


## Is Cryomodule Damping a Hoax?

### Background

There have been proponents, of whom this author may be the most irresponsible (responsible?), of the idea that by altering the phase of the PZT orbit<sup>1</sup> going into the cryomodule one can control its outgoing orbit amplitude and thus maximize the momentum damping by the cryomodule. The idea comes from the observation that the orbit goes through the first few cavities, where significant relative momentum increase occurs, at roughly the same phase and does not have time to “mix” position with angle as is normally the case with smooth acceleration. Thus the location in the phase space of this orbit can have a dramatic effect on how much damping in absolute amplitude it gets. In other words, if it is “pure angle”, it experiences maximal damping in its amplitude, while “pure position” will lead to no damping.

Considerable effort went into realizing this idea by characterizing the PZT orbit before the cryomodule, measuring the transport across the cryomodule, and numerically solving for the optimal quad settings at 5 MeV to create the right phase of the PZT going into the cryomodule. Several online tests have been carried out with varying degree of success, mostly limited to the area immediately after the cryomodule. The picture on the right<sup>2</sup> shows one such study where the red curve represents the original PZT orbit from the cathode to the Hall C target, and the blue one that after the optimal quad solution was downloaded. Indeed the blue orbit became much smaller after the cryomodule (arrow), but this effect did not persist. In the end the blue orbit usually had roughly the same amplitude as the red.



This prompted some rethinking. For an idea as simple as this it ought to work or at least show strong signs of promise on the first few tries. It did not happen. After scrutinizing the procedure going into the online tests and concluding that nothing was done wrong, I was forced to look hard into the basic premise of this idea. The following is what came out.

### Single Particle Emittance (or Courant Snyder Invariant) of an Orbit

The single particle emittance, or Courant Snyder invariant, of a particle trajectory  $(x,t)$  for position and angle is defined as


$$\mathcal{E}_0^2 = \mathbf{P} X^T \cdot \Sigma_D^{-1} \cdot X$$

$$X = \begin{pmatrix} x \\ t \end{pmatrix}, \Sigma_D = \begin{pmatrix} \beta_D & -\alpha_D \\ -\alpha_D & \gamma_D \end{pmatrix}$$

<sup>1</sup> This is the orbit excited by moving the laser spot on the cathode, thereby creating an almost pure position change in the launching orbit on the cathode.

<sup>2</sup> Caveats to this picture will be skipped here, but characteristically this is what I have been seeing in all the cases studied.

where  $\alpha$ ,  $\beta$ ,  $\gamma$  are the design Twiss parameters, and  $P$  the momentum, at the point of interest. The use of design Twiss parameters makes this definition dependent on the seemingly arbitrary choice of a design optics for transfer lines, as opposed to storage rings where the design optics has a more absolute significance. However as will be explained later, one can still bestow an almost unambiguous meaning to this quantity as a measure of the orbit amplitude when consideration of beam profile matching is incorporated into the picture.



In the picture to the right, the contours represent those defined by the design Twiss parameters at the point of interest at different Courant Snyder (CS) amplitudes. The trajectory represented by the red dot has a larger CS than the blue trajectory.

### *Absolute Orbit Amplitude*

This is worth mentioning only because we intuitively think there is something meaningful about the absolute amplitude of a particle trajectory  $(x, t)$ , possibly describable by something like

$$x^2 + c t^2$$

where  $c$  is an arbitrary position-angle conversion factor (roughly on the order of typical M12). Thus the smaller  $x$  &  $t$  are, the smaller the apparent orbit amplitudes. In the picture above the red trajectory has a smaller “absolute orbit amplitude” than the blue one. Part of the purpose of this note is to clarify that, without taking into account how the transport channel works, this notion of absolute orbit amplitude does not carry much significance. Any argument made based on this intuitive quantity implicitly assumes that the constant CS contours are round circles, or at least very well-behaved ellipses, at the point of interest, and the optics defined by such circles happens to be the right one for downstream transport.

One of the main themes of this note is to clarify when such intuitive interpretation is valid for predicting the behavior of a trajectory, and when this view is compatible with the Courant Snyder picture.

### *Transport of Twiss Parameters across an Impulse Acceleration*

The transport of Twiss parameters across an impulsive acceleration leading to momentum gain  $\delta P$  can be easily evaluated. The transfer matrix for such a system is

$$M_{acc} = \begin{pmatrix} 1 & 0 \\ 0 & R \end{pmatrix}, \quad R = \frac{P}{P + \delta P}$$

We can easily work out how emittance and Twiss parameters are transported across this based on

$$\Gamma' = M_{acc} \cdot \Gamma \cdot M_{acc}^T, \quad \varepsilon' = R\varepsilon$$

$$\Gamma = \varepsilon \cdot \Sigma = \begin{pmatrix} \langle xx \rangle & \langle xt \rangle \\ \langle xt \rangle & \langle tt \rangle \end{pmatrix} = \begin{pmatrix} \varepsilon\beta & -\varepsilon\alpha \\ -\varepsilon\alpha & \varepsilon\gamma \end{pmatrix}$$

Note that unlike Twiss transport at constant momentum, the emittance, which serves as an overall normalizing factor, is also changed, leading to the following transformation of the Twiss parameters

$$\begin{aligned}\beta &\rightarrow \beta/R \\ \alpha &\rightarrow \alpha \\ \gamma &\rightarrow \gamma R\end{aligned}$$

This is different from what one would conclude from an “adiabatic” damping scenario where the Twiss parameters are “unchanged” (after applying RF focusing and drifts) across the momentum gain and only the emittance is scaled by R. If an optics code does not model this correctly by invoking the above transform, a modeling error would result leading to mismatch between design and actual beam parameters. Given the definition of the Twiss parameters not merely as a formal bookkeeping device, but description of the moment distribution of the physical beam here, the correct transformation is crucial in modeling the real optics.

### *Courant Snyder Parameters of Pure Position and Pure Angle Trajectories across The Impulse Acceleration*

We can calculate the momentum independent CS parameters, or single particle emittance, as defined in the beginning, of a pure position trajectory  $(x,0)$ , and a pure angle trajectory  $(0,t)$ , before and after the impulse acceleration with no regard to the detail of the Twiss parameters other than requiring that they transform as specified above. Across the impulse acceleration, the pure position trajectory does not change, while the pure angle trajectory damps down to  $(0,tR)$ . It is trivial to apply the above algebra to show that

$$\begin{aligned}\mathcal{E}_{Position}^2 &= P_0 \gamma_0 x^2 \\ \mathcal{E}_{Angle}^2 &= P_0 \beta_0 t^2\end{aligned}$$

both before and after the acceleration. The subscript 0 indicates quantities before the acceleration. Thus the CS parameters do not change across the acceleration for a given trajectory, whether it is pure position or pure angle. Furthermore, if these 2 trajectories had the same CS values, such as originating from the same trajectory before going through different transport optics before the acceleration, they will have the same CS after the acceleration. Indeed this only demonstrates a well-known fact about CS invariants derivable from formal arguments. The formulas given above are sufficient for this. We are only working out the detail here for the special case of impulse acceleration since the intuitive picture is a little tricky.

### *Courant Snyder Parameters of Pure Position and Pure Angle Trajectories across The Impulse Acceleration Based on Incorrectly Calculated Twiss Parameters*

Suppose, instead of correctly modeling the Twiss parameter transform across the impulse acceleration, one used the wrong formula for adiabatic damping

$$\begin{aligned}\beta &\rightarrow \beta \\ \alpha &\rightarrow \alpha \\ \gamma &\rightarrow \gamma \\ \varepsilon &\rightarrow \varepsilon R\end{aligned}$$


then one can derive the wrong transformation of the CS parameter of the pure position and pure angle trajectories easily as

$$\begin{aligned} \mathcal{E}_{Position}^2 &: P_0 \gamma_0 x^2 \rightarrow P_0 \gamma_0 x^2 / R \\ \mathcal{E}_{Angle}^2 &: P_0 \beta_0 t^2 \rightarrow P_0 \beta_0 t^2 R \end{aligned}$$

The CS parameter for pure position is now larger than pure angle CS. I had some email communication with Valeri Lebedev recently on Optim modeling of transfer matrices with momentum gain. My impression is that Optim treats all acceleration as adiabatic in deriving Twiss parameters across cavities. If this is true and one takes the Twiss parameters calculated this way to perform downstream matching, there will be a discrepancy from reality. We need to be careful here.

### Trivial Case

Consider the trivial case where the orbits going into the impulse acceleration are given by  $(\sqrt{\beta}, 0)$  and  $(0, \sqrt{\gamma})$  where  $\beta$  and  $\gamma$  are the design Twiss at this point, and  $\alpha=0$ . After the acceleration the orbits become  $(\sqrt{\beta/R}, 0)$  and  $(0, \sqrt{\gamma R})$ . Using the formulas above for correct and incorrect Twiss transformation, one gets the picture on the right. The blue orbit undergoes damping by a factor of  $R$ , while the red orbit stays the same. The black circle is the constant-CS contour before acceleration. The two green ellipses are the correct constant-CS contours after acceleration at different CS values. The two final orbits sit on the smaller contour with radius reduced by  $\sqrt{R}$ . After the momentum normalization the CS values of these two orbits become the same as before, namely  $2P_0\beta\gamma$ . The magenta circle shows the constant-CS contour after acceleration calculated by the wrong Twiss transformation, which is a simple scaling of the original. In the last case one gets the wrong impression that the pure angle orbit has a smaller CS value than pure position after acceleration.




### Betatron Mismatch

Given a transport channel with design Twiss parameters  $\alpha_D, \beta_D, \gamma_D$  defined at every point, and those for beam spot  $\alpha_B, \beta_B, \gamma_B$ , one can define the dimensionless betatron mismatch factor  $\lambda$ :

$$\lambda = \frac{1}{2}(\beta_B \gamma_D - 2\alpha_B \alpha_D + \gamma_B \beta_D)$$

The physical picture of this quantity is as follows. At every point one can go to the “normalized” phase space where the constant-CS contours defined by the design Twiss parameters are concentric circles as shown in blue on the right. In this normalized space the constant-CS contours defined by the beam Twiss parameters are ellipses with semi-major axis length given by  $\lambda$  as shown in red. Thus  $\lambda$  quantifies the degree of mismatch between the beam and the design optics. The larger  $\lambda$  is, the worse the mismatch. Once  $\lambda$  is known at a point, it is fixed



throughout the rest of the line. The same picture persists downstream with only different orientation of the ellipse from point to point.

### *Significance of Design Twiss Parameters*

I have been using the design Twiss parameters freely in deriving my conclusions so far. One might ask whether this is justified since after all we are dealing with a transfer line and to some extent the design Twiss is arbitrary. I will argue here that this is true only if one does not care what the transport optics does to the beam spot itself. Assuming sanity of the person designing a transfer line, and keeping in mind that we are dealing with global transport from cathode to target in a single piece, Twiss parameters corresponding to the design beam spot is the most natural and useful choice with predictive power. All discussion of CS parameters should be based on such a definition, and I will be using this definition to illustrate my point.

I have shown in the above that the pure position and pure angle trajectories have the same CS values after impulse acceleration, or fall on the same constant-CS contours, if they started out with the same CS values. One might ask whether this is merely artificial since I have the freedom to define what the design Twiss parameters are after the acceleration. In other words, can't one always find a set of "design" Twiss parameters such that these two trajectories always lie on the same contour, or even lying on contours arbitrarily chosen for each trajectory respectively? The answer is yes and no. Up to a point allowed by physical consistency this is indeed the case, but it is so only if one ignores what happens to the global transport of the beam spot and gives up hope of predicting actual orbit behavior. Allow me to elaborate.

All transfer lines terminate at a point, say the target, or simply entrance to a FODO lattice, where a desired beam spot is unambiguously defined, and so are the design Twiss parameters. The argument in the following depends on one more "sanity" assumption, namely that the Twiss parameters at the end point should represent a well-behaved beam profile, or that there is no excessive position-angle correlation:

$$\alpha \approx O(1), \sqrt{\beta\gamma} \approx O(1)$$

This simply means the phase ellipse of the beam at the end is not grossly sheared due to bad matching, which should be easily satisfied by any sanely designed system. At any point **P** upstream, one can take the transport optics between this point and the end point, and unambiguously define as design Twiss parameters at this point those that would lead to the design Twiss at the end point given the intervening transport. One should use Twiss parameters thus defined to calculate CS invariants of a trajectory, since this CS value, being an invariant, can be readily used at the end point to safely connote the intuitive physical orbit amplitude mentioned earlier without much caveat. If the CS value thus defined at point **P** is small, it implies that the absolute orbit amplitude is most likely small, or at least bounded by a well-behaved ellipse in the phase space, at the end point. If on the other hand one chooses to use an arbitrary set of Twiss parameters as design at point **P** to calculate the CS, indeed one can almost get any arbitrary CS values, but such CS values are with little predictive power, since now the corresponding constant-CS contour at the end point is also arbitrary and very likely displaying gross position-angle correlation. A small CS value no longer means small absolute orbit amplitude and vice versa. Of course the real orbit amplitude at the end point is independent of the choice of design Twiss, but different choice of design Twiss gives different predictive power of the CS values.

If one always wants to have good predictive power with CS values calculated from an arbitrary set of Twiss parameters, the intervening optics has to be modified accordingly to map the arbitrarily chosen Twiss at point **P** to a well-behaved ellipse at the end point. This effectively amounts to re-matching the optics. Such

process can indeed bring an arbitrary set of trajectories to the desired amplitude at the end, but it will likely destroy the original propagation of the beam spot and cause it to be mismatched at the end.

The above arguments can be summarized as follows:

- Twiss parameters derived from desired phase space distribution at the end point and intervening transport optics serve as the best design Twiss for predicting behavior of trajectories at the end.
- This set of Twiss parameters should be consistent with the beam Twiss parameters in a sane optical design. The mismatch factor  $\lambda$  between the two should be of order 1.
- To see whether the pure position and pure angle trajectories reside on the same constant-CS contour at the end, one only needs to see if they do at any point based on the design Twiss defined above.
- If one uses a set of different Twiss parameters, the resulting CS invariants are not very useful in predicting behavior of trajectories at the end.

We will illustrate the above algebraically and with numerical examples below.

### *Twiss Parameters Leading to Arbitrarily Given CS Invariants for Two Arbitrarily Given Trajectories*

Generalizing on the example with pure position and pure angle, we consider two arbitrary trajectories  $\mathbf{X1}=(x_1, t_1)$  and  $\mathbf{X2}=(x_2, t_2)$  at point  $\mathbf{P}$  and ask what set of Twiss parameters will cause  $\mathbf{X1}$  to have a CS invariant that is  $\mathbf{K}$  times larger than that of  $\mathbf{X2}$ . Some simple algebra yields the relation that must be satisfied by such Twiss Parameters:

$$\beta(t_1^2 - \mathbf{K}t_2^2) = \alpha(\mathbf{K}t_2x_2 - t_1x_1) \pm \sqrt{(1 + \alpha^2)\mathbf{K}(t_2x_1 - t_1x_2)^2 - (\mathbf{K}t_2x_2 - t_1x_1)^2}$$

The first obvious thing is that, if  $\mathbf{X1}$  and  $\mathbf{X2}$  are linearly dependent, there is no solution for  $\alpha$  and  $\beta$  in real numbers. This is intuitively obvious since as long as  $\mathbf{X1}$  and  $\mathbf{X2}$  are linearly independent, we can imagine drawing an ellipse through them. But if they are linearly dependent, they must belong to ellipses of different radii and thus have different single particle emittances. The criterion on  $\mathbf{X1}$  and  $\mathbf{X2}$  is actually more stringent. Take the quantity inside the square root, basically the discriminant, and demand that it be greater than or equal to zero, we get, in order for  $\beta$  to be real:

$$\alpha^2 \geq \alpha_{Min}^2 = \text{Max} \left( \frac{(t_1^2 - \mathbf{K}t_2^2)(x_1^2 - \mathbf{K}x_2^2)}{\mathbf{K}(t_2x_1 - t_1x_2)^2}, 0 \right)$$

$$\beta_{Min}^2 = \left| \frac{(t_1x_1 - \mathbf{K}t_2x_2)^2 (x_1^2 - \mathbf{K}x_2^2)}{\mathbf{K}(t_2x_1 - t_1x_2)^2 (t_1^2 - \mathbf{K}t_2^2)} \right|$$

This further constrains the relation between  $\mathbf{X1}$  and  $\mathbf{X2}$  for a given  $\alpha$ , or if  $\mathbf{X1}$  and  $\mathbf{X2}$  are given, sets the lower bound  $\alpha_{Min}$  on allowable  $\alpha$ , with corresponding  $\beta_{Min}$ . Since a large  $\alpha$  implies strong position-angle correlation, or a strongly sheared phase ellipse, if  $\mathbf{X1}$  and  $\mathbf{X2}$  are such that they make  $\alpha_{Min}$  very large at point  $\mathbf{P}$  and if we want to force them to have the same CS value there, a very distorted phase ellipse will be needed. From this point on we have 2 options as alluded to earlier:


- Stay with the existing design transport optics from  $\mathbf{P}$  to the end. The propagated phase ellipse at the end

will be so distorted as not to have any predictive power on whether these 2 trajectories have comparable absolute orbit amplitudes. Nothing physical has changed in this case, of course.

- Change the transport optics to transform this distorted ellipse at point **P** to a well-behaved one at the end, where one can then ensure that the 2 trajectories do have comparable absolute orbit amplitudes. But in the process the beam spot, which presumably is well behaved at **P**, becomes grossly distorted at the end. Basically we have fixed the transport of the distorted trajectories at the expense of the beam spot.


From the expression of  $\alpha_{\text{Min}}$  one also sees that if the 2 trajectories are one pure position and one pure angle, there is no constraint on  $\alpha$  and one can always find an upright phase ellipse that satisfies the arbitrary CS ratio **K** between the 2 trajectories<sup>3</sup>.

What is shown above is the fact that one does not really have the freedom to choose design Twiss parameters to make 2 arbitrary trajectories have the same CS invariants at the end, and in turn demonstrate that they have roughly the same absolute amplitude. The fact that I did show that the pure position and pure angle trajectories have the same CS invariants after the impulse acceleration is not accidental, but due to the fact that the Twiss parameters used for calculating the CS invariants do coincide with the natural, design Twiss parameters that take the beam spot after the impulse acceleration to a well behaved spot on target<sup>4</sup>, as we must have by assuming sanity of whoever designed the transport optics.



It may be useful to give some numerical examples. Consider 2 trajectories  $X1=(0.1, 0.1)$  and  $X2=(-1, 1)$  at point **P** as shown on the top right corner. X1 inherently was generated at a much smaller amplitude than X2, but we want to find a set of Twiss parameters that makes X1 have 10 times larger (**K=10**) CS value than X1. This gives an  $\alpha_{\text{Min}}$  of 15.7956. The ellipses shown correspond to  $\alpha=16, 18, 20, \dots, 30$  that all satisfy this and all pass through X1 and envelop X2 deep inside. The numbers displayed are  $\alpha$  and  $\beta$  for these possibilities. They are all very skewed<sup>5</sup>.

We can look closer into the behavior of the phase ellipse corresponding to the minimal  $\alpha_{\text{Min}}$ . This is shown on the right in black. In the mean time by the sanity assumption since X1 was really generated at a smaller amplitude than X2, the beam spot itself, coming out of the same transport process, takes on a



<sup>3</sup> This does not mean this ellipse is the most natural, since the aspect ratio can still be off.

<sup>4</sup> Or simply to the entrance of the North Linac, which is sufficient for this argument.


<sup>5</sup> This is of course relative to the sanity assumption that the beam spot itself is well behaved at this point. If on the other hand the ellipses represent the real beam spot, then the upstream transport must be very skewed. X1 as is indeed may have been generated with amplitude 10 times larger than X2 earlier, and after some downstream fixing of the spot, X1 & X2 will lie on well behaved but different phase ellipses with a radius ratio of 10 (really square root of 10).

well-behaved ellipse shown in green with  $\alpha = -1$  and  $\beta = 1$ . This strongly suggests the true relative physical amplitudes of X1 & X2, with the former much smaller than the latter.

Now by the sanity assumption suppose this system is propagated to the end point by a gentle transport, without loss of generality, an identity transform, the picture at the end looks just the same. X1 still has 10 times greater CS than X2. But not much can be said about the absolute orbit amplitudes from this since the constant-CS contours are so skewed. We don't have much predictive power.


Now we want to force the skewed ellipse at point **P** to a well-behaved one at the end by changing the intervening optics, such that the CS values of X1 & X2 at **P** can be readily used to represent real absolute orbit amplitudes at the end. To do this we find the transform that takes the skewed phase ellipse at **P** with  $\alpha = 15.7956$  and  $\beta = 15.8272$  to an upright one with  $\alpha = 0$  and  $\beta = 1$  at the end. The matrix that does this is, for example<sup>6</sup>,

0.1, t1 -> 0.1, x2 -> -1, t2 -> 1, RR -> t




$$\begin{pmatrix} 0.251361 & 0. \\ 3.97039 & 3.97834 \end{pmatrix}$$

This transport indeed forced the skewed phase ellipse to a circle at the end, and X1 now indeed has both CS value and the true physical amplitude 10 times larger than X2 as shown on the right in true aspect ratio. But the price paid is obvious: the beam spot (green) has been virtually destroyed.



0.1, t1 -> 0.1, x2 -> -1, t2 -> 1, RR ->



Now if we take the more rational approach by recognizing the true origin of amplitude difference between X1 & X2, and look for Twiss parameters at **P** that make X1 have a CS value 20 times smaller than X2. This results in the phase space pictures at **P**, and at the end, shown on the left. Now  $\alpha_{Min}$  is only 0.3535, leading to a very well behaved ellipse (black) satisfying the requirement. The beam spot shown in green is still the same as before. Now an identity transform to the end will leave everything unchanged as in the previous example, but now since the constant-CS ellipses are much better behaved, they have more predictive power, and a small CS invariant can be safely taken to mean a small absolute orbit amplitude.

The matrix needed to transform this ellipse at **P** to a round circle at the end is<sup>7</sup>

$$\begin{pmatrix} 0.970984 & 0. \\ -0.343295 & 1.02988 \end{pmatrix}$$

This takes the phase space at **P** to the picture at bottom left. The beam spot (green) is still well contained.

<sup>6</sup> A somewhat useful concept in evaluating the transfer matrix property is the SVD condition number. Without going into detail, this number for a transfer matrix should be not much greater than the typical M12 of the system. In the example given the SVD condition number for the transfer matrix is about **32** while the implicit M12 was made of order 1, so the matrix is not entirely non-singular, meaning the transport has to work hard, depending on fine cancellations, to achieve the goal specified.

<sup>7</sup> In contrast the SVD condition number for this matrix is **1.41**, a very mild transport.


Finally we can use the invariant betatron mismatch factor defined earlier to evaluate these 2 choices of the design Twiss. In the first example it is 39.5, and in the second it is 1.94. Clearly the second is a much more natural choice for the design Twiss, both for interpreting CS values and for matching the beam spot.

All the above arguments serve to demonstrate that in making the statement about relative CS ratios between 2 trajectories one does not have too much freedom in choosing the design Twiss to suit his devious intent, if he wants to have any predictive power at the end. The Twiss parameters of the beam are most likely the right choice. The argument made earlier about the pure position and pure angle trajectories having the same CS invariants based on design Twiss parameters with real predictive power did not come about by accident. It is rather the only correct way of interpreting the underlying dynamics.

### Phase Ellipse through Pure Position and Pure Angle

In contrast to the previous example, consider the two trajectories  $X_1=(1, 0)$  and  $X_2=(0, 1)$  at point **P** as shown on the right. There is no limit to how small  $\alpha$  can be. One can use any of the  $(\alpha, \beta)$  pairs displayed on the right (with  $\alpha$  between  $-5$  &  $+5$ ) to define the constant-CS contour shown in the graph, for example. The most sane choice would be  $\alpha=0$ , although the aspect ratio of the ellipse should still be such that it is consistent with the beam spot at **P**, or in other words, the mismatch factor  $\lambda$  between the beam spot and the chosen ellipse should be very close to one. The algebra I used to show equal CS

-5.000000	5.099020
-4.000000	4.123106
-3.000000	3.162278
-2.000000	2.236068
-1.000000	1.414214
0.000000	1.000000
1.000000	1.414214
2.000000	2.236068
3.000000	3.162278
4.000000	4.123106
5.000000	5.099020



between pure position and pure angle after impulse acceleration used design Twiss that makes the mismatch factor exactly one, making that argument as non-pathological as possible.

### Can We “Shape” The Beam Spot at a Point to Influence the Orbit Amplitude at the End?

One might wonder, given the above arguments, if there is a way to shape the beam spot at point **P** based on the trajectory we want to selectively suppress, such that when the spot is restored to a nice round spot at the end, the transport needed would automatically make one trajectory much smaller in terms of absolute orbit amplitude than another. The answer is yes if we are talking about 2 completely independent trajectories. Unfortunately the problem we have here is one of trying to turn one (pure position) trajectory into another (pure angle) by varying upstream optics. The two originate from the same trajectory and had the same CS value to begin with. Changing the beam spot at **P** will not help in this case since it will disrupt this orbit conversion at the same time.

### What Might Actually Work (But Has Nothing to Do with Momentum Damping)

Despite the previous paragraph, the following scheme may actually lead to small PZT orbits at the end, although it has nothing to do with momentum damping. Since PZT orbits are mostly more position than angle in the phase space at the cathode, artificially making the design Twiss at the cathode in such a way that the constant-CS contour is stretched out along the position axis (i.e., large  $\beta$  & small  $\gamma$ ) and using

downstream optics to match it to a well-behaved ellipse on target can in principle greatly reduce the PZT amplitude on target. Maybe this is what we already do. Of course there is one caveat: if the beam spot has nontrivial angular divergence on cathode, it can be grossly mismatched on target. Again we are more or less reiterating the same theme of this note.

### *Why Does the Orbit Grow after Acceleration?*

We observed missing damping through the cryomodule manifested by the lack of reduction in absolute orbit amplitude after the cryomodule. But from the arguments above, if a pure angle orbit does not damp more in terms of CS, why should the pure position orbit damp less, as is suggested by observation? I suspect this is purely attributable to betatron mismatch, and to some extent XY coupling. As mentioned earlier, Optim modeling of impulse acceleration may not be adequate. This can lead to design optics not matching the real beam. The situation of XY coupling across the cryomodule is also murky at this point. My belief is that, to address this problem, modeling effort of this area at an unprecedented level is required.

### *Haven't We Already Matched the Beam Spot at 60 MeV into the North Linac?*

This is a tricky question. If the spot and the PZT have the most compatible orientations in phase space, wouldn't restoring one automatically rein in the other? I can think of two counter arguments, and may be missing the real, simple, correct one:

- Non-symplectic process is involved: To correctly address this problem the global transport from cathode to the 60 MeV line must be considered since we are looking at the transport of the PZT starting at the cathode. Beam profile matching at 60 MeV is a local process disregarding what happened before. All my argument depends on linear symplecticity over the entire range. Deviation from this (space charge, scraping, nonlinear focusing, longitudinal dynamics, .....) will likely cause the beam distribution (spot) to transform under a different map than the centroid (trajectory). In such a case the statement that "Trajectories starting with the same CS as defined by beam spot Twiss should end up with the same CS" is not valid. The matrix algebra leading to this statement does not hold any more. Physically if a mechanism can cause the orientation of the spot in the phase space to "break" from that of the PZT anywhere between the cathode and 60 MeV, then matching the spot may even have an adverse effect on the PZT amplitude.

Note this does not invalidate my argument for the futility of trying to alter the amount of damping by changing PZT phase entering the cryomodule at 5 MeV. We have the beam and a single PZT orbit at the beginning of 5 MeV, with very little space charge, scraping, and nonlinearity downstream. If we want to change this single PZT trajectory at 5 MeV and hope it will end up with a smaller CS value after the cryomodule, it will not happen because of violation of symplecticity. I am only saying that one cannot say by re-matching beam spot at 60 MeV, we have undone in a symplectic manner all the phase space distortion starting from the cathode on both the trajectory and the beam equally.

- The beam profile matching at 60 MeV pays no heed to XY coupling (in fact another violation of 2D symplecticity) in the beam. We know the cryomodule generates considerable XY coupling, and the picture is further complicated by the XY coupled PZT orbits. If we perform the naïve 60 MeV beam profile matching with no regard to the XY correlation in the beam, it is hard to say what happens to the centroid.

### *Phase Trombone at 5 MeV?*

Of course we can still build a phase trombone at 5 MeV to vary the PZT phase into the cryomodule without changing the spot matching. This will tentatively change the absolute orbit amplitude at the exit of the

cryomodule. But I hope I convinced you by now that this does not put the PZT at a constant-CS contour with a smaller radius, but merely moves its phase around in the same contour. The ultimate criterion still rests with what the absolute orbit amplitude is on the target, which depends on the final phase of the PZT on a fixed contour there. But then what is the point of constructing such a long-range phase trombone, when we can have one right before the target?

I have constructed a Mathematica-based phase trombone using the six 5 MeV quads. It can change the X & Y phase independently while maintaining the betatron match. I am waiting for the next MD to obtain realistic beam Twiss parameters for input to this procedure. However, I do not see the point for doing this.

### *More on Optim Modeling of Impulse Acceleration*

As I said my impression on how Optim treats impulse acceleration came from my communication with Valeri. The question I asked him specifically was about the way Optim incorporates an externally defined transfer matrix. His answer did include statement that “the user is responsible for making the transfer matrix symplectic”. If this means that for an impulse acceleration the user needs to scale the entire matrix by  $\sqrt{R}$  before incorporating it into Optim, then direct propagation of the Twiss parameters is indeed correctly handled. If this is the case, which we have to find out from the code itself or ask Valeri, then most likely my suspicion of Optim is premature.

This however does not mean that the Injector is well matched, and the problems we are seeing are not from this mismatch, of course.

### *A Thought Experiment*

To illustrate my points so far, I will present a thought experiment below. I hope someone can point out a flaw in the argument so I can plunge back into this task with renewed enthusiasm (no sarcasm intended). I will summarize the main points of this note as well as make recommendations after this example.

The thought experiment consists of a starting point where the beam is formed and a centroid trajectory is excited. It goes through a matching-phase-trombone region followed by a single point acceleration causing a momentum increase by a factor of R. The upstream matching was carried out so as to make the design Twiss describe a round circle after the acceleration. The beam then goes through a FODO channel with M12 roughly equal to 1 m, and a phase trombone again before reaching the target, where the Twiss parameters again describe a round circle.

I'll walk through steps A-E in each of the following scenarios:

#### **Scenario 1:**

- A.** Beam is represented by the orange dot below. The trajectory lies on a phase ellipse described by the chosen design Twiss. The design twiss and beam twiss have mismatch factor  $\lambda$  close to 1.
- B.** Before the acceleration matching was done to ensure a slightly elongated ellipse for the design Twiss such that after acceleration it is squashed down to a circle. In the mean time 2 different settings of the first phase trombone can make the centroid either pure position (blue) or pure angle (red).
- C.** After the acceleration. The two trajectories still lie on the same constant-CS contour.
- D.** In the FODO channel each centroid revolves around the same phase ellipse.

- E.** On target the second phase trombone brings both trajectories to exactly the same point (at two different settings).

Note that the entire manipulation could have been accomplished by the second phase trombone alone.

### Scenario 2:

Here nothing is changed physically, but we choose to reset the design Twiss after the acceleration to artificially make the pure angle orbit lie on an ellipse with a smaller radius than the pure position orbit.

- A.** Same as before.
- B.** Same as before.
- C.** After the acceleration, design Twiss is artificially changed<sup>8</sup> to make the red dot appear to have a smaller CS value.
- D.** In the FODO channel each centroid revolves around its own phase ellipse at different (artificial) CS amplitudes.
- E.** On target the physical locations of the two centroids are still the same as before, namely at the same phase space point. Since the matching section no longer maps the new artificial Twiss ellipse into a circle, the new Twiss ellipses at the target are skewed, from which not much information can be derived regarding absolute orbit amplitudes based on CS values. In the mean time the phase trombone is no longer matched to the new Twiss ellipses either, thus the two different phase trombone settings lead to different Twiss ellipse orientations on target. Nothing is gained or lost, except that we can no longer predict the behavior of these 2 orbits based on their CS values.

### Scenario 3:

Here instead of the FODO lattice after the acceleration, we implement special matching to bring the artificially defined Twiss in **scenario 2** to a round circle on target.

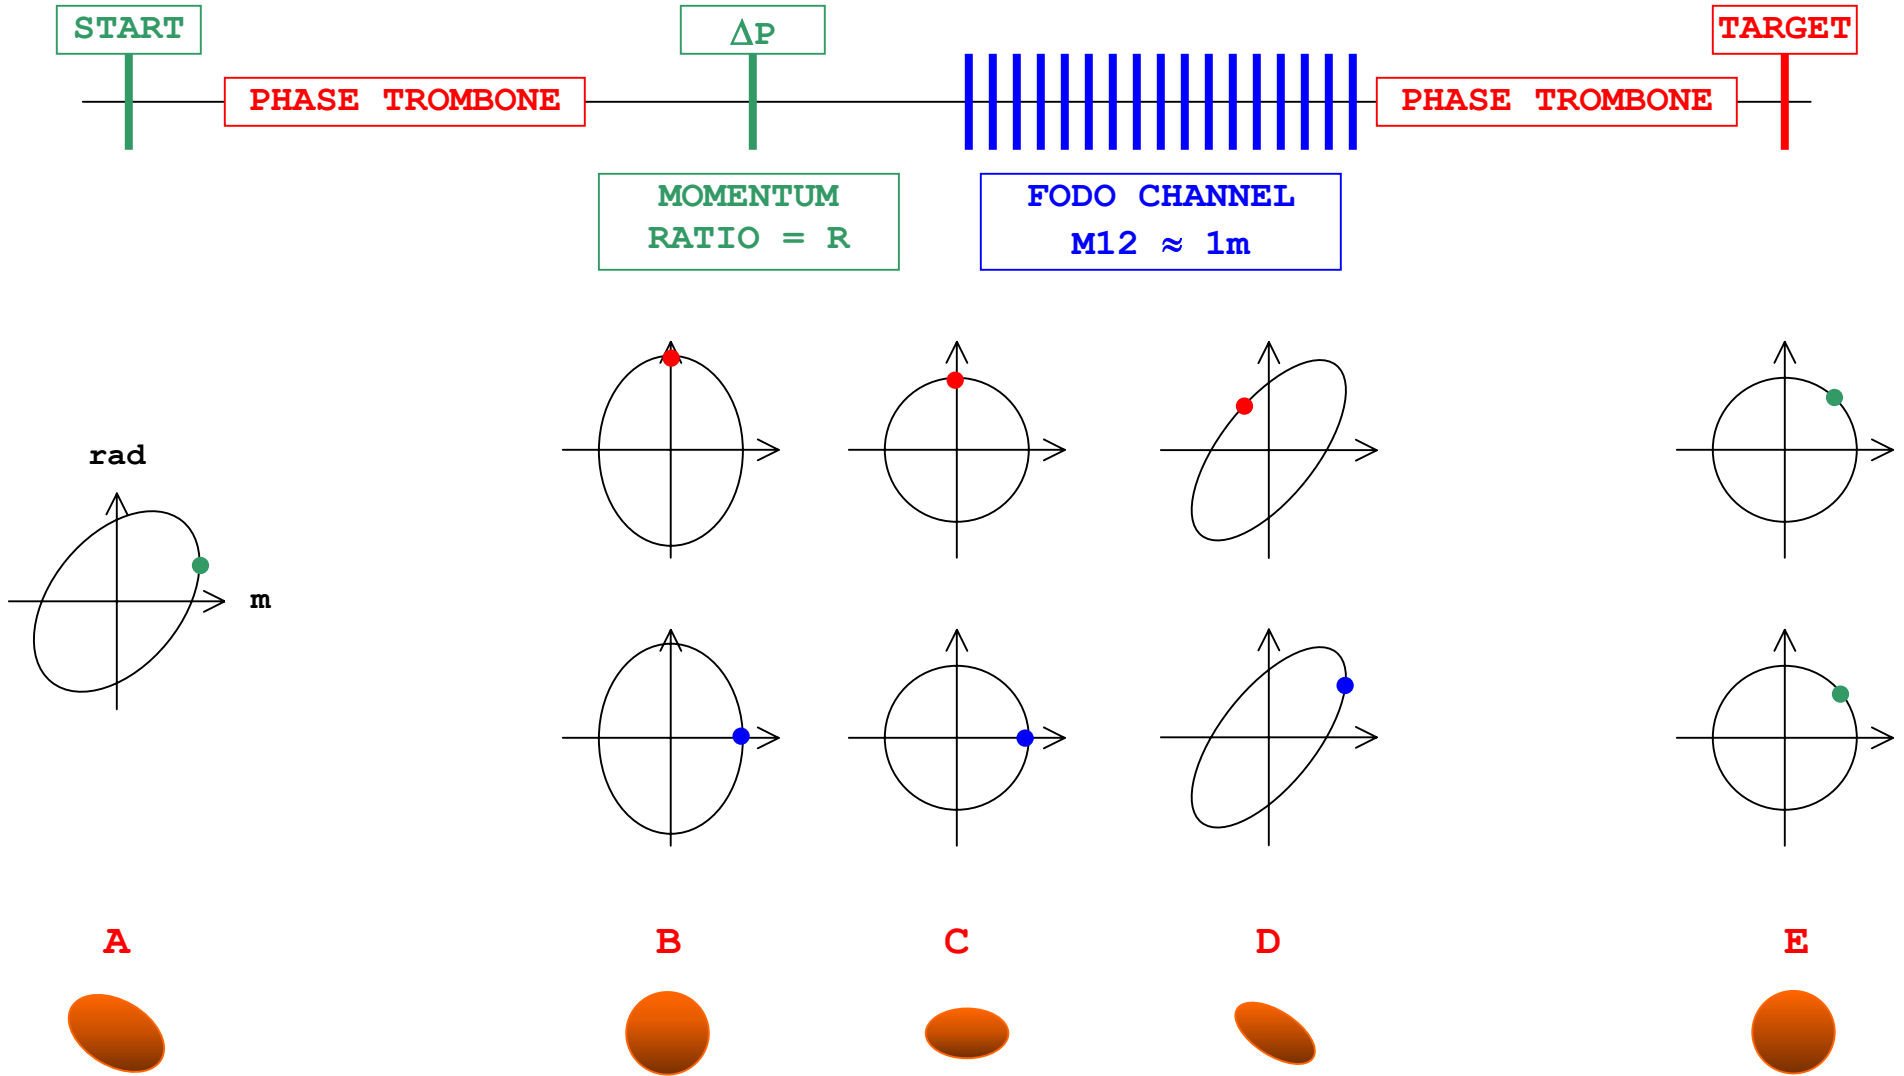
- A.** Same as before.
- B.** Same as before.
- C.** Same as **scenario 2**.
- D.** Special matching is performed to map the artificial Twiss ellipse back to a circle on target.
- E.** Indeed we have succeeded in reducing the real absolute amplitude of the pure angle orbit with respect to pure position, and the CS values are a useful indication for this since the final design Twiss is well behaved. However, at the same time the beam spot itself has become grossly mismatched by the sanity standard.

Note that in all the above discussions acceleration played no distinct role. The momentum change did not, as might have been wrongly expected, change the phase space contour a particular trajectory sits on. All that is brought about by the momentum change could be achieved by magnetic optics alone so far as the location of a trajectory in phase space is concerned.


---

<sup>8</sup> It is interesting to note that if the optics code transforms the Twiss parameters incorrectly across the acceleration as described earlier in this note, then it will do exactly what's outlined here. If this code is further used to match downstream optics, something similar to **scenario 3** below will happen.


# 1. A Thought Experiment



## 2. Same Optics, But with Design Twiss Reset Artificially after Acceleration



### 3. Downstream Optics Changed to Re-match the Trajectory at the End



## Conclusion & Recommendations

- Contrary to intuition, an impulse acceleration does not offer extra handles on enhancing momentum damping. Apart from the normal reduction in emittance after acceleration, any manipulation to exploit the betatron phase for this purpose is futile, unless one ignores what happens to the beam spot or sanity of the design phase space at the downstream match point (target, entrance to North Linac, etc.). Even in the latter case the outcome is indistinguishable from that achievable by magnetic optics alone.
- A phase trombone using 6 quads in the 5 MeV line will be tested once beam profile there is measured in the upcoming MD. According to what's discussed in this note, this will amount to no more than a long-range phase trombone controlling the position of the PZT signature along a (hopefully well behaved) constant-CS contour on target, but not moving it from contour to contour. It would act largely like a simple phase trombone in the Hall and serve a function redundant to the latter. Despite this, it is worth a try to empirically verify it one way or the other.
- The best place to devote our limited resource, in my opinion, is modeling, both theoretical and empirical, of the Injector transport up to North Linac with accuracy not achieved up to now, followed by detailed correction even if it may imply additional skew quads. This seems to be an absolute necessity given the PZT blowup across these sections, the difficulty encountered so far in accurately modeling and correcting the Injector optics, and inability to explain it otherwise. Empirical data shows that transport across the cryomodule is far from ideal both in-plane and cross-plane, as can be seen from the difference orbit measurement recently (next page). Recently measured transfer matrix across the cryomodule looks like this:

$$\begin{pmatrix} -1.52432 & 2.46252 & 0.012997 & 0.689894 \\ -0.0938983 & 0.0808547 & 0.00948916 & 0.030035 \\ -0.00311391 & 0.476694 & -0.848993 & 1.13378 \\ 0.011086 & -0.00799127 & -0.0173357 & -0.0685924 \end{pmatrix}$$

Situation of in-plane transport across the cryo-unit may need to be examined too. Efforts have been made to characterize these areas and put them in the model. This may need to continue at a deeper level. There are other areas in the Injector where increased level of modeling accuracy may be necessary in light of the unforgiving PZT transport observed. This will be a difficult task requiring not only hard work, but also innovative ideas. But I suspect this is the only solution to our problem.

- For acceleration at the low end, calculation of Twiss parameters does not merely amount to applying RF focusing and then scaling everything by the same damping factor. If Optim is indeed ignoring the effect of impulse acceleration on Twiss parameters<sup>9</sup>, at least we should make sure the Model Server does not. From discussion with Yves it appears that the Model Server is doing things from first principle. Maybe then at the low energy end the Model Server should be driving Optim, as opposed to the other way.
- Also, we have been traditionally following the practice of section-wise GOLD Twiss parameters. Namely, we have been resetting the GOLD Twiss to the design value at the beginning of each machine sub-section (Injector, Linac 1, Arc 1, etc.). Bridging gaps between sections at the high energy end has been taken care of by 30 hz Courant Snyder or Automatch, but little attention has been paid to such gaps at the low end. This is possibly why modeling issues such as discussed here never

---

<sup>9</sup> All I have is some email communication with Valeri. I can be wrong about Optim in the end.

caught much attention. With the need to understand damping of the PZT from cathode to target in a single piece, the luxury to do section-wise modeling is called into question.

I believe the last 3 points are important in the context of planning on an eventually completely model-driven accelerator operation as has been discussed in many recent meetings.

### *Explanation of Graphs on the Next Page*




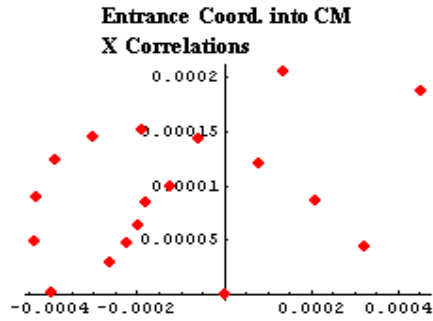
The graphs below show distribution of difference orbits in the phase space launched during a recent measurement. The top portion shows those launched before the cryomodules and propagated to the entrance of the first cavity of the first module. The bottom portion shows those derived from 60 MeV BPMs and back propagated to the exit of the last cavity of the second module. Plot labels are self explanatory, with XY-14 and XY-23 representing the X-Y' and X'-Y spaces respectively. It can be seen that with input of evenly covered trajectories in all dimensions, the outgoing trajectories are strongly coupled in the X-X' space and somewhat coupled in Y-Y', while significant XY coupling also developed across these cryomodules. All these can lead to blowup of difference orbits even if the phase space is correctly damped.

It is interesting to point out that not only the outgoing phase space distribution is skewed, so is the distribution going into the cryomodules. This comes from the fact that the difference orbit data were acquired using FOPT with Autoscale turned on. The original FOPT orbits were carefully tuned to cover a uniform semi-circle in the first 2 quadrants of the phase space. But when Autoscale is turned on, the pre-loaded amplitudes of difference orbits launched in the second quadrant of both X & Y spaces caused beam loss downstream and were automatically reduced by FOPT. Some orbits launched in the first quadrant at the same pre-loaded amplitudes apparently caused no beam loss. One can see that only a narrow slice of the phase space of the incoming orbit admitted larger amplitudes. This is actually a measure of the so-called **admittance** of the downstream transport system. The fact that it is also skewed is indirect indication of bad matching somewhere downstream<sup>10</sup>.

---

<sup>10</sup> Bad orbit, however, can cause similar effects.

## PHASE SPACE DISTRIBUTION BEFORE CRYOMODULE



## PHASE SPACE DISTRIBUTION AFTER CRYOMODULE

