

Horizontal Splitter Design for FFA@CEBAF: Focus on Geometry JLAB-TN-23-069

Ryan Bodenstein

October 2023

1 Document Description

This document will describe the horizontally separating beamlines required for the FFA@CEBAF energy upgrade (henceforth referred to as “splitters”), their purpose, and the process for designing the geometric layout. It also acts as a work survey document, as each step in the design process set new boundary conditions on the path forward. The document will also describe some of the progress on multipass simulations, but much of that discussion will be left for future documents, along with optics design and layout.

2 Background

For the FFA@CEBAF energy upgrade, fixed-field (permanent) magnets will be used to recirculate up to 6 passes in the same beam pipe on the East and West Arcs. The work in this document assumes there is a single FFA arc on each of the East and West Arcs, each containing 6 passes. This is the baseline as of the time of this writing. There will be four electromagnetic (EM) passes, and 6 FFA passes. These arcs replace the current ARC9 and ARCA in CEBAF.

Figure 1 shows a simplified block diagram to aid in visualizing the overall upgrade. A new, 650 *MeV* injector is placed into the current LERF vault, and the 650 *MeV* electron beam will transit along the ceiling (diagonally across from the LINAC) in the South LINAC tunnel, around the West Arc, and be injected into the North LINAC. From here, the first four passes are similar the current CEBAF, passing through each LINAC and the EM arcs on each side. After four passes, the rest of the passes will bypass the EM arcs, and instead pass through a Splitter, FFA Arc, and Transition lattice on each side of the machine.

In this document, the passes are described in multiple ways, which will be explained here. Sometimes, the CEBAF machine naming system is used, where the North LINAC (NL) and East Arcs have odd-numbered passes, and the South LINAC (SL) and West Arcs have even-numbered passes. In this system,

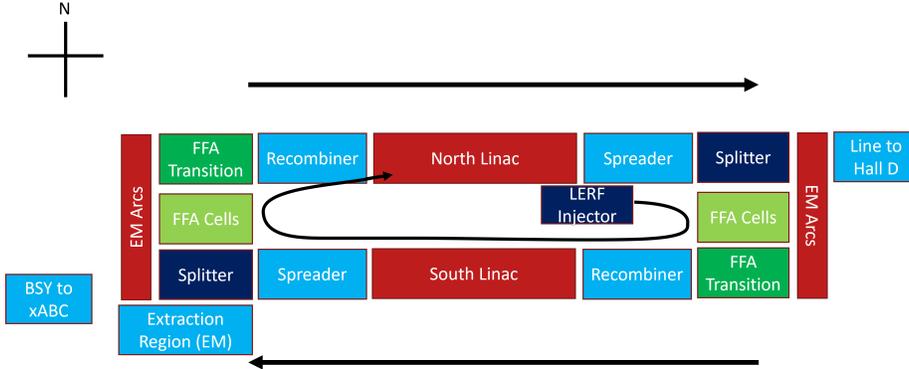


Figure 1: Simplified block diagram (created by Kirsten Deitrick) showing the current planned layout of the upgraded CEBAF.

for example, the FFA passes in the Northeast corner will be Passes 9, 11, 13, 15, 17, and 19, while the Southwest corner has Passes 10, 12, 14, 16, 18, and 20. At other times, for the sake of brevity or ease when speaking specifically of the FFA passes, they will be named according to the FFA-pass number in that area. For example, when discussing the NE Splitter, the lowest energy pass may be called Pass 1, and the highest-energy pass would be called Pass 6. In general, context makes this clear.

The Halbach magnets to be used in the FFA arcs will have Panofsky-style correctors, but otherwise will not be adjustable or tunable. In order to transport multiple passes, strict matching into the FFA arcs is required. Furthermore, since there will not be doglegs for the FFA passes, the Time of Flight (ToF) will need to be adjusted externally.

Based upon the experiences at CBETA, horizontal splitters are envisioned for use. These splitters must control the optics parameters entering the FFA arc, as well as be adjustable for ToF and R_{56} (sometimes referred to as M_{56}).

The splitters must independently control the optics for each pass in the FFA arc. This means that for six passes, six separate lines are required. This requires separating the passes from a co-linear transport in the LINAC first. The FFA passes exit the LINAC, pass through the vertical spreader (**NOT splitter!**) where they are separated vertically, and then recombined co-linearly. They then enter the splitter, where they are separated horizontally. Once fully separated and no longer sharing magnets, each line must match the Twiss parameters (α, β) , dispersion and dispersion prime (η, η') , horizontal position (x) , time of flight, and R_{56} . At minimum, this requires seven quadrupoles, though it is highly recommended that eight or more are used. For this work, we have assumed a minimum of eight independent (not shared) quadrupoles are required per pass/beamline. This will change if the beams do not enter the splitter with non-zero vertical dispersion, however, as we will also need to correct for this. Currently, we assume the splitters are at LINAC height, and the spreader design

will nominally leave the beam entering the splitters with no vertical dispersion or dispersion prime. Of course, it may be important plan for the ability to correct errors in this nominal design as well, which would require more quadrupoles.

3 Geometrical Constraints

If this were a green-field design, we would be able to easily accommodate six independent beamlines with chicanes that could control all of the above-listed parameters. The design could be tidy and elegant, easily controlling the beam parameters throughout each line, and cleanly matching them into the FFA arcs.

However, since this is an upgrade, we must remain within the pre-existing bounds of our tunnel, and accommodate health and safety requirements as well. This means no (or very minimal) alterations to tunnel walls, and no infringement upon the personnel access requirements. Furthermore, equipment access must be maintained as much as possible, especially access for large equipment.

The latter point provides our first major limitation: limiting the splitters to the upstream side of each FFA arc. Given the equipment access points located only in the Northwest (NW) and Southeast (SE) corners, and the transverse space requirements of the splitters, we cannot place two splitters on each FFA arc. The use of ramps over the splitter lines was discussed, and may be possible, but doing so will have a very large impact operationally and logistically. Since the splitters are at LINAC height, and the magnets are currently envisioned as approximately 0.25 m in radius (or half-width), the clearance height of the ramp would be quite high. This would necessitate very long ramps to reduce the incoming angle, and may also put the heads of personnel into the Oxygen Deficiency Hazard (ODH) zone near the ceiling of the tunnel. Given these concerns, ramp(s) should be considered as a last resort, if reasonable solutions cannot be found otherwise.

3.1 Transverse Space Restrictions

Chase Dubbe was able to provide measurements for how much space we have. For the conceptual design, the following assumptions were made:

- The measurements provided remain relatively constant through the whole region under investigation. This means that we assume a constant distance between the tunnel walls throughout the spreader/extraction region. In reality, this may change, and future iterations will need to address this more precisely.
- The measurements are the same on the NE and SW corners of the machine. This may not be exactly true, but for a conceptual design, it is accurate enough. Future iterations will require full wall/beamline measurements for the both regions in full.

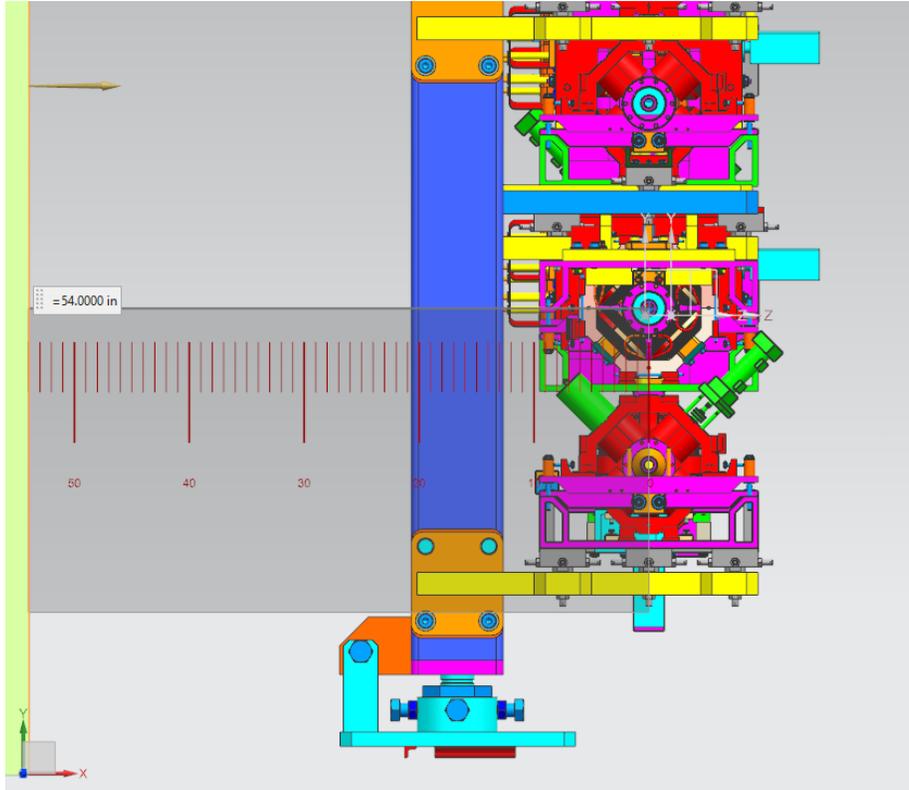


Figure 2: Measurement (inches shown) from the near wall to the beamline center. In meters, this is 1.3716 *m*.

- There are no items/pipes/cables/etc... on the wall in these areas. In reality, this is likely not the case, and these items will likely need to be relocated.
- We are able to adequately run cables/cooling/etc... in this area to power and cool the magnets.

Figure 2 shows the distance from the near wall to the center of the beamline in the SW spreader region. This translates to 1.3716 *m*. Figure 3 shows the distance from the beamline center to the far, aisle-side wall in the same region. This translates to 2.6845 *m*. For completeness, Figure 4 shows the distance from the floor to the current Pass 5 beamline center, where the splitter will be. This translates to 0.68581 *m*.

An additional constraint on transverse space comes from the necessary personnel clearance requirement. According to Harry Fanning, this requirement is 44 inches (1.118 *m*). This must be subtracted from the aisle-side measurement, leaving a total of 1.5665 *m* of allowable space from the beamline center to the

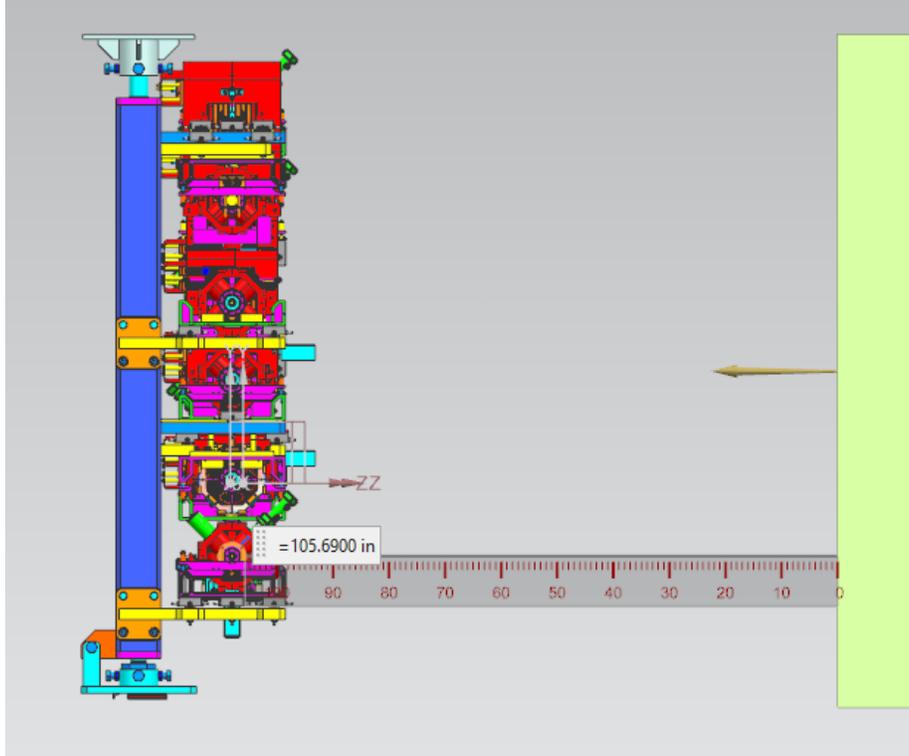


Figure 3: Measurement (inches shown) from the beamline center to the aisle-side wall. In meters, this is 2.6845 m .

required clearance limit. Adding the total allowable transverse space, the splitters can fill a **total of 2.939 m** horizontally, from the near wall to the personnel clearance limitation. All Splitter beamlines must exist within this space. It is important to note that there is more space on the aisle-side from the beamline center.

Finally, in the vertical dimension, having the splitter at LINAC height allows for the possible extraction of the beam downward toward the floor. As the splitters are the only place where the FFA passes are significantly separated, this may be the only feasible manner in which to extract the beams.

3.2 Longitudinal Space Restrictions

Assuming the splitter starts immediately at the exit of the spreader for both arcs, and that we will not be using RF extraction in the same manner in which it is currently employed for the FFA passes (necessitating the long drift space

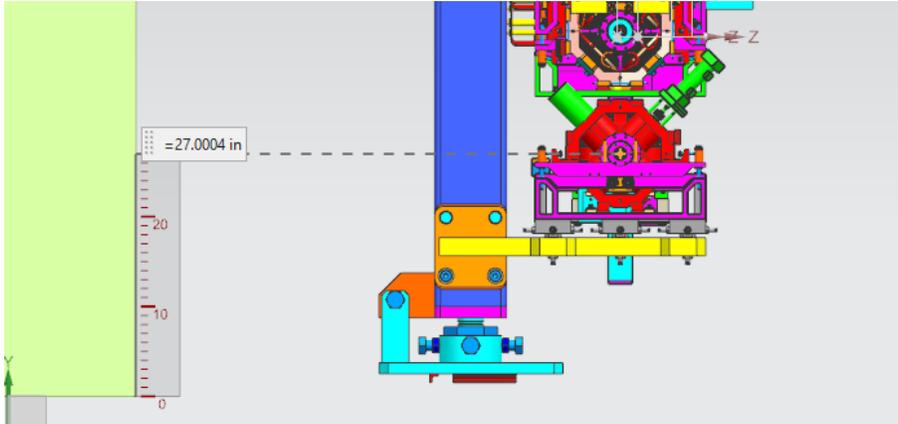


Figure 4: Measurement (inches shown) from the floor to the beamline center of the current Pass 5 beamline. In meters, this is 0.68581 m .

in the extraction region), we can bound the longitudinal extent of the Splitter to the end of the Extraction region (see Figures 5 and 6). This will keep the splitter in a straight section of beamline (allowing us to be more free with bend directions), and allow us to interleave the magnets of the different passes together, conserving transverse space at the expense of longitudinal space.

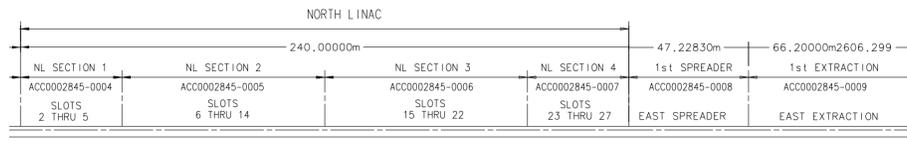


Figure 5: Engineering songsheet for the North LINAC, NE Spreader, and NE Extraction regions.

As currently designed, the NE spreader is 16.357 m in length, and the SW spreader is 19.413 m . From the CEBAF songsheets, the spreader and extraction sections are 47.2283 m and 66.2 m long in the NE, respectively. In the SW corner, they are 45.999 m and 66.2 m , respectively. Adding the lengths of the spreader plus the extraction sections together, and subtracting the length of the currently-designed spreader for the FFA passes, this gives a total available longitudinal space of 97.1 m for the NE corner, and 92.8 m for the SW corner.

In an effort to keep the geometric layout the same for both the NE and SW corners, the conceptual design presented in this note will assume no magnets will extend beyond 92 m in the z direction. Drift space after the 92 m may exist, as it can easily be cut when the splitter section is matched into the FFA arcs.

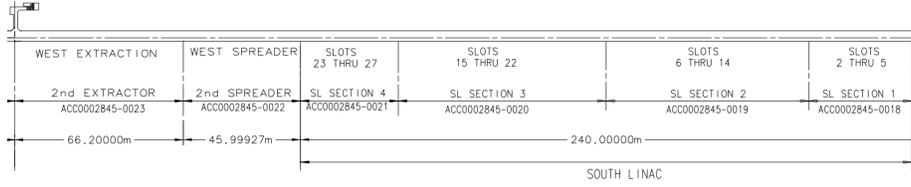


Figure 6: Engineering songsheet for the South LINAC, SW Spreader, and SW Extraction regions.

3.3 Summary of Physical Constraints

Table 1: Physical Constraints - Limits on Splitter Geometric Extent

Name	Plane	Value	Units
Wall to Beamline Center	Horizontal	1.3716	<i>m</i>
Beamline Center to Personnel Clearance Limit	Horizontal	1.5665	<i>m</i>
Total Available Transverse Space	Horizontal	2.939	<i>m</i>
Beamline Center Height (LINAC Height)	Vertical	$y = 100$	<i>m</i>
Total Length in Z (to End of Final Magnet)	Longitudinal	$z = 92$	<i>m</i>

For ease of reference, Table 1 can be used to easily check the physical constraints being placed on the splitter design geometry. Please note, in the machine coordinate system, LINAC height is $y = 100$ *m*. According to Figure 4, this is 0.68581 *m* off of the floor of the tunnel.

The transverse constraints are also shown graphically in Figure 7. The beamline center is labeled with the black line. The walkway is in dark grey, and the walls are labeled in black. The distances are shown from the beamline center to the limitation in each direction.

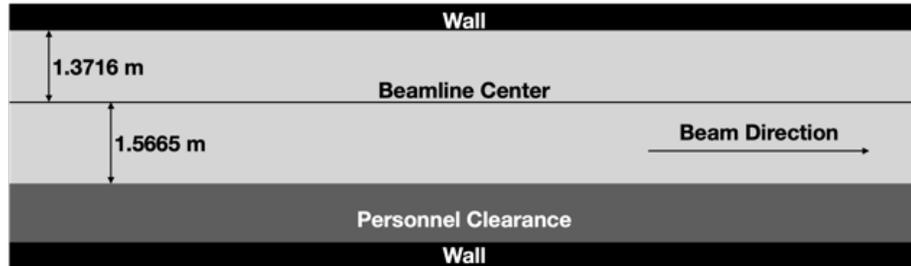


Figure 7: Transverse constraints.

4 Design Rationale

Before continuing, it is important to understand the guiding rationale used in the creation of this conceptual design. (Pardon the switch to first person.)

Firstly, I wanted to place as many realistic design restrictions and constraints as possible on the design. I wanted to be (mostly) pessimistic, with the exception of places I could not afford to be. I wanted to only use electromagnets for the first design, and thick ones at that. I opted to use dipoles that are half a meter in both transverse dimensions, having seen similarly-sized dipoles used elsewhere. The thought is that, if the layout can work with thicker magnets, it may then be possible to re-design them with a smaller transverse footprint (either as pure EM magnets, multi-function magnets, and/or permanent/hybrid magnets) and save space.

I also wanted to focus on **flexibility, and operational robustness and simplicity**. I want to make sure that we can adjust the ToF as the machine “breathes” throughout the year (as the temperature varies and the overall length of the machine changes). I wanted to make sure that we could still deliver beam if we drop a C100. I wanted operations to be able to independently control each pass in the splitter as much as possible (and keeping this control as simple as possible), to be sure that it can be well-matched into the FFA arcs. I want to make sure plenty of space is left for diagnostics, pumps, and other auxiliaries. The more constraints that one can place early on, the easier it is to iterate with fewer or eased constraints later.

In short, “measure twice, cut once.”

Summarizing the guiding design rationale:

- Start with as many realistic constraints as possible.
- Remain a realist/pessimist: only make assumptions and technological leaps when necessary.
- Keep in mind real, designed and/or functioning magnets. We might be able to do better, but we shouldn’t start there.
- It’s easier to remove constraints later than add them.
- **KEEP IT SIMPLE!** Simple magnets, simple designs, **simple(r) operations!**
- Maximize flexibility (optics, time of flight correction, etc...) within limits set by constraints.
- Remember, this is going in a real machine. The machine is imperfect, and the splitters will need to compensate for most of these imperfections for the FFA passes.

4.1 CBETA Splitter

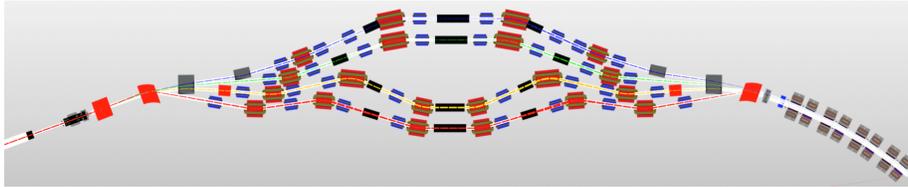


Figure 8: CBETA Splitter

Only one splitter of this sort has previously existed: the one at CBETA (see Figure 8). CBETA has a four-line splitter (for up to 8 total passes in ERL mode - 4 up and 4 down). The energies were MeV -scale, and it was geometrically limited in a manner similar to our situation. During the design process, it became clear to the designers (Chris Mayes and Scott Berg) that the geometric constraints are the main limitation and difficult aspect to overcome. As stated previously, with a green-field design, one has “space to play” and can make any design that works. This is not the case with a highly-constrained system. The first piece of advice given was to “fit the pieces in the box.” Essentially, make sure you can fit all of the necessary magnets and components into the allotted space before attempting to focus on optics design.

The recommended order of design is as follows:

1. Lay out all of the magnets that one expects to need in rough lines to make sure that they fit. This requires realistically-sized magnets to be sure they do not collide with other magnets in the same or adjacent beamlines.
2. Adjust for ToF. Attempt to correct for the time of flight by adjusting the chicanes and adjusting the path length as needed while paying close attention for magnet collisions.
3. Place quadrupoles and start adjusting for R_{56} and the Twiss parameters.
4. Iterate ToF and optics corrections until a solution is reached. This may require re-starting with alternative geometries.

Most of the work in this document focus on the first two points from this list. Future notes will describe other aspects. To start, realistic dipoles, including dimensions, strengths, and shapes are required.

4.2 Dipoles

To base this design in reality, efforts were made to use realistic dipoles from the start. After some other FFA-related magnet designs by Jay Benesch [1, 2, 3, 4, 5, 6, 7, 8], he was asked to check if a $0.5 \times 0.5 \text{ m}$, 3 m -long dipole could feasibly deliver the necessary strength needed to control beams of over 10 GeV in the

Splitters. I wanted to know the roughly-estimated limits on magnet strength for such a dipole so that I was sure not to exceed this value in the design process.

Jay made a preliminary engineering design a very similar, 3.8 m dipole [9]. This work placed a conservative rough limit of 1.825 T on the magnet strength given cooling concerns. Subsequent discussions indicated that a rough limit on a 3 m dipole of those transverse dimensions could perhaps be pushed a little higher, but more detailed studies would be needed in order to be sure exactly where that limit is. A 3 m dipole would be easier to cool than a 3.8 m dipole, and so we could push it a bit higher [10].

For the conceptual optics design, I had to make a few assumptions. Jay’s design is for an H-steel magnet, but I assumed that a similar strength could be achieved with both C-dipoles and septa as well. Furthermore, for the first conceptual design, the orbits may pass out of the good-field regions of the magnets (sometimes very close to the edge, for example). Once this conceptual design is finalized, further iterations will be required to design magnets that meet the specific needs of our design, and to check that the optics and the magnet designs are compatible. For the first attempt, magnet size and strength were used. It is likely that, once the specific magnets are designed, changes will be necessary to the overall optics layout and design as well, to better match reality.

4.2.1 Extraction Dipoles

Given that the splitters are the only location where the FFA passes are significantly separated, it is likely that they will need to be extracted to the halls from within the splitters. One simple way to do this is using C-dipoles and/or septa that vertically bend the beam down at a specific pass.

For this conceptual design, 3 m -long dipoles of 1.5 T are used. They are assumed to be 0.3 cm wide in the X-plane, with no specific definition in the Y-plane. Visual markers are also used to designate the downstream “stay clear” area. These markers assume the next element downstream is 0.5 m wide transversely, so smaller items can be placed in the “stay clear” area, but care will need to be taken to be sure there is still beamline clearance.

4.3 Quadrupoles

The first geometric considerations only considered dipoles, and quadrupoles were only considered after the initial geometric layout was nearly complete. I mention them here for completeness, and will describe some of the details. But it must be noted that they were not considered until after the dipoles were “fit into the box.”

While quadrupoles are not used for the overall geometric design, attention to their realistic sizes is still important. The quadrupoles must be strong enough to control beams between approximately 10 and 22 GeV , but not take up much transverse space. Quads of the necessary strength exist, but most are very large, transversely. It would be preferable to have a longer quadrupole with a smaller transverse size in most places. After calculating the necessary

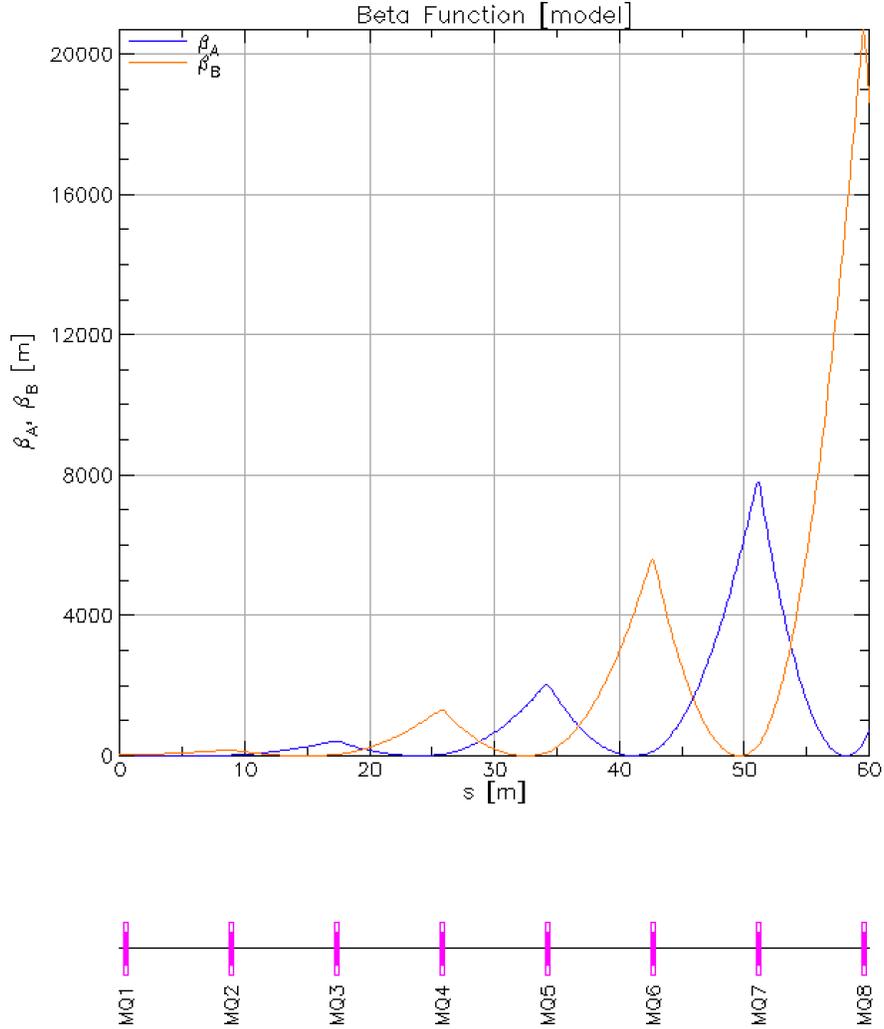


Figure 10: Bmad simulation showing β -functions down a 60 m line.

imately a 4 T field, which is not reasonable. This means the angles will need to be reduced, and the geometry extended into the Extraction region. Setting the first common dipole to 1.5 T gives much smaller angles for each pass, as well as smaller separation between the passes. Looking at the NE splitter, the approximate angles for each pass through the 3 m, 1.5 T common dipole are summarized in Table 2. To maintain the same separation angles in the SW Splitter, the common dipole can be scaled according to energy (with each pass 1.1 GeV higher than in the NE).

Given these small angles, separation of the passes becomes challenging. In-

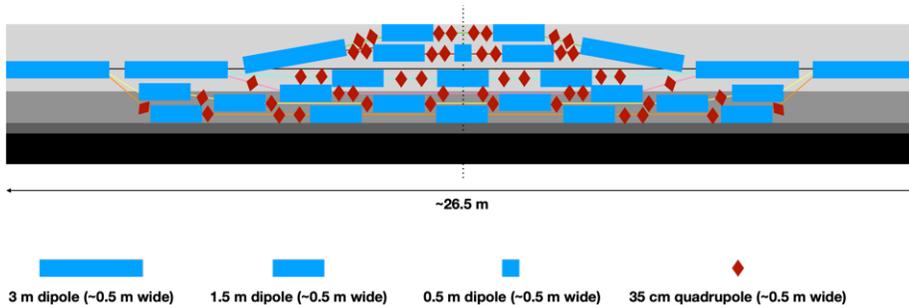


Figure 11: Early splitter layout based upon space and geometry.

Table 2: NE Splitter 3 m , 1.5 T Dipole Separation Angles

Pass	Angle (Degrees)
9	7.3
11	6.1
13	5.3
15	4.5
17	4
19	3.6

creasing the drift space between the magnets increases the space between passes, but also adds to the total length of the Splitter. This also forces the magnets from neighboring passes to be closer together, increasing the need to interleave magnets from adjacent passes.

Seeing the need to use all of the available transverse space, placement of the first dipoles is straightforward. Place the lowest-energy beamline so that the edge of the first non-common dipole is at the personnel clearance limit (since this is more space in this direction than toward the wall). This choice, along with initial common dipole choice, bounds much of the rest of the design. The initial angles for all six FFA passes are set, and the transverse offset is bound by the lowest energy pass. All subsequent passes will depend upon these initial decisions.

Doing this work one beamline at a time would make sense if this were a green field. However, given the tight constraints, it becomes crucial to do all of the lines simultaneously. This avoids situations where the magnets of adjacent beamlines collide. Bmad is capable of performing these multiple-pass simulations simultaneously. Figure 12 shows an early attempt at designing these beamlines separately. The transverse beam sizes were not yet included, and so are not accurate. However, the lengths are accurate, and the simulations use the “patch method” (as opposed to the MAD-style lattice), which ensures the rectangular bends are preserved. Pass 9 and Pass 11 (the two lowest energies for the NE Splitter) were designed in separate files, then overlaid graphically. It

is clear from this picture that the magnets are already colliding with beamlines and each other, even without the proper sizes included.

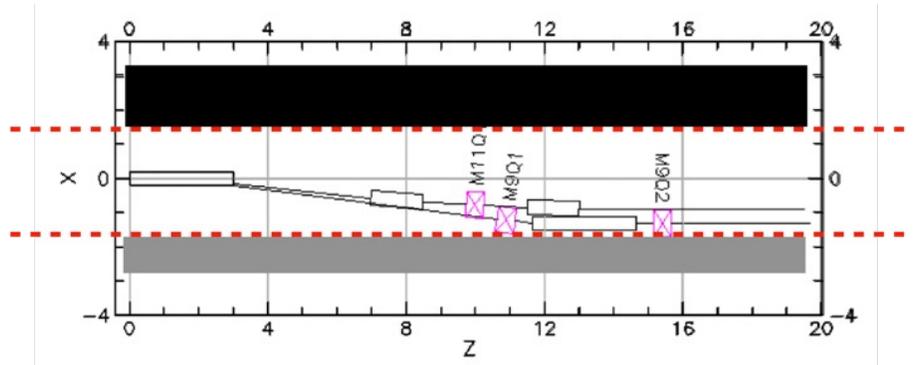


Figure 12: Start of initial Bmad simulations, with the two lowest-energy passes performed independently. The transverse sizes are not properly set yet in this. The black bar is the wall, and the grey bar is the walkway. The red dashed lines are roughly the limits of where the beamlines can extend.

5.1 Alternative Splitter Placement

As an aside, during this time an alternative Splitter was proposed in the center of the FFA arc. This alternative has, at the time of this writing, not been simulated or pursued. But the general idea is shown in Figure 13. The beam will exit the LINAC and go into the Spreader. From here, a small “mini-match” will bring the beam into the first FFA arc, and then into the Splitter. After the Splitter, the beam is rematched into the second FFA arc, then the Transition, Recombiner, and into the other LINAC.

The reason for this is that there is a bit more space between the beamline center and the wall in the arcs. Other benefits include:

1. A mid-arc place for diagnostics, pumping, and mid-course correction.
2. Would not interfere with the current extraction system.
3. Could take advantage of the curved beamline.

However, this would mean that the overall splitter design would need to be on a curved line, and that the direction of the chicanes will need to be taken into account, since this will change the R_{56} .

This idea could be pursued, if it were deemed feasible. But for now, it is only mentioned as an aside.

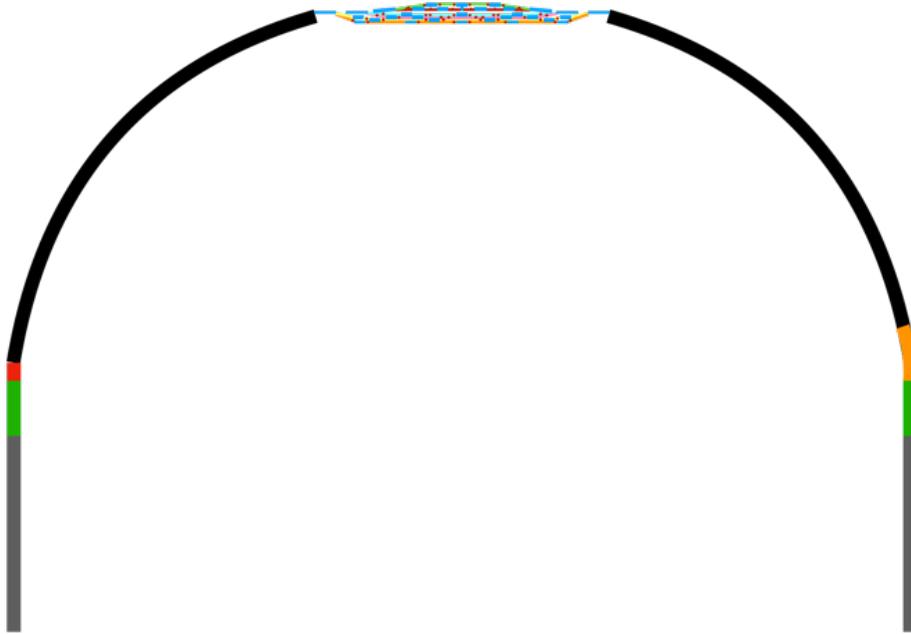


Figure 13: Mid-arc splitter. The beam goes from LINAC (Grey) > Spreader (Green) > “Mini-match” (Red) > FFA Arc 1 (Black) > Splitter > FFA Arc 2 (Black) > Transition (Orange) > Recombiner (Green) > LINAC (Grey).

6 Multipass Simulations

After realizing that the six FFA passes must be simultaneously designed, all efforts went into doing this properly in Bmad. This required using the correct geometrical equations for the incoming and outgoing patches for the dipoles, including common dipoles and the relationships between the passes. A separate tech note will go into the details and specifics of the patches and Bmad coding required to accomplish this. However, a general description would benefit the reader here.

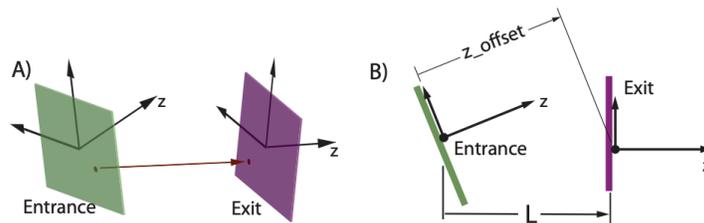


Figure 14: Graphical representation of patches from the Bmad manual [11].

In Bmad [11], patches are used to change coordinate systems between elements (see Figure 14). One way to ensure that dipoles remain the shape that they are supposed to be (i.e. rbends remain rectangular) is to change coordinate systems upon entering and exiting the dipoles in a manner which takes into account the edge angles. An example of the code for these patches is shown in Figure 15. The first dipole is the common dipole for all six passes. The second dipole flattens out the beam at the personnel clearance limit, and the third dipole starts the first chicane.

```

!!!!!!!!!!!!!!!!!!!!!!
!!!!!! Pass 9 !!!!!!!
!!!!!!!!!!!!!!!!!!!!!!

pi_9CB1_xpitch: patch, x_pitch = 0
pi_9CB1_xoff: patch, x_offset = 0

po_9CB1_xpitch: patch, x_pitch = asin(((mCB1[db_field])*(c_light)*(mCB1[L]))/(p9))
po_9CB1_xoff: patch, x_offset = (mCB1[L])*tan(0.5*po_9CB1_xpitch[x_pitch])

pi_9B2_xpitch: patch, x_pitch = -po_9CB1_xpitch[x_pitch]
pi_9B2_xoff: patch, x_offset = po_9CB1_xoff[x_offset]

po_9B2_xpitch: patch, x_pitch = 0.0
po_9B2_xoff: patch, x_offset = 0

pi_9B3_xpitch: patch, x_pitch = -po_9B2_xpitch[x_pitch]
pi_9B3_xoff: patch, x_offset = po_9B2_xoff[x_offset]

po_9B3_xpitch: patch, x_pitch = asin((((m9B3[db_field])*(c_light)*(m9B3[L]))/(p9))-sin(pi_9B3_xpitch[x_pitch]))
po_9B3_xoff: patch, x_offset = -(pi_9B3_xoff[x_offset]+m9B3[L]*tan(0.5*(pi_9B3_xpitch[x_pitch]-po_9B3_xpitch[x_pitch])))

```

Figure 15: Patches used for the first three dipoles in Pass 9. Input patches begin with “pi” and output patches begin with “po”.

In this example, one can see that the beam is coming into the common dipole with no offset and no angle, meaning it is coming into the dipole perpendicular to the edge, and in the center of the magnet. It then exits the magnet at an angle and offset from the center, and enters the second dipole at the opposite sign angle that it exited from the first. To make the second dipole flat against the boundary, the beam enters at an offset, but then exits at the center of the dipole with no angle.

This concept is maintained for all of the dipoles in all of the lines. For patches of shared elements, care must be taken to scale according to energy and placement to make sure that the shared element (such as a common dipole) is placed on the floor identically for each pass.

6.1 Separating the Passes

Separating the passes from each other is the first major challenge of this design. Figure 16 shows the initial separation of all six passes. This work was presented at IPAC23 [12]. Two passes at a time were able to be separated out, then further separated until each is independent. The two lowest energy passes (1 and 2, or in CEBAF convention Pass 9 and Pass 11), at the bottom of the figure (or on the right if looking down the beamline), are able to be independently controlled

by the second dipole in each pass. The four remaining passes are bent through a septum to “level off” the trajectories, and then passes three and four are further separated by two more dipoles. The fourth pass is then separated from the third pass by a septum.

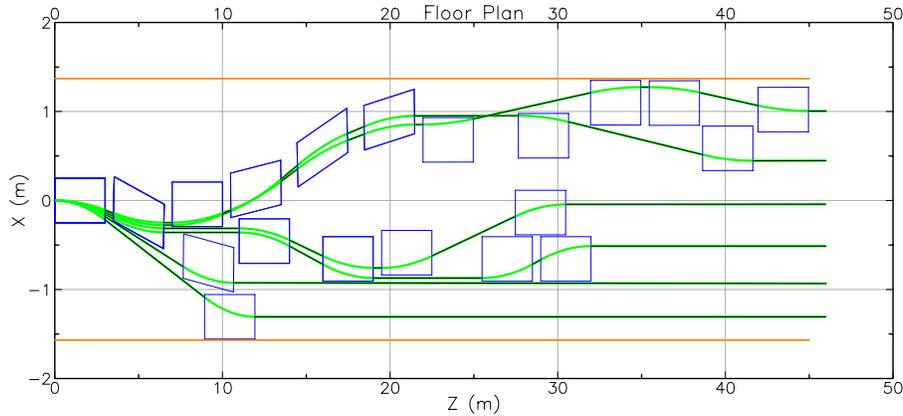


Figure 16: The first six passes, separated out by energy. This was presented at IPAC23.

The two highest energy passes were more challenging. Given the space concerns, and the conservative, pessimistic design rationale, these passes had to be bent to the left (upward in the plot) to clear the other magnets. In bending down and then bending up, the fifth pass (being less rigid than the sixth) crosses over the sixth pass inside of one of the dipoles (the fourth dipole in those shared lines). In order to have the passes in the correct order in the FFA arc (lowest energy to the right/bottom, highest energy to the left/top), the fifth and sixth orbits needed to cross back over. This is accomplished after the $z = 20\text{ m}$ mark, with two septa (or C-dipoles) and the beams crossing over in between (at about 25 m).

It was considered to do this after the right/down bend of passes 3 and 4 by making the common dipole for passes 3 and 4 a septum, and allowing passes 5 and 6 to pass through and then be bent up/left. However, the impact would be the same on the orbits, and the overall design length would increase. Splitting passes 5 and 6 off earlier allows the orbit flips to happen earlier, and reduces the overall length.

Once separated, the patches were used to off-center the dipoles as needed, ensuring that none are colliding. As stated earlier, there are places where the beam may be too close to the edge of the magnet, outside of the good-field region. But this needs further investigation and iteration. For this conceptual design, the assumption is that magnets can be designed to make this work within the overall size of the magnet used in the simulations.

The manner in which the two highest-energy passes are separated places a further constraint on the rest of the design. When adding the chicanes for path

length correction, the direction of the bends for the highest energy passes is in the opposite direction to those of the four lower-energy passes. This can limit some of the flexibility of the chicanes, especially for passes four and five in the splitters, as any movers placed on the chicanes will be limited in the amount that they can be moved. Movers will be discussed a bit more later in this note.

6.2 Adding Chicanes

In a green field design, a simple 4 dipole chicane would be added in each pass to correct for, or aid in correcting, time of flight and R_{56} . These would be stacked on each other or slightly nested into each other, and be adjustable both through the strengths of the magnets and through movers (in a manner similar to that done at CBETA).

However, assuming six passes and half-meter widths on all of the dipoles, simple math shows that stacking won't fit: $6 \text{ passes} \times 0.5 \text{ m} = 3 \text{ m} > 2.939 \text{ m}$ of allowable space. So already, nesting will be required to avoid magnet collisions. Even with smaller magnets in the transverse direction, it is likely beneficial to nest these chicanes, since the fringe fields may also end up being a concern.

The simplest method of stacking the chicanes is like a set of Matryoshka dolls: the lowest energy line will have the shortest drift between the dipoles, and each successive pass will have longer drifts between the dipoles to go around the lower energy pass beamlines. This can be done for the first four FFA passes in this splitter, but the two highest-energy passes will need to be separate and bend the other way, due to the manner in which the orbits cross and uncross. This will limit how far passes 4 and 5 can extend horizontally, as the extension of the chicane in one will necessitate the shortening in the other.

In order to increase the ability to adjust for changes in path length throughout the year, each chicane needs to be adjustable in a manner similar to that of CEBAF's doglegs. Operationally, the doglegs are not used often, but they are needed at times to adjust on the order of centimeters of path length throughout the year. They are capable of far more adjustment.

CEBAF's fundamental frequency is 1497 MHz , meaning that an RF wavelength is 20.0262163 cm . (Pardon all of the decimal places in the upcoming paragraph - this is copy/paste from a spreadsheet.) This gives a period of $6.68002672010690\text{E-}10 \text{ m}$. One degree of RF wavelength is 0.000556284 m , which translates to $1.85556297780747\text{E-}12 \text{ s}$ when converting to time. Currently, we are capable of correcting to better than a quarter of a degree, or $4.63890744451868\text{E-}13 \text{ s}$. The splitters should also be able to control the time of flight to this level, although there have not been specific requirements written to address this. For now, it is an assumption.

Stacking the chicanes allows all of the beamlines to fit in the allowable transverse space. However, the "height" of each chicane is limited, and therefore its flexibility and ability to adjust path length. However, the physical constraints longitudinally are not as strict, and it may be prudent to make multiple, smaller amplitude chicanes for each pass in order to regain some of this operational flexibility.

Starting with the lowest energy pass, a “mini-chicane” was added prior to the main chicane. The thought was that by adding this, a few centimeters of path length could be added/subtracted by adjusting the strengths of the dipoles and moving the beam around in a thick beam pipe. Figure 17 shows an early iteration of this process. Looking at the lower left of the figure, one can see a small “mini-chicane” added with 1.5 m dipoles. Downstream in the same beamline, the main chicane for the lowest energy pass is under construction. The patches had not been added yet, but the magnet strengths were adjusted to show the intended extent of the beam. Later iterations improved upon this, and re-ordered some of the magnets from the adjacent passes as well. One iteration after this also added a second mini-chicane after the main chicane for the lowest energy pass. However, this was later removed, as it was unnecessary, and the space is needed for other elements.

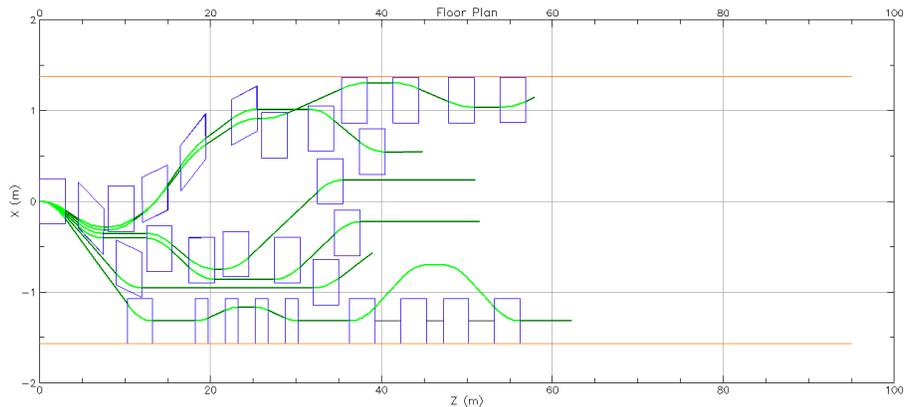


Figure 17: An early step in the construction of the chicanes. Please note: not all of the patches were updated at this point. This floor plan is an output from Bmad. The orange lines on the plot represent the wall (top orange line) and the personnel clearance limitation (bottom orange line). The sizes of the components are accurate, but the scales for each axis are different.

For all of the other passes, however, there is no space upstream to add mini-chicanes. Furthermore, as the energy increases, larger magnets are required to achieve the desired orbit changes. Adding chicanes downstream of the main chicane for each pass is the only option. Doing this will add length to the overall design, as well as extra magnets. But the added control and operational flexibility is a worthwhile trade-off. For the four lowest-energy beamlines, this idea was pursued. For the two highest-energy passes, only the main, four-bend chicanes were added, partially due to space concerns, partially due to the very high energies, and partially due to them bending in the other direction.

To add the downstream, smaller chicanes, the beam will no longer go back to the original “base” of the chicane main chicane of that beamline. After discussions with Scott Berg, care was taken to make sure that movers (more on

this later) could be added to one or both of the chicanes in each beamline. This means that the pair of dipoles at the “peak” of each chicane will need the ability to move together in the same direction with the same magnitude (as well as the beamline in between them). This rules out ideas such as two-step chicanes (as in our Spreaders).

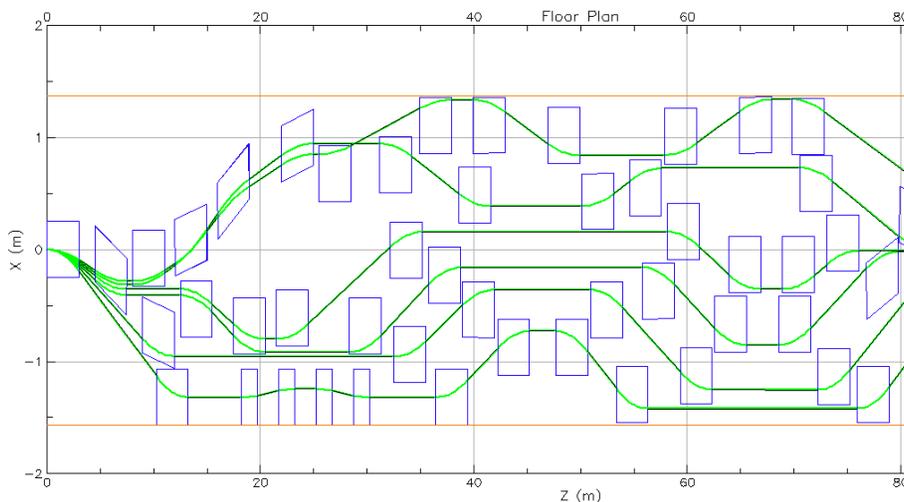


Figure 18: Floor plan output from Bmad showing the series of chicanes for each Splitter pass. Please note, some patch adjustments are still required to even out magnet spacing between passes.

Each main chicane was designed to bend the beam back toward the base of each chicane, but not to the same transverse offset. The first pass beam is brought back to a point further right/down than the upstream side of the chicane in order to allow more space for the other passes. The second pass also returns to a point further to the left for the same reasons. Passes three and four stack onto each other, but next into the chicane of the second pass. Passes five and six bend back to the left/up, but only to complete their main chicanes. For these last two passes, the added flexibility can be found by adjusting these latter magnets, but the flexibility is more limited than the lower energy passes. Figure 18 shows this on the floor plan.

It is important to note that this decision required many compromises and trade-offs, but these are operationally beneficial. The first of these trade-offs is that the beam will no longer be closed using common dipoles. This will be explained further later, but in the end, this leads to more flexibility and more robust operations. This also requires more magnets and more power supplies for the magnets. However, until the operational needs are studied in greater detail, the above-stated rationale for maximizing flexibility guided this decision. Furthermore, this decision allows for quadrupoles to be placed more easily, and the optics to be better controlled for each pass.

6.2.1 Movers

As a quick aside, the addition of movers may or may not be required, but the option should be left open. For CBETA, using the movers to change the path length was beneficial, and added a level of flexibility and control to the path length. While much of the path length could be corrected without their use in CBETA, having this option available during this early design stage is important. It is easier to cut the option out later than to add it.

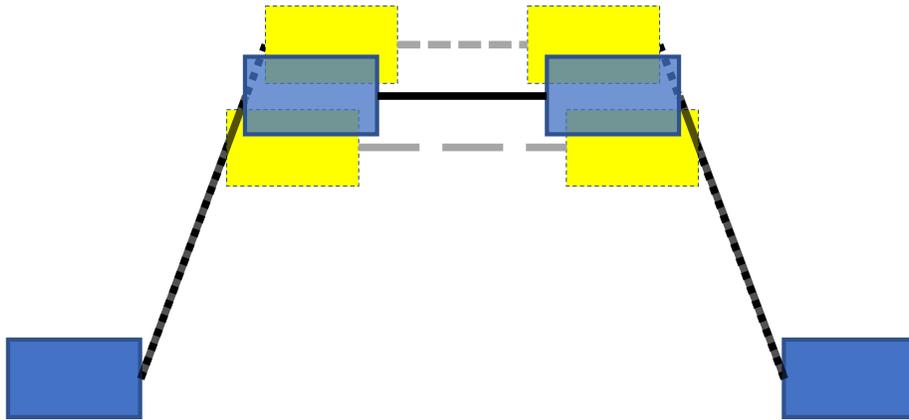


Figure 19: Simplified cartoon of the mover system for the Splitters. Blue are dipoles, and yellow are possible positions of the dipoles after moving. The grey lines show the orbit for the different paths.

The movers would be mounted on the dipoles at the peak of each chicane (as well as the beamline and girder in between). This way, if more path length is needed, the height of the chicane can be increased by essentially increasing the length of the drifts on each side of the chicane. Vice-versa if less path length is required. In the design above, movers can be added to either or both chicanes in each beam line. Operationally, limits will need to be placed on them so that element collisions do not occur. Figure 19 shows a very simplified diagram to express this graphically. The dipoles at the peak are able to move up or down to the positions indicated in yellow. The black lines are the nominal orbit without the movers moving. The grey lines are where the orbit will be after using the movers. In between the two peak dipoles, the beamline itself will need to move with the dipoles.

Mover systems can be placed appropriately on either or both of the chicanes in each line, assuming the systems do not encroach upon the transverse space too much. A system like this exists at CBETA. The engineering aspects of this system are beyond my expertise, but knowledge that they exist is adequate for their possible inclusion in our Splitters.

6.3 Adding Extraction

While no technical work has yet occurred for the FFA@CEBAF extraction system, general concepts and ideas have been discussed. During these discussions, it became clear that the only place where the FFA passes are significantly separated from each other are in the splitters themselves, and this may be the best (or only) place where extraction can occur.

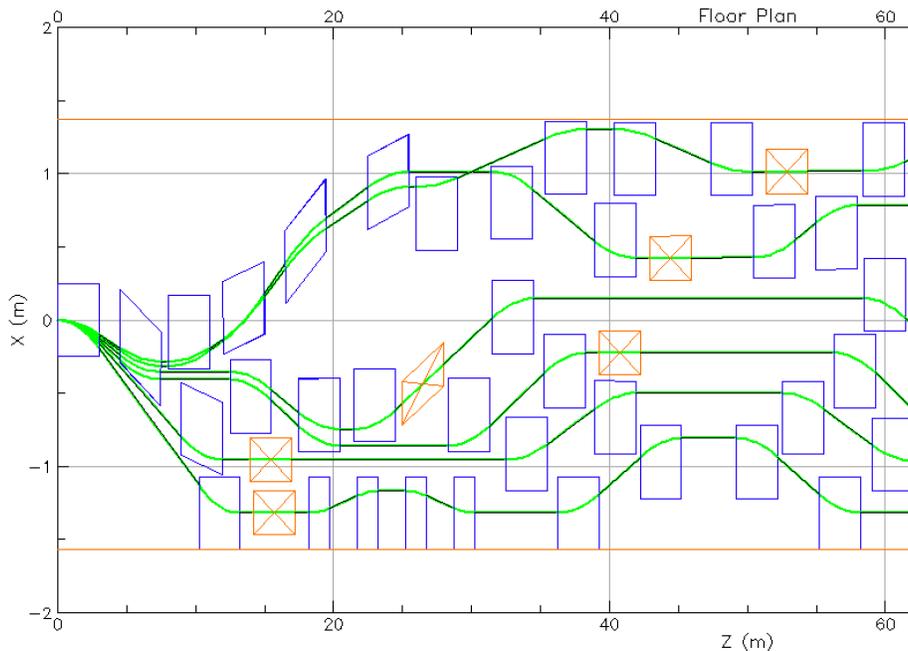


Figure 20: Initial placement of the extraction dipoles, indicated in orange.

In order to accommodate this extraction, space was reserved in each Splitter line for beam extraction. The assumption is that a 3 m , 1.5 T C-dipole or septum could be used to bend a given pass downward toward the floor, and then transported to the Beam Switchyard (BSY) region. Essentially, these magnets would be unpowered unless the beam for that pass is to be sent to the halls. In that case, the magnet would be turned on, and the same energy beam would be sent to all three halls, A, B, and C. Hall D would need a separate extraction system, or may be excluded from operation while Halls A, B, and C are in operation. The details of how the beams will be extracted need to be investigated, and this may necessitate significant changes in the placement and type of these dipoles.

Figure 20 shows the initial placement of the extraction dipoles. These are estimated to be 30 cm wide in the horizontal plane, and 3 m long. Later

iterations moved these locations a bit, and used a marker to indicate a “stay clear” region for any element with a half-height of 25 *cm*. In Bmad, these magnets and markers are placed using the superpose functionality, which allows elements to be superimposed onto other elements, such as magnets onto drifts.

Placing these extraction dipoles on the upstream side sacrifices space for quadrupoles in some lines. However, this leaves more space on the downstream side of the Splitters to concentrate more quadrupoles, and allow better matching into the FFA arc.

6.4 Closing the Orbits

In the end of the Splitters, the orbits must be brought back together into the FFA arcs. It is important to reiterate (as this is frequently overlooked or forgotten) that the beams must NOT be recombined co-linearly. Although they enter co-linearly (nominally), there is significant separation required of the beams as they enter the FFA arc (again, forgive the extra decimal places). Table 3 summarizes the requirements for the beam centroid positions as they enter the East FFA Arc. The full distance between the entrance points of the lowest and highest energy beams is approximately 3.5 *cm*. While initial error studies indicate some level of flexibility on this requirement, a final requirement has not yet been written. Good design practice aims to meet the nominal requirements as exactly as possible, so these numbers are being used for all of this work.

Table 3: Required X-Positions Entering East FFA Arc

FFA Pass Number	X-Position Into East FFA Arc (m)
1	-2.90804649998540E-02
2	-2.50580958006650E-02
3	-1.92036478877850E-02
4	-1.18325916016940E-02
5	-3.19002105495230E-03
6	6.52988940656190E-03

As mentioned previously, opting for multiple chicanes per line limits the ability to recombine the orbits using common dipoles without extending the splitter into the arc region. However, since each pass has an independent horizontal recombination coordinate, and controlling multiple passes through common dipoles can be incredibly complicated operationally, this option leads to more control and better matching into the FFA Arc.

Furthermore, closing the orbits in separate lines without the use of common dipoles allows a final match section to be added to most of the lines. A series of independent quadrupoles in the final stretch of beamline can be used to more accurately control the beam into the FFA arcs.

Figure 21 shows the first simulation where the orbits were closed in Bmad, with all elements (excepting the final drift) within the geometric constraints

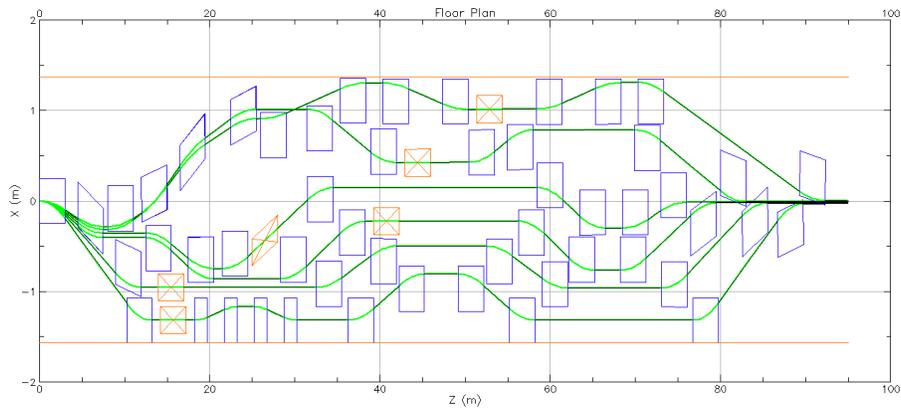


Figure 21: First closure of the orbits in Bmad.

placed upon the design. Each line brings the orbit to the correct horizontal position. It is important to note that the final dipoles/septa in each line will likely need to be specially designed, and the corners which overlap with the common beamline too much will need to be chamfered. However, angling the magnets as they are shown allows for interleaving of the magnets, as well as giving space for small correctors to be installed to account for the crosstalk that will occur at each magnet with the common beams. As this design progresses past the conceptual stage, these details will become more important, and may indicate necessary changes.

6.5 Adding Quadrupoles to Check Clearances

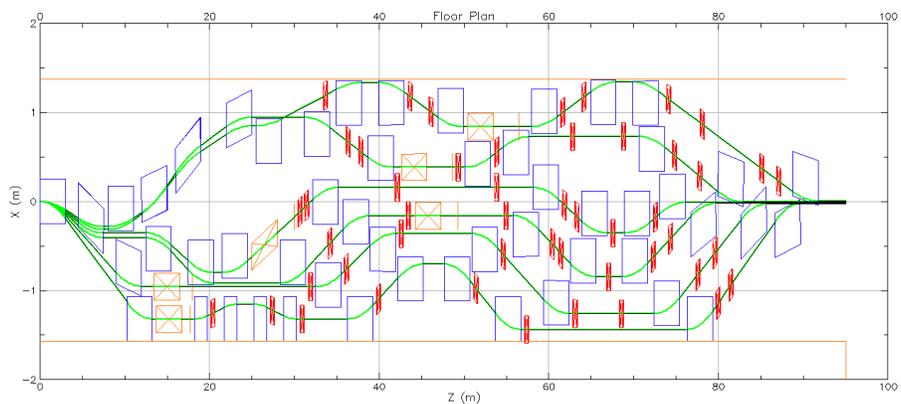


Figure 22: Placing quadrupoles to check for clearances.

Recalling the minimum requirement of eight *independent* quadrupoles per line, QR quadrupoles were placed on each line of the splitter to check for clearances and possible fit. Please note, there were not placed for optimal optics control, but rather to see where and how they fit, geometrically.

Figure 22 shows an example of how eight or more quads can fit on each beamline. Some of the adjacent beamlines are so close together that the quadrupoles must be interleaved. It is important to note where these tight spaces are, as they will dictate where quadrupoles can be placed, and thus the overall control of the optics. In the Bmad code, I opted to place the quads by using the superposition functionality, where one can superimpose the element onto another element, and use the placement position as a variable. In this case, this eases the movement of quadrupoles along the drifts, as one does not need to manually change the drift lengths to accommodate the move.

While some lines may not be capable of accepting quadrupoles in the standard, textbook locations commonly associated with four-bend chicanes, there is quite a bit of space in some areas. The placement of the quadrupoles in Figure 22 are only to demonstrate possible locations, and how tight some spaces may be. The number, placement, and strengths of all of the quads will need to be determined by the optics needs, with the physical constraints in mind as boundary conditions.

6.6 Adjusting R_{56} Manually

Prior to using optimizers to correct for the time of flight, a rough, by-hand adjustment was performed based upon a quadrupole-free Splitter. For an order-of-magnitude approximation, in the absence of quadrupoles, the $R_{56} \approx \Delta s$. **Again, once quadrupoles are part of the equation, this is no longer true.**

The adjustments were made by hand so that a better understanding of the interplay between elements could be gained. This helped to better understand how to write an optimizer for later work.

Using this approximation, the path lengths of all six passes in the Splitter were adjusted to get close to the R_{56} needed to compensate for that in the FFA Arc, the Spreader, and the Recombiner. This was done roughly, so that the beamlines will be closer to a reasonable solution when using the optimizer to correct for time of flight.

Table 4 shows the values of R_{56} for each section of the East FFA Arc in meters. In this table, the CEBAF machine pass naming convention is used. Since the Recombiner is an exact mirror image of the Spreader, the R_{56} value is doubled. The Transition lattice is not complete, and so for now it is ignored. The Total Arc column is the full measurement of all FFA Arc sections *except for the Splitter section*. The Splitter column shows the values of the R_{56} for the initial geometric Splitter layout. These are the values that were corrected by hand.

Table 5 shows the adjustments in the R_{56} required to cancel it out for the full arc (minus the Transition). The “Exact” column shows the adjustments

Table 4: East FFA Arc Section R_{56} Values (m)

Pass	FFA Arc	Spreader $\times 2$	Transition	Total Arc	Splitter
9	-0.005344593	0.07623084	0	0.070886247	0.389353268
11	-0.096832448	0.052138893	0	-0.044693555	0.228843571
13	-0.165990039	0.037899254	0	-0.128090784	0.211129281
15	-0.220688224	0.028788014	0	-0.19190021	0.16462742
17	-0.264936009	0.022607879	0	-0.242328131	0.263299145
19	-0.299034716	0.018223956	0	-0.28081076	0.236980915

required to keep all of the FFA passes in the same RF wavelength. All of the other columns are the adjustments required if they will be an integer number of wavelengths out. At times, the adjustments (drift length and dipole strengths) “fought” each other. This was especially true for the fourth (Pass 15) and fifth (Pass 17) FFA passes, as the chicanes bent in the opposite directions. So while a small adjustment may be required to compensate, that would cause magnet collisions, so the next wavelength out, while a larger adjustment, might be necessary.

Table 5: R_{56} Adjustments Required in Splitter (m)

Pass	Exact	2π	4π	6π
9	-0.318467021	-0.118204858	0.082057305	0.282319468
11	-0.184150016	-0.073274962	0.126987201	0.327249364
13	-0.083038496	-0.138957902	0.061304261	0.261566424
15	0.02727279	-0.156265467	0.043996696	0.244258859
17	-0.020971014	-0.305365112	-0.105102949	0.095159214
19	0.043829845	-0.317529512	-0.117267349	0.082994814

By hand, adjustments were made to the path length of each line in the Splitter to get reasonably close to the values listed in Table 5. Passes 9, 11, and 13 (first 3 FFA passes) were adjusted to so that they were two integer wavelengths out, giving errors of $-2.27166\text{E-}05$ m, $9.05552\text{E-}05$ m, and $1.33744\text{E-}06$ m, respectively. Passes 15, 17, and 19 (last three FFA passes) were adjusted to the same wavelength, with errors of 0.00094377790358741 m, -0.00083255737487281 m, and -0.00007791454945300 m, respectively. Again, please pardon the extra digits in these numbers, they are a quick copy/paste for the note.

After these by-hand adjustments were made, the focus was changed to time of flight. Now that the pathlength is reasonably close, and a better understanding of how the geometries of different lines limit each other, writing an optimization for time of flight adjustment is more straightforward.

6.7 Adjusting Time of Flight

As of the time of this writing, the simulations for ToF correction are still ongoing. After the by-hand adjustments mentioned in the previous section, the focus changed from ballpark R_{56} adjustments to real time of flight corrections. For now, these simulations can only be approximations as well, since the full beam transport between the LINACs is not complete. We have beam transit times for the Spreaders, Recombiners, Splitters, and the FFA Arc itself. However, the transition section on the downstream end of the FFA arc is not complete. Its length is still unknown, and so it is not yet included in the numbers used for the ToF correction. It is possible that the length of the transition may encroach on the nominal FFA arc, and we may need to cut or add cells to the FFA arc to compensate. Once this section is complete, we will then have more complete numbers for our calculations.

Table 6: East FFA Arc Section Time Values (s)

Pass	FFA Arc	Spreader \times 2	Transition	Total Arc	Splitter
9	7.878894531243210E-07	1.093683364999980E-07	0	8.972577896243190E-07	3.169229878607050E-07
11	7.879087202062690E-07	1.093519983455030E-07	0	8.972607185517720E-07	3.175382776501070E-07
13	7.879571939810150E-07	1.093423379742990E-07	0	8.972995319553140E-07	3.178892563957920E-07
15	7.880283847613040E-07	1.093361553853070E-07	0	8.973645401466110E-07	3.175249238605530E-07
17	7.881170811162100E-07	1.093319611440570E-07	0	8.974490422602670E-07	3.172118001821740E-07
19	7.882187546861550E-07	1.093289856429910E-07	0	8.975477403291460E-07	3.151564314054270E-07

Table 6 lists the beam transit times through each section of the East FFA Arc for each pass in the same manner described for the R_{56} section. The FFA Arc proper is added to the Spreader and Recombiner (and the Transition Lattice) times to get a Total Arc time, excluding the Splitter. The values shown in the Splitter column are for the Splitter as set from the previous design step (setting the R_{56} by hand). All of the digits are left in here, partially due to laziness, and partially because some of the differences fall several digits out. Not all of these digits are significant.

Next, a target time for each pass needed to be set. As long as all of the passes transit the from the end of one LINAC to the beginning of the other with either the same time, or an integer number of RF wavelengths from the others, then things are ok. One method to do that is to add the current Splitter time to the transit time of the rest of the FFA Arc. This gives the total transit for each pass. Then a target can be chosen: one could choose a specific pass, or some value that they can all be adjusted to. For the first attempt, the average total transit time was set as a target (or $n \times 2\pi$ from the target). The path length of the Splitter is then adjusted to change the time so that all passes reach the next LINAC in the same amount of total time (or an integer number of wavelengths out).

In Bmad, overlays are set up to control different aspects of the machine. For this optimization, overlays were written to adjust dipole strengths and drift

lengths for specific elements in each Splitter line. The optimizer then uses these overlays as variables, with constraints added to avoid element collisions. Each pass is simulated to target the time, the x-coordinate entering the FFA arc, and the z-coordinate at the end of the line. The optimization procedure is still being refined. Currently, simulations are capable of meeting two of the three targets nearly exactly (namely time and either x or z), but not all three adequately. Refinements are underway to improve this.

7 Current Status

As of the time of this writing, there is a reasonable geometric layout of the Splitters. Care has been taken to make the geometry identical for both the Northeast and Southwest Splitters, with the magnets being scalable from the NE corner design. Minor changes are expected to be necessary as the time of flight is adjusted to compensate for the transition lattice. Refinements of the ToF optimization are underway so that all three targets can be met for all six FFA passes while maintaining adequate transverse space to install quadrupoles and other beamline elements.

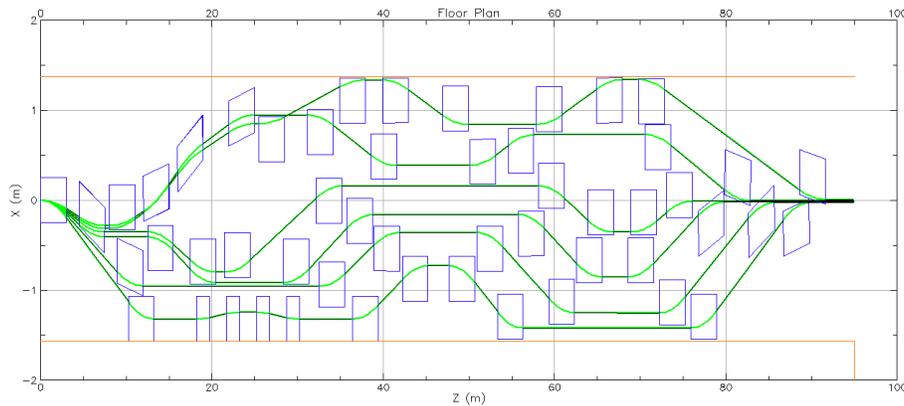


Figure 23: Current Splitter design, without extraction dipoles or quadrupoles turned on.

Once the ToF optimization is complete, quadrupole-based optics calculations and simulations will begin. It is important to note that, with the recent baseline change necessitated by the transition lattice design, there are no accurate input Twiss parameters to input into the Splitters. All optics are currently based on the longstanding baseline with weakly focusing triplet optics. Until the new baseline, based on strongly focusing triplet optics, is propagated throughout the machine (including matching through each arc, the spreaders, and the recombiners), there are no reliable numbers to input for optics studies with the new baseline.

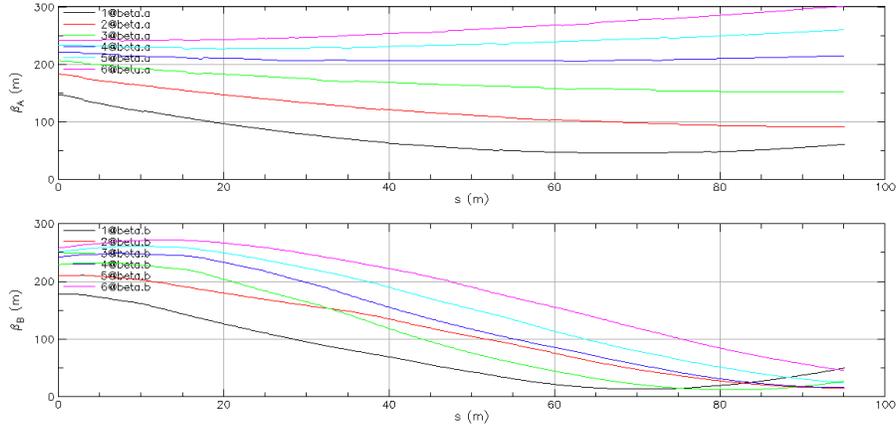


Figure 24: Current β s, without quadrupoles turned on.

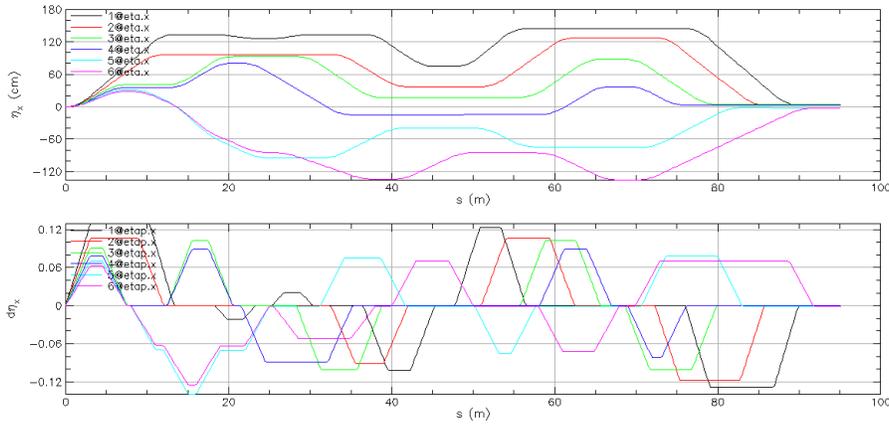


Figure 25: Current dispersion, without quadrupoles turned on.

Figure 24 shows the current β functions in the absence of quadrupoles (rather, they are not powered). A is the x plane and B is the y plane in these plots. The dispersion and η' is shown in Figure 25, again in the absence of powered quadrupoles. It is worth noting that the dispersion nearly closes at the end, despite the orbits not coming back to co-linearity. Closure of the dispersion should be relatively straightforward to accomplish. However, the dispersion throughout the lines is relatively large, reaching over 1 m for the passes which are bent furthest from the original beamline center. Care will need to be taken to control the dispersion, while ensuring that it still closes at the end of the splitters.

Significant efforts will be required to match the β s into the FFA arc, though

with the new strongly focusing LINACs, this may be easier to accomplish. Re-matching the new optics through CEBAF should be a high priority, given the dependencies throughout the machine.

Table 7 shows the parameters required to match into the FFA arcs. Please note, more precise numbers can be acquired from Alex Coxe, who provided the numbers, or on our GitHub repository [13]. Also note, the pass numbers have two different labels. The numbers in parentheses are labelled according to CEBAF’s pass convention, where the North LINAC and East Arcs get odd numbered passes, and the South LINAC and West Arcs get even numbered passes. It will clearly be challenging to match the current incoming Twiss parameters to those required by the FFA arcs. With the reduced β s that the strongly-focusing triplet LINACs provide, this should be easier, although still non-trivial.

Table 7: Entrance Match Parameters for FFA Arcs

	E [GeV]	X [m]	Px	β_x [m]	α_x [m]	β_y [m]	α_y [m]	η [m]	η'
East Arc:									
Pass 1 (9)	10.55	-2.908E-02	1.116E-02	4.157	3.049	6.515	-3.190	0.027	-0.020
Pass 2 (11)	12.75	-2.506E-02	9.765E-03	2.951	1.822	6.477	-3.037	0.046	-0.026
Pass 3 (13)	14.95	-1.920E-02	7.564E-03	2.718	1.539	6.995	-3.206	0.061	-0.031
Pass 4 (15)	17.15	-1.183E-02	4.695E-03	2.602	1.399	8.035	-3.636	0.073	-0.034
Pass 5 (17)	19.35	-3.190E-03	1.264E-03	2.521	1.311	10.132	-4.549	0.084	-0.035
Pass 6 (19)	21.55	6.530E-03	-2.643E-03	2.455	1.247	16.840	-7.524	0.093	-0.035
West Arc:									
Pass 1 (10)	11.65	-2.336E-02	1.008E-02	3.254	2.693	6.444	-3.457	0.028	-0.021
Pass 2 (12)	13.85	-1.955E-02	8.529E-03	2.510	1.739	6.200	-3.201	0.043	-0.026
Pass 3 (14)	16.05	-1.427E-02	6.277E-03	2.360	1.492	6.490	-3.279	-0.030	-0.030
Pass 4 (16)	18.25	-7.763E-03	3.430E-03	2.286	1.369	7.160	-3.571	0.066	-0.032
Pass 5 (18)	20.45	-2.090E-04	7.594E-05	2.234	1.290	8.396	-4.153	0.075	-0.033
Pass 6 (20)	22.65	8.245E-03	-3.716E-03	2.190	1.233	11.045	-5.435	0.083	-0.033

8 Ongoing and Future Work

Currently, the optimizations required for time of flight correction are being refined and improved so that all of the targets can be reliably reached. With the way this is written in Bmad, quadrupoles (and the extraction dipoles) are superimposed onto drifts. This allows them to be moved and changed easily along the drift. However, this complicates the ToF correction, as the drifts on which they are superimposed no longer exist after the lattice is read by the program (they are split automatically and renamed accordingly). So to run the ToF optimization, I disable the superpose functionality, essentially deleting the

elements. Once the ToF correction finishes, the elements are then superimposed again to check for element collisions, etc...

Once the ToF correction is complete (at least temporarily until the transition section is complete), focus will turn to optics control. The main focus will be matching into the FFA arcs, while attention will also be paid to controlling dispersion, \mathcal{H} , and radiation losses within the Splitters. Placement of quadrupoles will be based upon optics control, using the geometric limitations as boundary conditions. Furthermore, care must be taken to leave adequate space for diagnostics, correctors (at least dipole, but perhaps higher-order magnetic correctors as well), pumps, and other auxiliaries.

This will be an iterative process. However, the hopes are that a complete conceptual design will be provided so that Start-2-End simulations can be achieved. Since the overall design of FFA@CEBAF is in constant flux, it is important to be sure that changes are propagated throughout the whole machine each time, and that a working model is agreed upon at each stage.

References

- [1] Jay Benesch - JLAB-TN-22-033, *First attempt at a conductively-cooled superconducting septum magnet for the FFA upgrade*
- [2] Jay Benesch - JLAB-TN-22-037, *Second attempt at a conductively-cooled superconducting septum magnet for the FFA upgrade*
- [3] Jay Benesch - JLAB-TN-22-052, *Third attempt at a conductively-cooled superconducting septum magnet for the FFA upgrade*
- [4] Jay Benesch - JLAB-TN-22-041, *Extended Lambertson for FFA*
- [5] Jay Benesch - JLAB-TN-22-051, *ZA modification concept*
- [6] Jay Benesch - JLAB-TN-22-054, *Water-cooled copper septum for FFA*
- [7] Jay Benesch - JLAB-TN-23-040, *A Simple Modification of DC Current Septa to Reduce Current Density by Half*
- [8] Jay Benesch - JLAB-TN-23-041, *Rectangular common dipoles for FFA*
- [9] Jay Benesch - JLAB-TN-23-016, *Conventional Dipole for FFA Splitters*
- [10] Further discussions and email exchanges helped to clarify the details and assumptions.
- [11] Bmad website: [Bmad \(Clickable Link\)](#)
- [12] R.M. Bodenstein *et al.*, “Designing the spreaders and splitters for the FFA@CEBAF energy upgrade”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 965-968. doi:10.18429/JACoW-IPAC2023-MOPL183 ([Clickable Link](#))
- [13] FFA@CEBAF [GitHub \(Clickable Link\)](#)