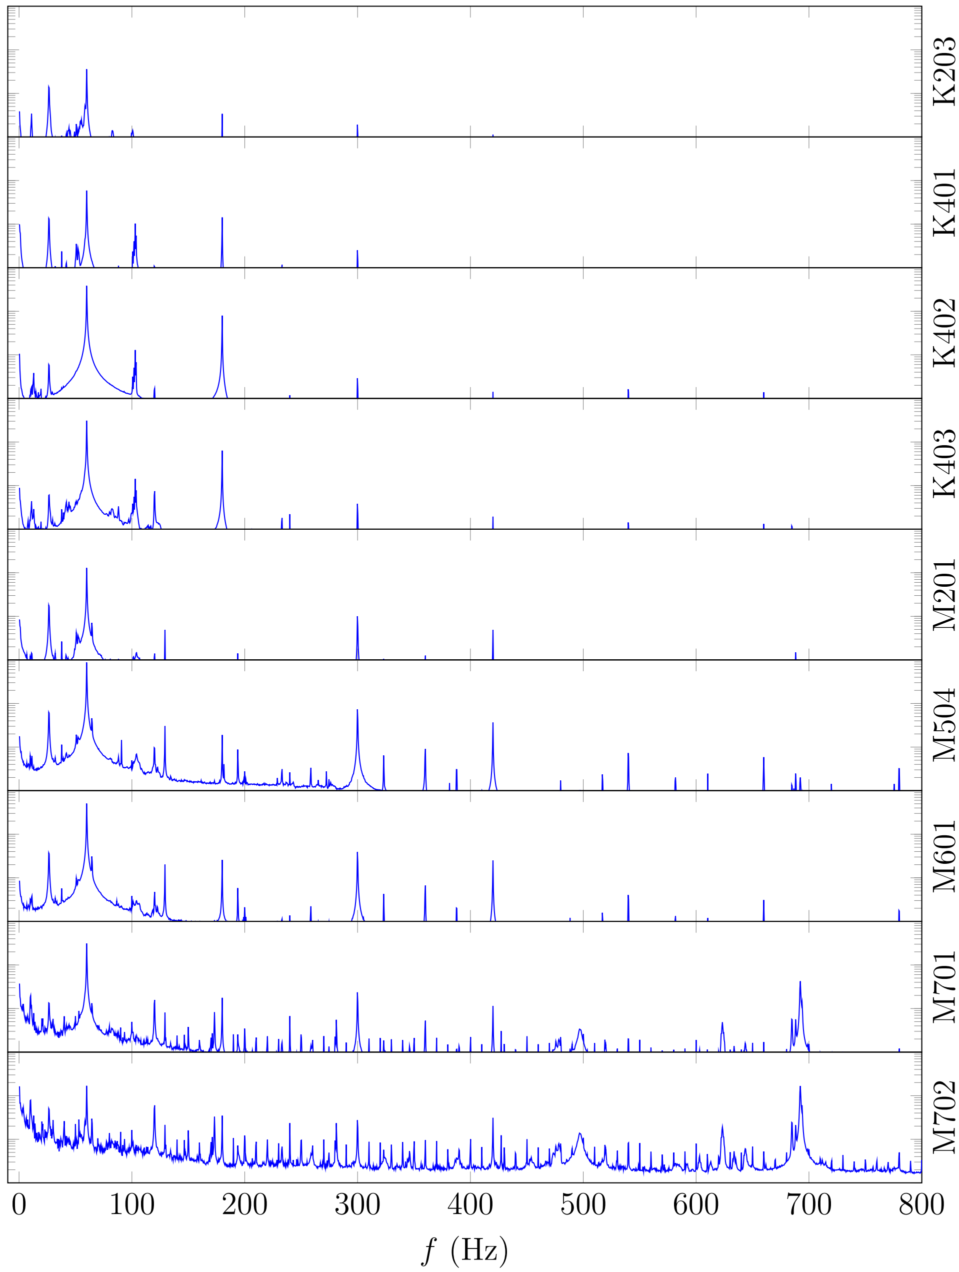

Figure 17: Average spectra of all BPMs going to the M703 dump, all in the same (arbitrary, logarithmic) units. The non-dispersive spectra mostly contain mains harmonics, 27 Hz, and 103 Hz. If the latter are mechanical, it may be interesting to try to damp them away.

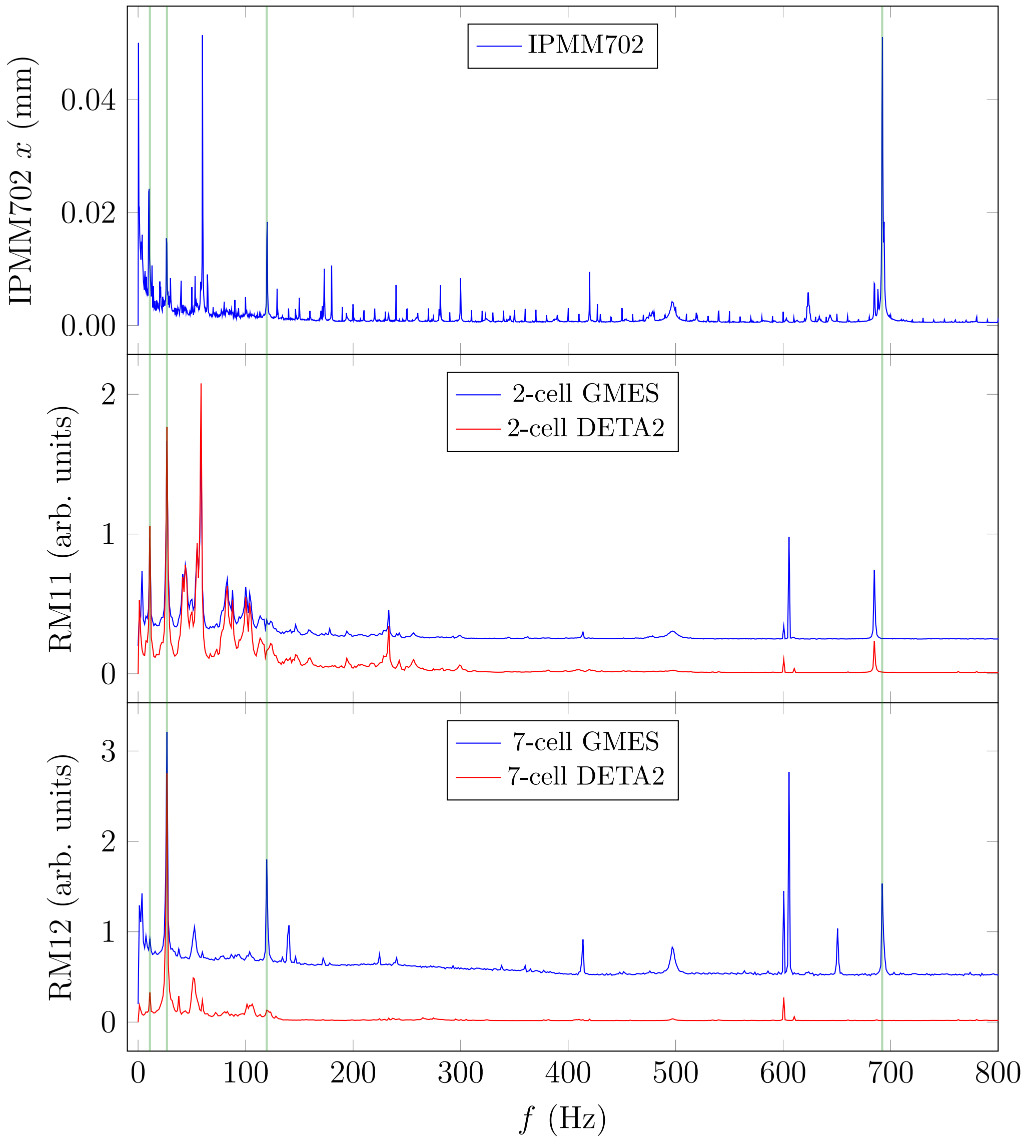

Figure 18: Spectrograms of dispersive BPM, field probes, and detune angles averaged over time; lines are shifted vertically for visibility. The green lines indicating correlation between a dominant peak in the momentum spectrum and one in the cavity spectra are located at 11, 27, 120, and 692 Hz, respectively.

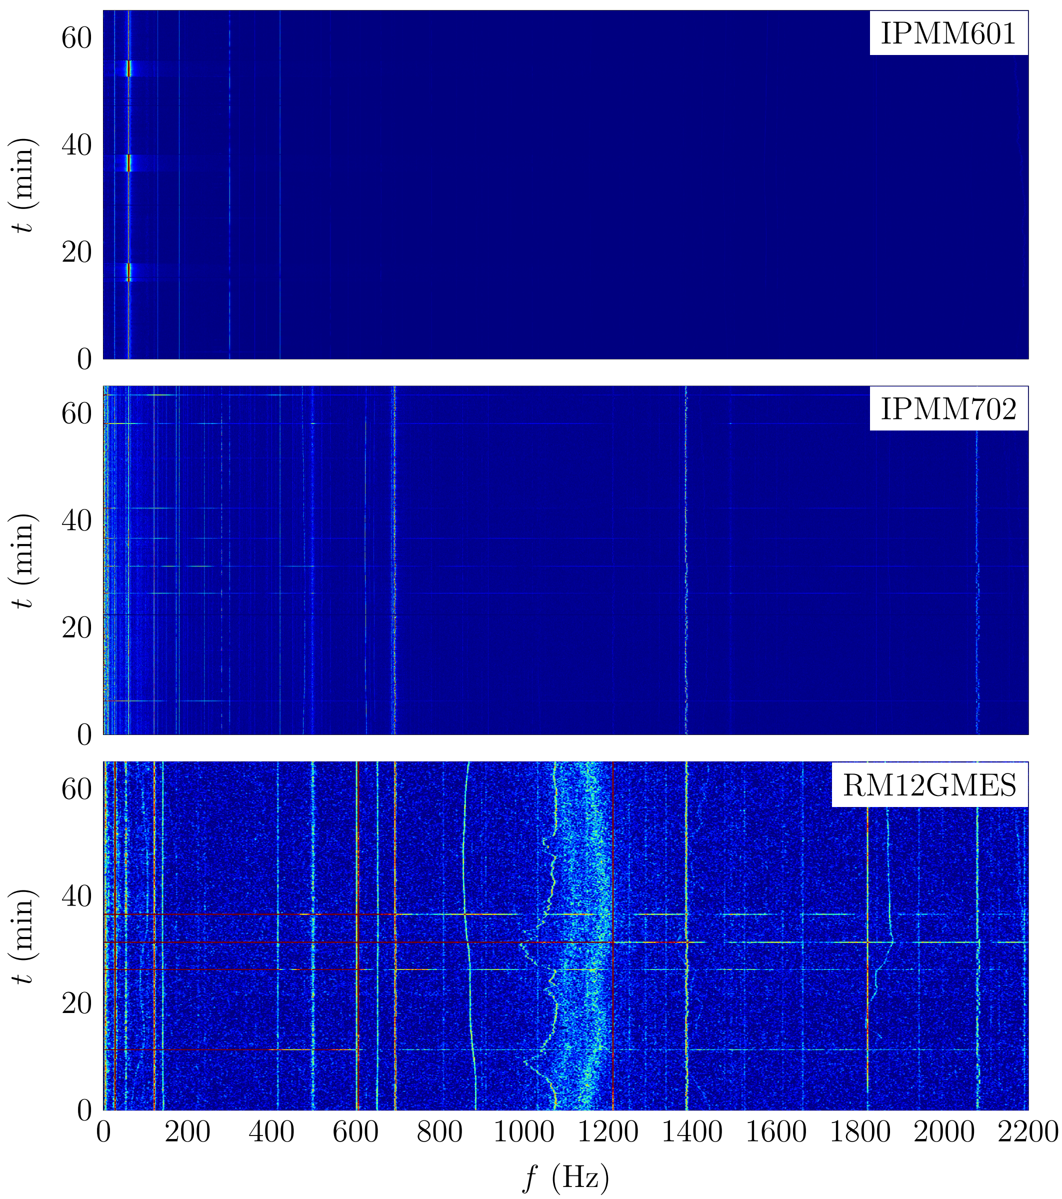

Figure 19: Spectrogram of non-dispersive and dispersive beam position, respectively, and field probe signal from the 7-cell cavity. Linear color scale in arbitrary units.

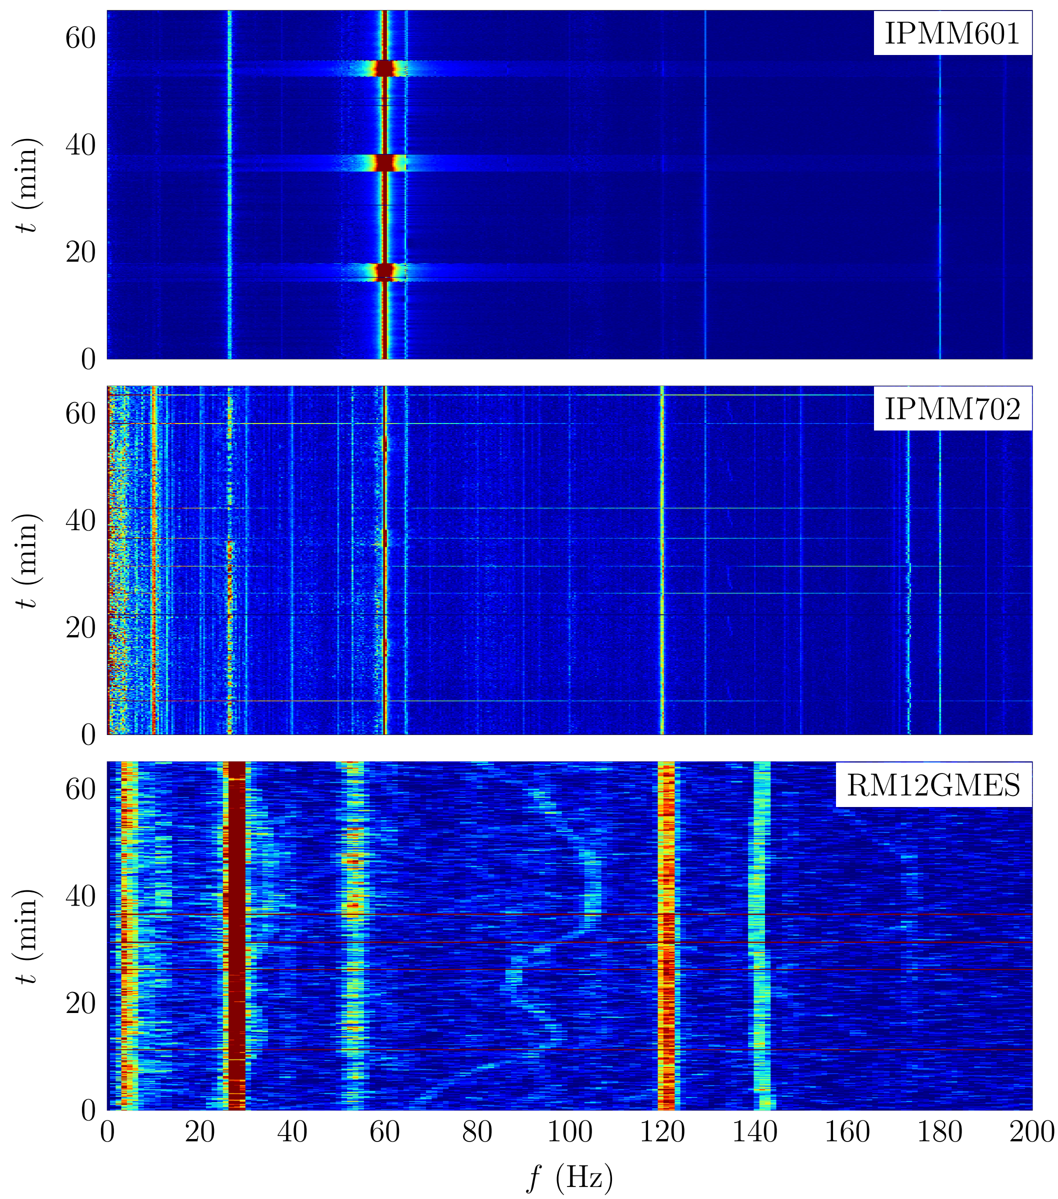

Figure 20: Spectrogram of non-dispersive and dispersive beam position, respectively, and field probe signal from the 7-cell cavity. Linear color scale in arbitrary units. Narrower frequency scale than Fig. 19.

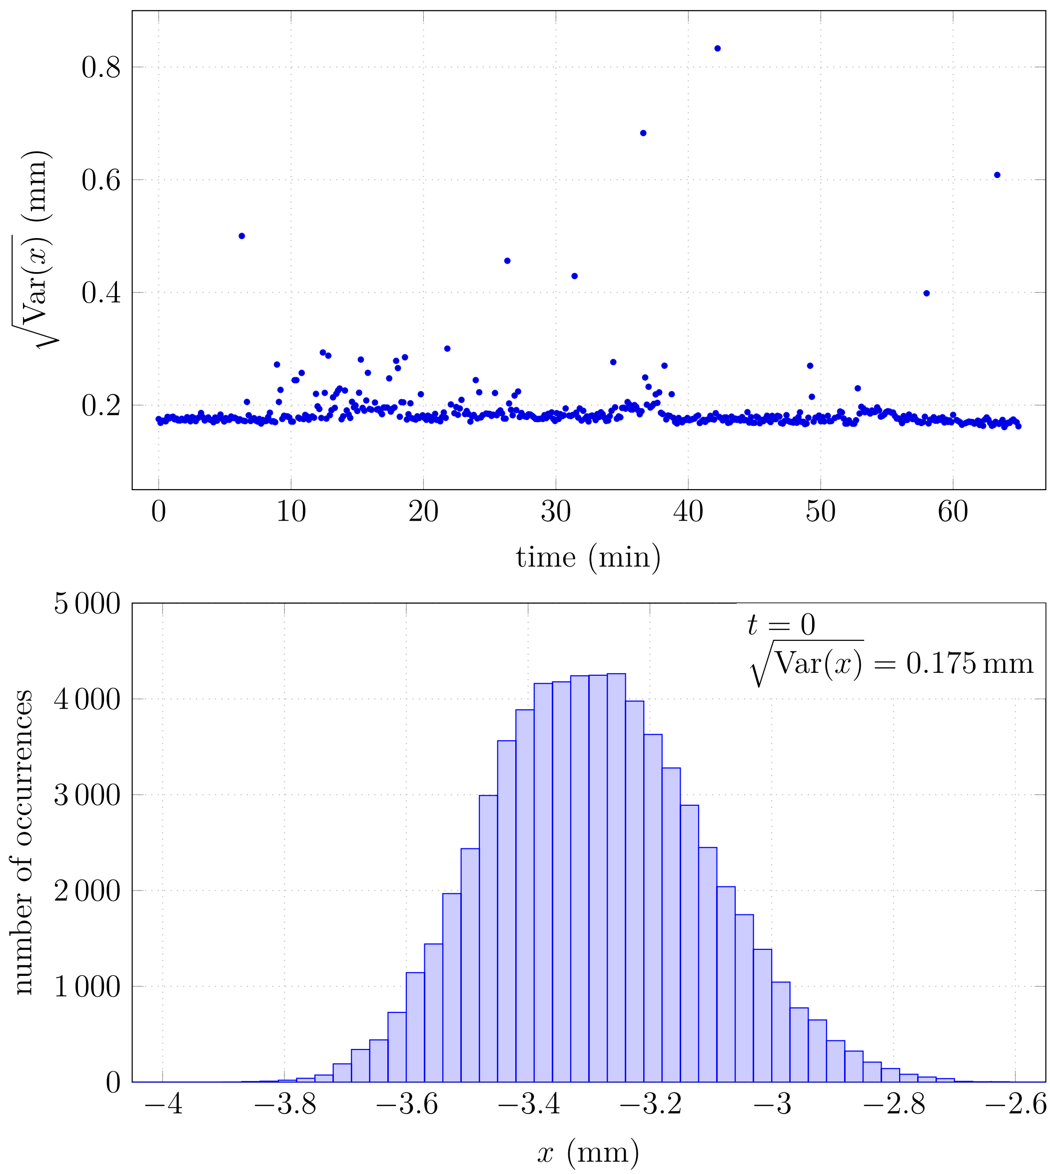

Figure 21: Top: Second-time-scale RMS beam position fluctuations at IPMM702. Each data point represents a duration of 2.6 s. This plot accounts for the rotated BPM axes. Except for the outliers, the usual value seems to be 0.18 mm, i.e., $\delta p/p \approx 2 \times 10^{-4}$. Bottom: Histogram of the distribution in one example data point.

2.5.2 Booster field modulation

We wanted to do a modulation experiment to make sure the spectra of the BPMs are real and well-understood. The original idea was to modulate the beam current to be able to use a lock-in technique, but this would have involved an unnecessary level of complexity and could not be done within a week. Instead, we added amplitude modulation at 42 Hz to the field setpoint of the 7-cell cavity. The resulting measurements of the field amplitude and beam position are shown in Fig. 22. While the frequency resolution of the cavity signal is not optimal because of the choice of sampling rate and buffer size per frame, the comparison between the two spectra offers enough insight into how well we understand the system.

The 42 Hz peaks in the G_{mes} and BPM plots have an amplitude of $a = 0.030\,98(6)\ [G_{\mathrm{mes}}]$ and $b = 2.03(6)\,\mathrm{mm}$, respectively. The latter should actually be $b' = b\sqrt{2} = 2.87(8)\,\mathrm{mm}$ to account for the CW data having been recorded with a 45° rotation that was internal to the receiver electronics, so for a purely horizontal displacement, the BPMs would show the same value in both axes. With ξ from Fig. 9, the dispersion at the BPM is

$$\eta_x = \frac{b'}{a\xi} = 1.03(4)\,\mathrm{m}, \tag{8}$$

suspiciously close to expectation.

The harmonic content of the two spectra is very different: second-harmonic and third-harmonic distortion of the 42 Hz peak shows up at relative levels of $-45\,\mathrm{dB}$ and $-49\,\mathrm{dB}$, respectively, in the cavity field, while they are at $-27\,\mathrm{dB}$ and $-44\,\mathrm{dB}$ in the BPM. Maybe the BPM adds some distortion of its own due to the massive signal. This observation is curious but not relevant enough to warrant any further study at this time.

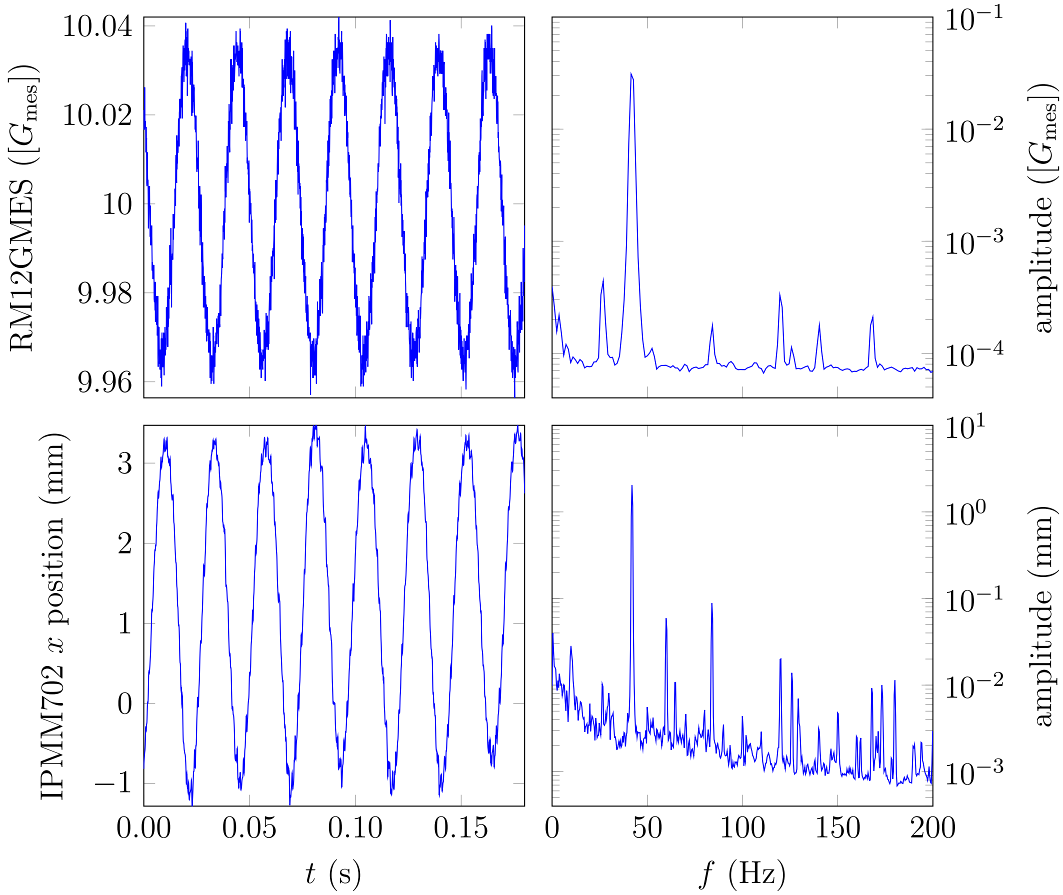

Figure 22: Time-domain and frequency-domain representations of the modulated cavity field amplitude (top) and the resulting beam displacement (bottom). The FFT is computed from the time-domain data of each frame with a Hann window and with DC suppressed, normalized to the sample length, and averaged over 10 min.

2.5.3 Long-term stability with tune beam

Because we have always had to deal with issues of orbit and momentum reproducibility—not only after cold start-up, but also in the middle of running—it is worthwhile to measure the orbit as a function of time on a thermal/human time scale. This would ideally be done for multiple days without interruption to see all effects related to the mechanical and thermal environment of the Test Lab. However, we had not understood how to make the BPMs work properly in tune mode until the last week of operation, so allocating enough time and staffing was difficult on short notice. We will have to be content with a 4-hour-long measurement for the time being. This test will be very simple to repeat at the start of the next run.

Figure 23 shows the average BPM readings while delivering tune beam to the M703 spectrometer dump throughout this 4-hour period. A PID feedback loop was used to keep the average beam current constant at about 80 nA. Neglecting any upstream orbit change, the horizontal displacement at M702 is proportional to momentum change ($\eta_x \approx 1\,\mathrm{m}$, so 1 mm of displacement means $\delta p/p \approx 1 \times 10^{-3}$). The short-term momentum changes have already been discussed; what can be seen on top of those is a 1×10^{-3}-level fluctuation on a time scale of hours. One has to assume the drift is even worse when observed over a longer time scale.

In the non-dispersive part of the machine, we see some slow drifts of the beam position, for which no explanation exists at this time, but more prominently, there is something evil going on: almost periodic jumps that originate somewhere between K401 and K402. One might be inclined to guess that these jumps have something to do with the apertures, e.g., part of the beam being scraped away in response to some periodic change, causing the effective centroid to move. While such an effect may exist and the current signal intercepted on A4 indeed shows a dependency (not shown here), we should find the root cause.

Let us look at the buncher, which is the only interesting element in this region apart from the apertures. The period of the jumps is on a thermal time scale; Figure 24 indeed shows a highly suspicious correlation with the thermal detuning of the buncher, but the jumps themselves are too quick to be directly caused by something thermal. My first idea was that an amplifier in the buncher loop might be clipping at a certain power level, but the forward power looks benign. However, the buncher is directly heated by filaments controlled by what seems to be a two-point regulator supplying switched mains voltage. I am going to venture a wild guess, which is that these heaters make a magnetic field that steers the beam. I would not be at all surprised if the intermittent 60 Hz honking in the CW BPM spectrograms (see Fig. 20) also had something to do with this. If we are leaving this buncher in the machine and my guess about its heater is correct, we will need to find a way to heat it without an electric load directly attached to it. In any case, there is something to be learned here.

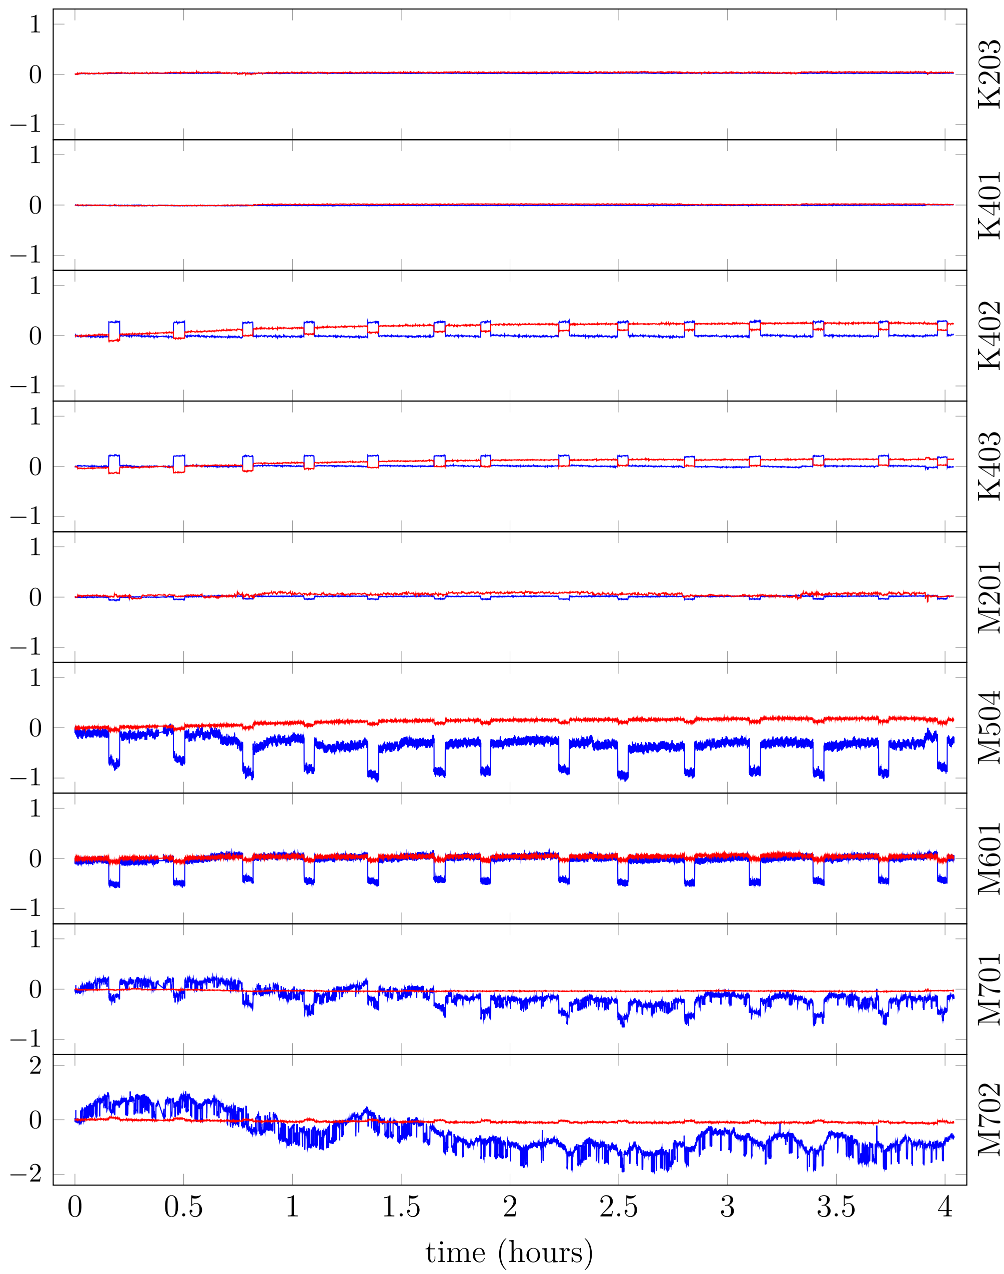

Figure 23: Average relative beam position (in mm) over time. Blue is x, red is y. $t = 0$ refers to 12/09/2021, 23:42. With $\eta_x \approx 1\,\mathrm{m}$ at IPMM702, the 4-hour drift of the momentum (neglecting AC components) is about 2×10^{-3} peak to peak.

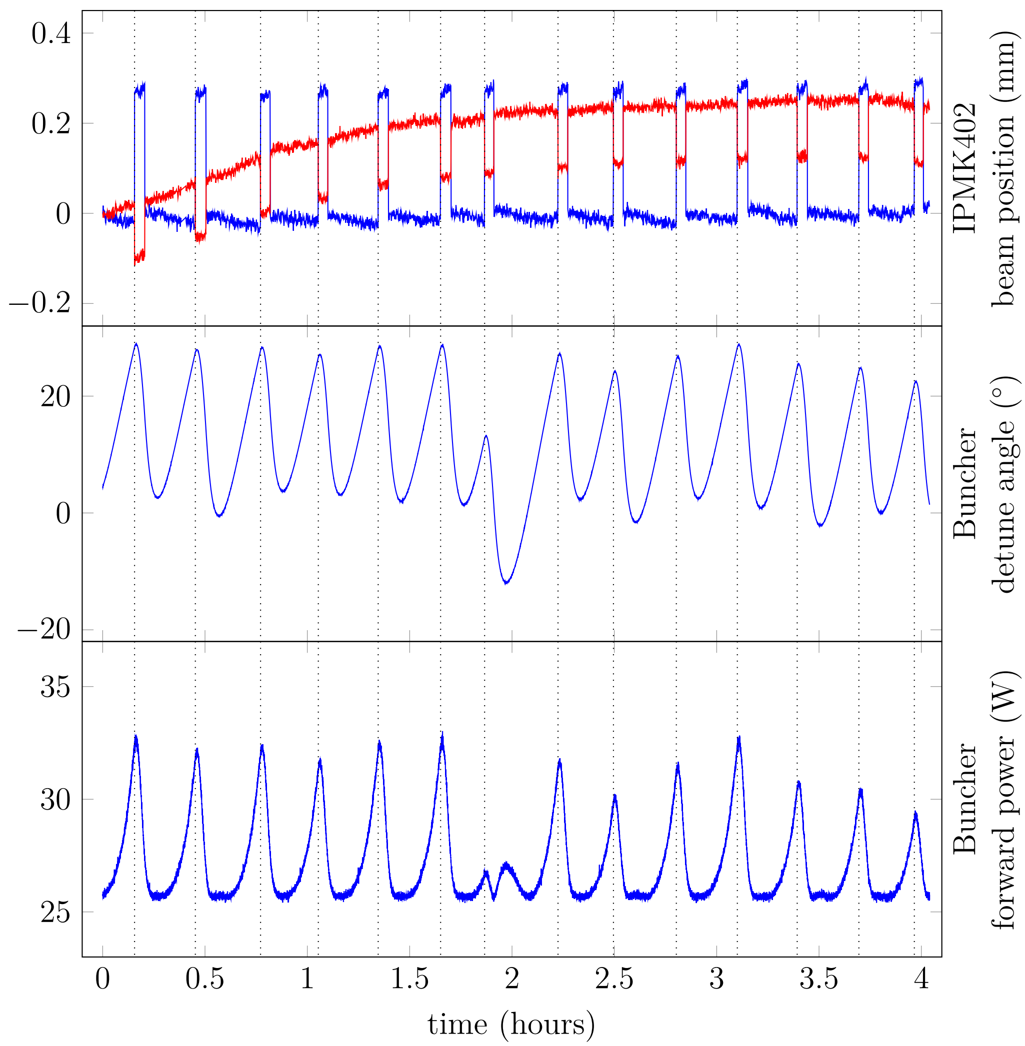

Figure 24: Average relative beam position over time compared to the detune angle of the buncher (i.e., its temperature). Note how the inflection points of the temperature coincide with the jumps (dotted lines), yet no amplifier clipping or other discontinuous effect is visible.

3 Transverse measurements in the straight-ahead line (M600)

The current UITF beam line design features a set of four consecutive QJ quadrupoles intended to transform the (a priori undetermined) beam ellipse at the exit of the booster cryomodule into whatever may be needed at the downstream end. Figure 1 shows the relevant components. In theory, the phase space at the exit of the booster can be simulated without any beam-based input, given accurate field maps of all components. However, because one has to be prepared for the possibility of this prediction being inaccurate, the idea behind this design is to measure the beam ellipse by way of quad scans so it can be used as a well-known input parameter for an optical model of the MeV line without any uncertainties of upstream components. Such a measurement, possibly at a variety of momenta, will also help compare the performance of the booster with predictions, which was initially the primary incentive for the whole study. Unfortunately, a number of real-life complications are currently preventing this methodical approach from working properly, so we will need to understand them first. While the beam spot at the downstream end has been set up empirically with some amount of success in the past, this situation is less than optimal, especially if the facility is going to be used for precision studies requiring accurate and reproducible beam parameters.

Figure 25 shows the results of different quad scans performed to study their robustness. The beam at the entrance of this beam line is nominally the same in every case, i.e., the true Twiss parameters at any upstream point are constant, and if the method works, the results should reflect that. In principle, if the Twiss parameters at any one position in the lattice are known, those at any other position can be calculated, provided all transfer matrices in between are known as well. However, the data suggest at least one of these premises is not met in practice: it is obvious that there are systematic dependencies we must understand and remove if we want the numbers to mean anything.

Because my knowledge of optics is still perfunctory, I will not try to give an authoritative explanation just yet. It would be nice to predict the Twiss parameters in some ab-initio way to get a sense of which scan, if any, might be

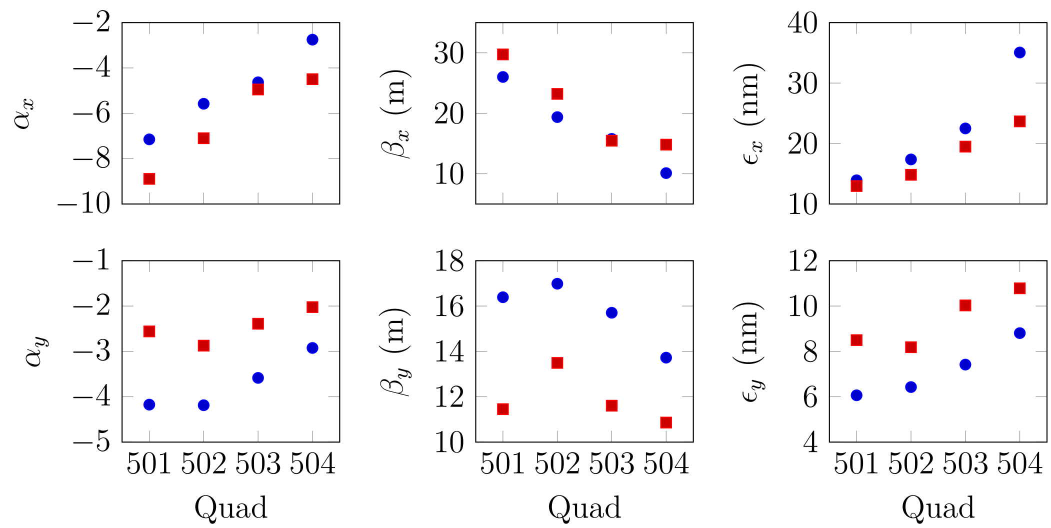

Figure 25: Twiss parameters obtained from different quad scans after backpropagation to the upstream end of MQJM501. For each scan, one quad was used at a time, the others being set to zero. The ● points were measured with the nominal orbit, i.e., centered in all quads at the same time, whereas the ■ points involved the smallest possible corrector strengths, only centering in the quad being used. $p_0 = 8\,\mathrm{MeV}/c$.

right, but this is prevented by a hen-and-egg-type problem in that all predictions based on a clearly incomplete model will be flawed as well.

The discrepancies we are seeing can be caused by a combination of multiple factors, some related to the measurement itself and the computations based on it, others related to beam manipulation. If the Twiss parameters change along the lattice in a way unaccounted for by the model, this invalidates any backpropagation as well as each individual measurement involving the respective part of the lattice; it also means that the beam quality (i.e., emittance) obtained at the exit may be irrevocably deteriorated, especially in the case of parasitic transfer matrices other than quads and drifts.

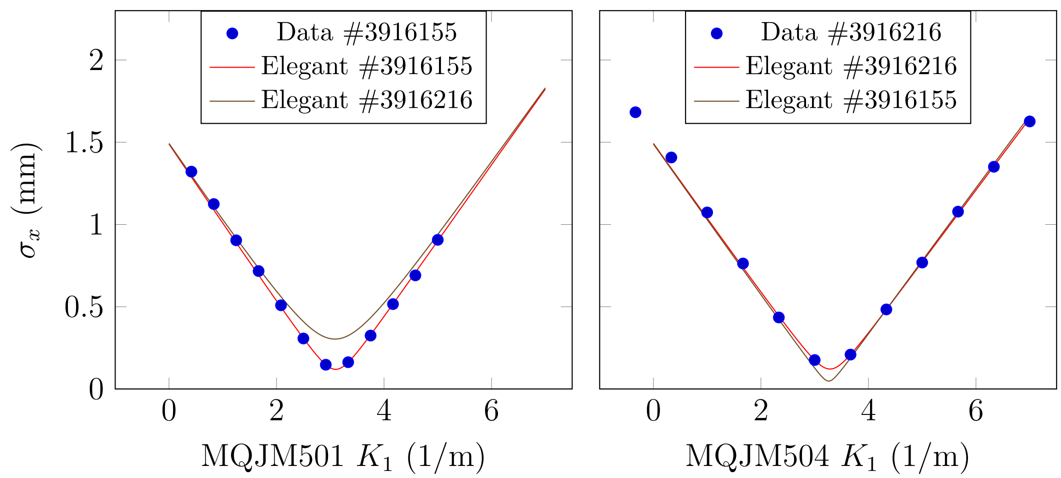

Figure 26: Predictions of the quad scan results for MQJM501 and MQJM504 based on the Twiss parameters backpropagated from the other measurement, respectively. The numbers in the legend refer to the logbook entries that have the data files in them.

The way in which the results are inconsistent is peculiar and not immediately obvious from Fig. 25. Figure 26 shows what would happen to the quad scan results of MQJM501 and MQJM504 if we pretended that the initial Twiss parameters were those obtained from the other data set, respectively (simulated with Elegant based on the same lattice as was used for backpropagation by qsUtility). It is clear that with this configuration of quad scans, the correlation between the Twiss parameters extracted from the fits makes the result ambiguous if one is not careful to explicitly measure the trough, which is where the curves differ, while the asymptotic behavior is the same. This ambiguity can unphysically trade ϵ for (α, β), which I had not been able to appreciate enough until Yves and Michael Tiefenback patiently explained it to me. Herein lies one potential source of error (which will be avoided next time), but it may not be the only one.

As can be seen in Fig. 1, there are more elements in the beam line than there would be in the idealized lens-drift-harp picture. In particular, there are two dipoles, which are nominally off, and a large number of correctors. The latter are needed to center the beam in all lenses at the same time, but, with the exception of the last one right before the harp, all correctors between MQJM501 and IHAM601 can be set to zero in y (small values in x) and still transmit the beam to the harp if only one quad is being used and one is willing to accept a minor amount of extra steering upstream. The impact this change has on the measured Twiss parameters can be seen in Fig. 25. Correctors can have multipole moments due to their geometry and additional alignment issues; also, even if they are ideal, they will always introduce some dispersion. The same is true for any residual field that may be in the dipoles. In a purely transverse optics model, uncontrolled dispersion will increase the transverse emittance (through the angular spread), which can only be predicted if the momentum spread is known; this is the main reason why the studies in Section 2 were a prerequisite.

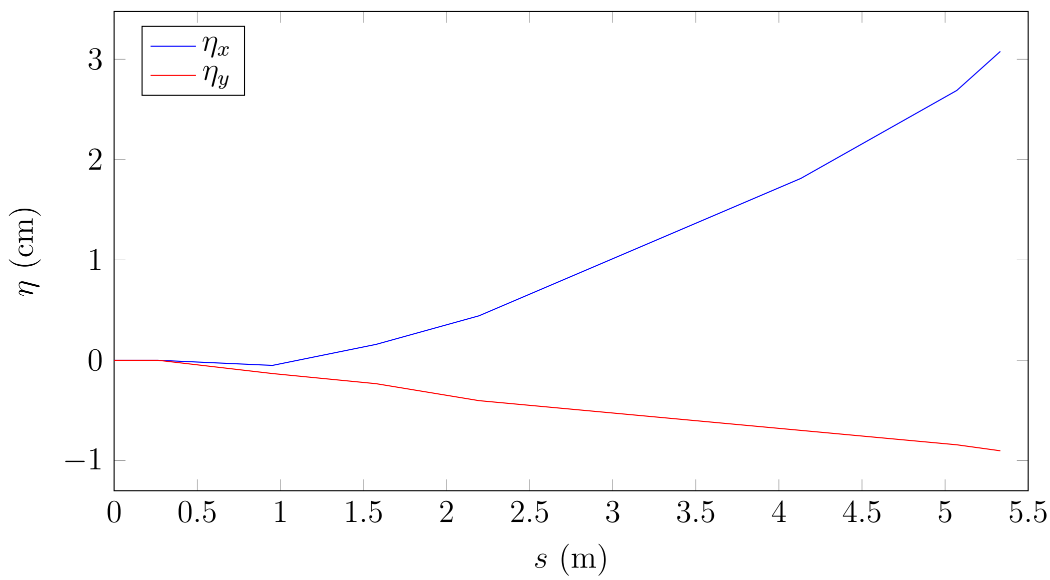

Figure 27: Elegant prediction of the dispersion from MQJM501 to IHAM601 based on the empirical corrector values for nominal orbit. Upstream dispersion contributions are not included as they should not affect the agreement between the quad scans. External fields are not accounted for.

Figure 27 shows the dispersion functions in the region of interest assuming the only steering fields are those of the correctors and all quads are turned off. With example values of $\beta_x = 20\,\mathrm{m}, \alpha_x = -5, \epsilon_x = 10\,\mathrm{nm}$, the horizontal emittance grows to 11.7 nm over the length of this section at $\delta p/p = 5 \times 10^{-4}$. As this growth is basically a convolution, it scales by quadratic addition rather than linearly, so the contribution would be disproportionately larger at higher $\delta p/p$. As it stands, the dispersive emittance growth from correctors alone is not enough to explain the issue.

What is hard to determine is the influence of the earth's field on the dispersion. The notion that it has no appreciable effect on an 8 MeV/c beam is a misconception: at this momentum, $p_0/e = 37.5\,\mu\mathrm{rad}/(\mu\mathrm{T\,m})$, i.e., 1.65 mrad of horizontal steering per meter of distance traveled at 44 µT vertical field and about half as much in the other plane (values from [12]). So far, I have been unable to reliably measure the field in the UITF enclosure and do not know its value or whether it is even there. Regardless, the conceptual problem with all this is that any dispersion from the earth's field will not necessarily add to that from the correctors; it is more likely to subtract from it because the corrector fields are meant to compensate for the orbit change due to external fields, not add to it. All these statements notwithstanding, I have tried to artificially increase the momentum spread to see how it would impact the emittance measurement, and it did so noticeably (not shown here because it was not done systematically).

Other assorted and poorly documented error sources include:

1. The (hard-edge) quadrupole model used in Elegant may not be accurate enough. Since a field map was not readily available, I did a 3D simulation with CST but have not made an effort to include it in the Elegant model so far; for lack of explicit field-map support, this would need to be done via slice-wise multipole decomposition.

2. The quadrupoles may have different field calibration factors and hysteresis offsets. However, Elegant tests have shown that any remotely reasonable combination of such errors will not reproduce the measurement.

3. Any reasonable and even unreasonable guesses of hypothetical corrector multipoles have no effect. I have tried to do a rough CST simulation of the Haimson coils as measured with a caliper, but to do it right, we would need to cut one in half to measure the true geometry of the current path. It currently does not seem worth the sacrifice and effort.

4. A miscalibrated harp axis has no effect, except that it changes the total emittance.

5. Measuring the distances between the elements wrong even by significant amounts will not reproduce the measurement.

6. We have observed that when degaussing a dipole from the cycled, nominally zero state, both its steering and, curiously enough, its focusing will change [13]. Then, cycling again will not go back to the original state. This strange behavior has not been investigated systematically but may contribute to both dispersion and focusing. Can the two dipoles be too close together so as to influence one another's remanence during cycling?

7. The impact of space charge has not been calculated. However, we have tried repeating the quad scans at 10 % of the peak beam current without seeing a big difference in the disagreements. The Twiss parameters were different, but the changes most likely take place in the keV beam line.

8. The beam does not usually stay centered in the quadrupoles during a quad scan no matter how carefully the centering is performed initially. This may be due to upstream orbit jumps and drift (see Section 2.5.3) but has not been thoroughly investigated so far.

9. Hall probe measurements in close vicinity to the beam pipe show that there are parasitic magnetic fields from various sources, most notably the body of the gate valve VBVM601 and multiple parts that attach the pneumatic assembly of viewer ITVM504 to its flange, which seems to include unnecessary magnetic parts (a quick test with a permanent magnet indicates regular steel, not stainless); Phil says the updated viewer assembly, which cannot be swapped in without venting, has no magnetic parts close to the

flange. Some stainless-steel nuts and washers are also mildly magnetic, which is to be expected and probably inconsequential. Whether or not any of these fields reach into the beam pipe defies direct measurement.

Once we have beam time again, the first thing that needs to happen is a BPM-based measurement of the dispersion throughout the MeV section (i.e., orbit change as a function of central momentum after the booster), allowing us to find out if there is anything grossly wrong with the steering. With all correctors and quads off (as far as possible), one can see the dispersion from any external fields (including earth) directly. One can also try to measure the orbit change in response to an upstream kick to find any parasitic focusing/defocusing fields. The same method can be used to compare the response of the quads with model-based predictions (it is basically a direct measurement of the β function). Such measurements are especially valuable seeing as we have more BPMs than we do other diagnostics.

In conclusion, on the one hand, there are modeling uncertainties preventing the quad scans from giving consistent results, and, more importantly, there may also be errors in the machine that are detrimental to the beam in a way that cannot be compensated for with empirical tweaks. Identifying the latter is an underappreciated benefit of a model-based approach. Ultimately, the goal must be to find and resolve the dominant errors in both categories. Systematic studies of the optics will become much easier if fluctuations (both slow drifts and AC components) in orbit and momentum can be substantially reduced. Especially when it comes to the booster, one must view its stability requirements in the context of the UITF, not that of CEBAF.

References

[1] M. McCaughan and E. Forman. *UITF Quick Reference Drawing rev. 14*. https://wiki.jlab.org/ciswiki/images/9/9c/UITF_quick_reference_rev.14.pdf.

[2] S. Gregory. *UITF MeV beam line as built*. https://wiki.jlab.org/ciswiki/images/2/25/Updated_MeV_songsheets_Shaun_Gregory_drawings_of_%22as_built%22_beamline.pdf. Jan. 2019.

[3] M. Bruker. *Some random musings about the gun kick and related issues*. unpublished. Aug. 2021.

[4] M. Bruker. *180 keV at 703*. https://logbooks.jlab.org/entry/3899879. Sept. 2021.

[5] M. Bruker. *Simplistic QCM energy gain model*. https://wiki.jlab.org/ciswiki/images/1/17/Max%27s_thoughts_on_booster_GSet_calibration.pdf. Feb. 2021.

[6] M. Bruker. *Simplistic QCM energy gain model, this time including the buncher*. https://wiki.jlab.org/ciswiki/images/6/6c/Max%27s_thoughts_on_impact_of_buncher_on_booster_GSet_calibration.pdf. Feb. 2021.

[7] M. Bruker. *QCM cavity power vs. field setpoint measurement*. https://wiki.jlab.org/ciswiki/images/7/79/Max%27s_RF_power_measurements_related_to_nonlinear_behavior_of_2-cell_GSet.pdf. Apr. 2021.

[8] M. Bruker. *(Ab)using the QCM as a virtual cathode*. https://wiki.jlab.org/ciswiki/images/8/8c/20210408_Max_Virtual_Cathode.pdf. Apr. 2021.

[9] M. Poelker. *Chopper setup and e-bunchlength measurement*. https://logbooks.jlab.org/entry/3937600. Nov. 2021.

[10] R. Bachimanchi et al. *New CEBAF Booster Performance and Test*. Tech. rep. JLAB-TN-21-012. Jefferson Lab, 2021.

[11] M. Poelker. *booster 7 cell wandering GSet*. https://logbooks.jlab.org/entry/3940018. Nov. 2021.

[12] National Centers for Environmental Information. *Compute Earth's Magnetic Field Values*. `https://www.ngdc.noaa.gov/geomag/magfield.shtml`.

[13] M. Bruker. *Comparison MDLM504 cycled vs degaussed*. `https://logbooks.jlab.org/entry/3940227`. Nov. 2021.