

Status update on R30 electrode studies

M. Bruker

12/31/2023

corrected 1/29/2024

Going into the new year, I'll summarize what we learned in the past couple of weeks so everyone can be on the same page again. Don't mind the ugly figures and such. References are also missing.

Contents

1	Strategy	2
2	Poisson models	2
3	Focal length and spherical aberration	4
4	Anode shape	11
5	Astigmatism and kick	13
6	Random ideas	13

1 Strategy

In an ideal world, we would like to conceive the perfect R30 gun that produces the smallest possible emittance, Twiss parameters that make the injector happy, and minimum sensitivity to space charge. However, we only have another couple of weeks to come up with a design because the gun needs to be ready for installation by May, and it is unlikely we can do a beam test at UITF beforehand. On the one hand, this means we have to get it right at the first attempt without much room for iteration. On the other hand, the initial conditions are not unfavorable. The R30-3 gun was rejected mostly because of its very short electron-optical focal length, which led to a large beam diameter and high sensitivity to random laser-spot motion, but it performed according to expectations otherwise. The idea is, thus, to start with the R30-3 design and replace only a handful of parts that are quick to fabricate and do not add excessive risk. It is possible the only modification necessary turns out to be the cathode cone insert. For now, we will try to find a reasonable cost/benefit balance to ensure success. However, should it become clear that significant improvements can be gained by making more invasive changes, we can keep developing those for some point in the future.

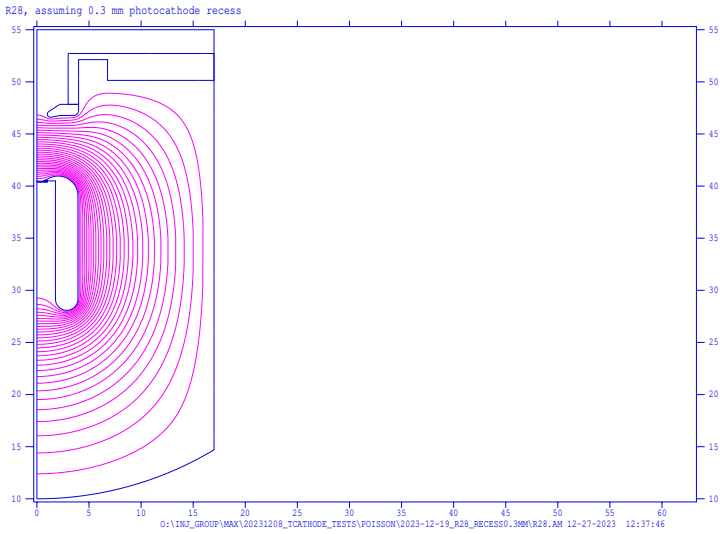
2 Poisson models

I wrote some code to generate POISSON input files for the R28 (T-shape) and R30 geometries to allow for quick tests of how changes in geometry will impact the gun performance, and also as a cross-check to make sure there is nothing grossly wrong with the CST models. Created in an age when most simulations were done by rearranging rocks¹, POISSON is for cylindrically symmetric problems but very fast compared to CST, and the input file format makes it “easy” to play with parameters that would require a big modeling effort in CST². Realizing that the actual geometry is not cylindrically symmetric and therefore cannot be modeled accurately in POISSON, I note that while the numbers will not be accurate enough to cut metal right away, the point is to let us figure out which parameters drive the performance.


Figure 1 shows the basic R28 model, Figure 2 the basic R30 model. All models assume -200 kV cathode voltage. One could probably remove the bottom half of the model for faster computation without losing much, but this is what it is now.

¹ <https://xkcd.com/505/>

² Easy in the sense that once you have a piece of code that writes the file for you, it can be automated. Writing that code involves some quite obnoxious trigonometry exercises.



← Figure 1: R28 gun, Poisson model. Axes in centimeters, axis of radial symmetry on the left, beam goes up.




← Figure 2: R30 gun, Poisson model. Axes in centimeters, axis of radial symmetry on the left, beam goes up.

3 Focal length and spherical aberration

In any radially symmetric electric field, the radial component $E_r(r, z)$ (which is what focuses or defocuses the beam) depends on the longitudinal component as

$$E_r = -\frac{r}{2} \frac{\partial}{\partial z} E_z(r=0) + \mathcal{O}(r^3). \quad (1)$$

This means one only needs to look at the shape of $E_z(z)$ on axis to get an idea of how the gun will focus. An accelerating field E_z with increasing absolute value (i.e., falling in the plots because negative fields accelerate) will focus, one going toward zero will defocus. Figure 3 compares the as-built R28 and R30 guns.



← Figure 3: Axial (top) and radial (bottom, at $x = 3$ mm) electric field in the R28 and R30 guns simulated with Poisson.

— R28 as built
— R30-3 as built

Both models assume 0.3 mm photocathode recess.

We readily see something is awry: the R30 gun has way less photocathode surface field, which is undesirable in itself due to the surface-charge limit³ but also hints at different focal properties. Although the slope, $\partial E_z / \partial z$, is initially the same in both cases, the decline of the R30 field lasts longer, giving it a larger integrated focusing strength in the first few centimeters. What happens after that has relatively little effect on the beam optics. Looking at the fields only, while the focusing can be redistributed along z , the integrated field must always be defocusing⁴. However, interpreting the radial field as a proportional change in beam angle would only be valid if the particles were at constant r and constant momentum. From

³ G. Mulhollan et al. “Photovoltage effects in photoemission from thin GaAs layers.” In: *Physics Letters A* 282.4 (2001), pp. 309–318. DOI: 10.1016/S0375-9601(01)00202-X

⁴ Because $E_z < 0$ at the cathode and $E_z = 0$ at $z \rightarrow \infty$, one always gets $\int E_r(z) dz < 0$.

the point of view of the beam, the effect of the defocusing field is actually weak compared to any focusing that happens at the cathode because

1. the beam is smaller after having been focused and
2. its momentum is much higher after acceleration.

In the case of a very short focal length such as that of the R30-3 gun, the field geometry at the anode is almost irrelevant because the beam naturally has a waist there (assuming no space charge).

In general, what we need to understand is that

1. the anode always defocuses, no matter how it is shaped, and
2. the cathode can focus or defocus depending on its shape.

Engwall argues in his thesis about the FEL gun⁵ that no focusing should take place at the cathode because the resulting dependency between axial field and radial position will lead to a correlation between transverse position and time of flight, i.e., non-emittance-preserving debunching. I do not recall this ever being talked about in the context of CEBAF, and it may well be worth simulating to see how big the effect is, but in practice, we cannot design the CEBAF gun without focusing anyway because the first lens is too far away and the beam size must be managed. Everyone seems to agree that as a baseline objective, we should aim for focal properties similar to those of the R28 gun, which, assuming no emittance and no space charge, has a beam waist about halfway between the cathode and the first solenoid. Points to consider include:

1. Too big a beam at the first lens will cause lens aberrations to increase the emittance.
2. Too small a beam at the first lens will cause the lens not to focus.
3. Too big a beam anywhere, even if technically within the acceptance, means laser-spot position fluctuations get amplified correspondingly. Conversely, as far as I can tell, this means there is no way to engineer the gun such that the beam position will be insensitive to the laser-spot position because this would imply a collimated, zero-diameter beam envelope. Even if we knew how to produce that for the zero-emittance case, it would not seem like a good idea.


There is a tradeoff between photocathode surface field and focusing strength in that for there to be a focusing field at the cathode,

⁵ D. A. Engwall. "High-brightness electron beams from a DC, high-voltage, GaAs photoemission gun." PhD thesis. University of Illinois, 2000

$E_z(z=0)$ cannot be maximally negative; it has to keep falling for as long as we want E_r to be positive. Everything else being equal, we should always maximize $|E_z(z=0)|$ to combat the surface-charge limit, but the tradeoff with focusing means that for a given focal length, the only available knobs are voltage and gap length, which are in turn both limited by field emission from surfaces other than the photocathode, most notably the tip of the cathode cone. When asked about our pain threshold, Carlos asked me to avoid field strengths in excess of 10 MV/m at the envisioned *conditioning* voltage, which translates to 8 MV/m at 200 kV if we plan to condition the R30 at 250 kV.

It should be noted that the choice of focusing strength does not, in itself, affect the emittance, so it is mostly a practical consideration related to matching the beam to the injector. Space charge makes things more complicated, but because our injector operates with bunches of different and unpredictable bunch charges in any given beam, we cannot use emittance compensation nor design for a particular value of charge and must accept some amount of deterioration, although the bunch charges of interest for CEBAF are not high.

What does increase the emittance is optical aberrations; the most notable issue in our guns is the focal length depending on the position on the photocathode. This is not so much an issue of the optics changing upon a laser spot move as it is about the particles within any given beam not being imaged to a single point. Our laser has an RMS spot width of about 0.5 mm, so 99 % of the particles will be generated within a radius of 1.5 mm, which is not negligible compared to the radius of the active area. As an example, Figure 4 shows the situation for particles on a 2 mm \times 2 mm grid for the case of the R28 gun (POISSON model). At $x = y = 2$ mm, the focal length is only half of what it is in the center.




← Figure 4: Particle trajectories started on a 2 mm \times 2 mm grid in the R28 gun, POISSON model.

In measurements of the R28 gun at CEBAF and UITS, this aberration has always appeared to be much smaller, about 10 %. Whether this is an issue of the measurement or the simulation is not clear

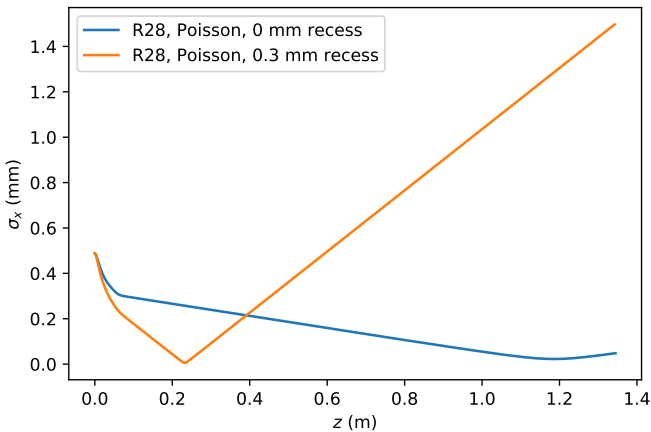
yet and may be revealed at UITF in January when we can try the measurement again with a larger active area. The measured central focal length is also shorter, about 0.4 m in x . To reconcile the discrepancies from the point of view of modeling, we looked for features that may be missing in the model, be they alignment/machining uncertainties or neglected design features. One feature that the optics are very sensitive to is the recess of the photocathode surface behind the electrode hole, which is determined by the mechanical stack-up of the cathode and tantalum ring in the puck. Because the cathode cannot be allowed to touch the back of the electrode, some recess is unavoidable, but based on the puck geometry, it seems to be limited to a few tenths of a millimeter. In principle, this sensitivity is well-known; Sajini studied it for a different gun some time ago⁶. However, the CST models of our guns that we have been using for everything do not include any recess and therefore give wrong numbers for the focal length. Once we agree on a gun design, it will be important to quantify how much variation in beam optics the uncertainty between different pucks can cause. Joe gave me a set of drawings that will let us predict a reasonable range of values; for now, I am working with 0.3 mm recess as a default choice for my simulation models, which appears to be an overestimate. Figure 5 shows how the recess is modeled. The model includes the correct edge radius at the cone hole (according to the drawing), but this radius is difficult to interpolate with any reasonable mesh density and likely of little importance as long as the hole diameter is correct. Figure 6 illustrates the sensitivity to the recess for the examples of 0 mm and 0.3 mm.

⁶ sajini



← Figure 5: Recessed photocathode in the R28 T-shape electrode.

Poisson axes are in cm, axis of radial symmetry on the left, beam goes up.




← Figure 6: GPT comparison of focal properties between 0 mm and 0.3 mm photocathode recess, measured by looking at the RMS width of a Gaussian beam for simplicity. Initial beam size 0.5 mm RMS; the effective focal length depends on this value due to aberration.

For any given puck, we must hope that the recess, while indeterminate, at least remains fixed, i.e., the puck has no degrees of freedom once inserted. Due to the sensitivity of the focusing to the photocathode location, the photocathode absolutely cannot be allowed to move relative to the electrode due to vibrations and such. The wiggly behavior we saw during the hunt for retroreflection when the puck was poked from the back worries me a little in this regard, but it may not mean anything.

Transitioning from the R28 model to the R30 model with the same anode and the same cathode-cone angle, we saw that the focal length indeed dropped to a few centimeters, as observed in the measurements. One reason for this change is the shape of the surrounding electrode, which was not noticed in the design phase because one would think it is so far away from the accelerating gap that the field there should not be affected. While this turns out to be wrong, there is a second smoking gun⁷ that, easy to miss, was overlooked at the time: the hole in the cathode cone has a different diameter, which ends up having the same effect as changing the cone angle by multiple degrees. The cone insert for the spherical electrode, drawing ACC-200-3000-0461, specifies 10.97 mm inner diameter (extrapolated to conical surface, i.e., pretending the sharp edge is not broken), whereas the T-shape electrode, drawing ACC-200-3000-0350, has 12.47 mm. Carlos does not remember how this change came about. Figure 7 illustrates that the hole diameter and the cone angle are interdependent for a given focal length, and that a larger hole diameter is beneficial for field homogeneity. GPT simulations of focal length vs. laser spot position support this interpretation as shown in Figure 8. The focal length is chosen to roughly match that of the R28 T-shape gun for the same cathode recess.

⁷ No pun intended.



← Figure 7: First (top) and second (bottom) derivative of the radial field with respect to the radial coordinate, taken at $r = 1$ mm.

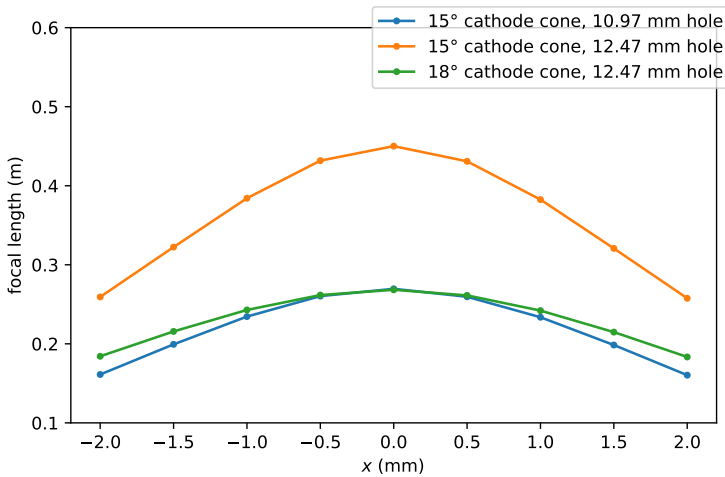
The first derivative represents the focusing strength, whereas the second one is an aberration (E_r should ideally be proportional to r without any curvature).

Comparison between the two hole diameters of interest and two cone angles shows that the larger hole needs a larger cone angle for the same focusing strength but generally gives much lower aberration.

This raises the question of how big we can make the hole. Because of the interdependence between the parameters, this diameter must be fixed (ideally at the maximum value that makes practical sense given the puck geometry) before anything else gets optimized. The larger hole in the T-shape electrode clearly worked, so we should be able to use the same diameter in the R30 design, but can we make it larger still? Is there a possibility of exposing a gap between the GaAs and the Ta ring? If not, is there a downside to exposing the Ta ring? I have not thought about this enough but would certainly like the experts' opinion.

In parallel, let us look at what options we have available to modify the cathode electrode. Assuming we cannot modify the sphere in this design iteration, but we can make a new insert, I experimented with different shapes. The current shape is unfavorable in terms of peak surface field because the radius at the tip is much smaller than that of the sphere, owing to the steep angle at which it joins with the edge of the sphere, see the red detail in Figure 9. The angle at which the cone insert intersects the horizontal at the joint is about 62° as built, whereas for the sphere it is only 27° . This is likely on purpose, but I do not understand the reasons yet.

My parametric model of the cone insert takes the angle between

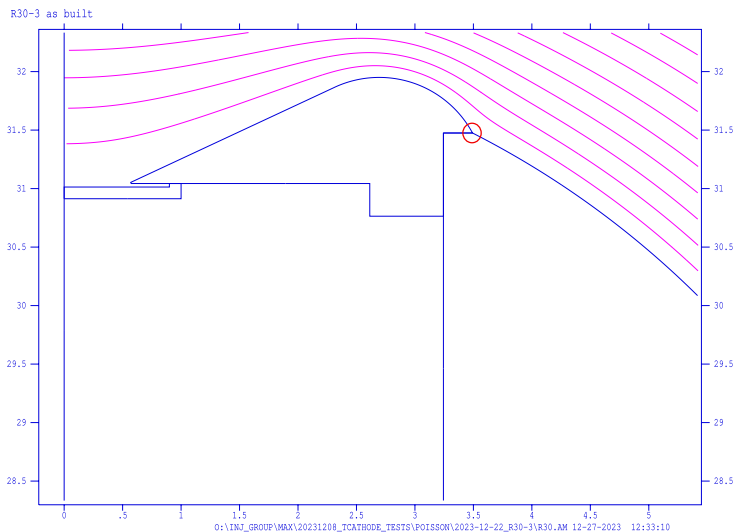


← Figure 8: Focal length as a function of x on the cathode at $y = 2$ mm. Simulated with GPT with a zero-emittance, Gaussian beam with $\sigma_{x,y} = 0.1$ mm.


The larger hole (as in the T-shape gun) needs a 3° higher cone angle for the same central focal length as the smaller hole but shows less aberration.

insert and horizontal as an input such that the radius and center of the tip arc are implicit; the angle between sphere and horizontal cannot be changed as long as we want to keep the sphere as it is. In this way, I am ignorant of any potential difficulty in fabricating or assembling the part that may arise if the angle and radius change; to know what degrees of freedom we have in redesigning the insert, we will need to understand how the parts fit together without causing field enhancement at the joint, as well as any other implications there may be. Figure 10 illustrates how we are free to choose the joint angle, at least on paper. The red, shallow geometry has a tip radius of 2 cm as opposed to 0.9 cm for the blue one, giving a peak surface field of 5.6 MV/m compared to 7.1 MV/m. Because this peak field ultimately determines the lower limit of the gap length (or the upper limit of the operating voltage, depending on how you look at it) and thereby limits the photocathode surface field, making an improvement here seems worthwhile if there is no showstopper. While it may seem unfavorable to introduce curvature in the cone so close to the photocathode from the point of view of field homogeneity, starting at $r = 1.65$ cm in the red geometry, let us put this into perspective by noting that in the T-shape electrode it was even closer than that (1.23 cm).

The final optimization of the cone angle can take place using CST once the other features of the geometry are fixed, particularly the hole diameter and the tip radius. It seems likely that *if the insert can be fabricated with the optimized geometry*, this will be the only part we absolutely have to modify in order to tune the optics the way we want. However, we shall explore if there is something to be gained from changing the anode as well, including potentially



← Figure 9: Both the cone insert and the sphere are specified to have a sharp edge where they join. How does the assembly procedure ensure there can be no field enhancement due to a protruding edge in this location?



← Figure 10: Comparison of different joint angles at the edge of the insert. A smooth joint increases the tip radius, benefitting the surface field at the tip, but it also decreases the straight length of the cone.

moving it toward the cathode a little more if we can afford it (more photocathode field, less asymmetric external field bleeding into the gap).


4 Anode shape

Most reasonably recent publications⁸ assume that the anode should always be conical and have the same cone angle as the cathode electrode. The reason for this must be so obvious that nobody cares to explain it, and Erdong's paper even shows a figure that openly defies the description next to it with no explanation as to why (his cathode cone is 22° and the anode front 37°). The CEBAF anode is evidently a copy of the one described in Dunham's thesis⁹ 30 years ago, which in turn had been copied from a thermionic gun and never revised to meet the needs of a photogun. Seeing as the anode

⁸ J. Maxson et al. "Optimization of DC photogun electrode geometry." In: *Proc. NaPAC 2011*. 2011. URL: <https://accelconf.web.cern.ch/PAC2011/papers/wep245.pdf>, E. Wang et al. "High voltage dc gun for high intensity polarized electron source." In: *Phys. Rev. Accel. Beams* 25 (3 Mar. 2022), p. 033401. DOI: 10.1103/PhysRevAccelBeams.25.033401

and cathode are not strictly individual parts but form the gap field together, it would not be surprising if there were some potential for improvement here, but this is not our first order of business because the cathode is far more important to get right, as explained above.

Shown in Figure 11 (top left), the as-built anode has a front cone angle of 11.25° up until $r \approx 2.2$ cm and then a flat face up to $r = 4$ cm. The angle of the inner cone is 33° and likely of little consequence. The hole radius is 1 cm and considered fixed for now, although it could be increased if deemed beneficial. Decreasing it would not only bring about more defocusing but also a risk of dumping halo on the anode.



⁹ B. M. Dunham. “Investigations of the physical properties of photoemission polarized electron sources for accelerator applications.” PhD thesis. University of Illinois, 1993

← Figure 11: Different anode shapes. The center of the tip arc is the fixed point of reference.

Top left: as built

Top right: cone angle matching cathode, no flat face

Bottom left: no cone

Bottom right: cone as built, but extended flat face

Gabriel and I have tried some different anode shapes and cone angles and independently found the effects on the beam optics to be very small compared to those of the cathode geometry. In principle, we can optimize the cathode independently and postpone the beam-optical optimization of the anode to after that is done. However, optimizing the length of the accelerating gap may play a role here: the surface field on the anode is strongly dependent on its shape and may become the limiting factor. First, the edge radius of the anode at the outer circumference is unduly small and could probably be increased without much harm to the optics. Alternatively, one could consider increasing the outer diameter to move the edge out of harm’s way (Figure 11, bottom right). Second, there is always going to be a lot of field at the edge of the hole, and this gets worse as the cone angle increases. I have not explored the beam-dynamical implications of these choices enough yet but intuitively expect a flat

anode with a generous radius on both edges (like Figure 11 bottom left, but bigger) to give us the least amount of trouble in terms of surface field.

5 Astigmatism and kick

Because our inverted guns lack azimuthal symmetry, some of the field from the HV stalk, NEG screen, etc. will bleed into the accelerating gap, causing different focal lengths in x and y as well as a net kick in y . The R30-3 gun tried to cancel both by tilting the anode. Multipole analysis of the fields shows that the dipole component largely integrates to zero due to the tilt, whereas the quadrupole, which causes the astigmatism, is almost unchanged. The true reason why the R30-3 gun has very little astigmatism is because its focus is close to the location where the quadrupole field peaks such that the particles are not affected by it. This means the astigmatism will reappear once we restore the focal length of the R30 gun to that of the R28. We can choose to ignore it or find some other way to combat it, such as adding some more azimuthally symmetric metal to keep the asymmetric fields out of the gap. But such solutions, if viable, are likely outside the scope of this design revision. At this time, I suspect we will keep the anode tilt if it does no harm, although I personally do not consider the kick a problem as long as it can be steered out, but we will not be able to cure the astigmatism in 2024. With some luck, the Wien quads can make the beam round.

6 Random ideas

Some parts, like the anode standoff, are pretty massive and may not have to be: electrostatics only cares about the surface. As far as I remember, thick metal (more than a couple of millimeters) causes hydrogen outgassing. Can we make Marcy happy by removing some metal that is not structurally needed?

Are the bore shield tubes long enough to fully screen any stray field from ceramics?