

Barrel Imaging Calorimeter Light Guide Simulations

Tegan Beattie, Aram Teymurazyan, Zisis Papandreou - Winter 2025

The goal of these tests is to estimate the transmission efficiency (the percentage of thrown photons that reach the sensitive volume of the SiPM) and degree of light mixing (some measure of how uniformly the light illuminates the SiPM surface) for BIC light guides of various designs and lengths. Since space in the ePIC detector is extremely limited, knowing how short the light guides can reasonably be is important. In the past, Elton Smith had done these studies for the GlueX light guides using Geant3, results we wish to verify using Geant4 and extend to BIC light guides.

1 Elton's Studies

Over the years, Elton conducted various studies testing different light guide geometries before settling on the final GlueX design. These studies are detailed in various GlueX DocDB documents, the latest of which is GlueX-doc-1784 [1]. Table 1 of this document shows the dimensions and efficiencies of the 80-mm-long GlueX production light guides, called 'g1' through 'g10,' with the number referring to the layer, ordered radially out from the center of the detector. Using these documents and a somewhat outdated copy of the his Geant3 code (found at <https://halldsvn.jlab.org/repos/trunk/home/elton/geantguide/>), we were able to piece together the details of Elton's simulations.

The large trapezoidal base of a simulated light guide lies in the x-y plane and the z-axis lies along the length of the guide. Optical photons (particle type 50 in Geant3) were thrown uniformly from the large trapezoidal base of the light guides in the z-direction flat in ϕ (azimuthal angle) and flat in $\cos\theta$ to approximately mimic the distribution of photons emerging from an optical fiber matrix. The maximum polar angle for emerging photons, θ_{\max} , depends on the material makeup and geometry of the fibers. For his studies, Elton used θ_{\max} values found in the documentation for Kuraray's single- and double-clad fibers [2]. The relevant figure for the Kuraray double-clad fiber is reproduced in Figure 1.




Figure 1: Visual representation of the maximum polar angle of scintillation light that escapes the end of a double-clad Kuraray fiber [2].

For a double-clad fiber, $\theta_{\max} = 26.7^\circ$, and for a single-clad fiber, $\theta_{\max} = 20.4^\circ$. Since GlueX eventually settled on using double-clad fibers, Elton's final results use $\theta_{\max} = 26.7^\circ$. The generated photons pass through the light guide, which is modeled as acrylic with an index of refraction $n = 1.49$ and attenuation length $\lambda = 240$ cm, through a 0.5 mm air gap ($n = 1.00$, $\lambda = 30000$ cm), then into the SiPM. The SiPM is modeled as a 0.45-mm-thick acrylic window, 12.7 mm \times 12.7 mm in surface area, in front of a silicon sensitive volume of the same shape. A rendering of this setup using Geant4 is shown in Figure 2. The positions of any photons which make it to this sensitive volume are recorded on entry.

We set up the same simulation in Geant4, copying the GlueX light guide and SiPM geometries and the optical properties of the light guide, air, and SiPM window. We used `G4EmStandardPhysics_option4()` and `G4OpticalPhysics()` for shooting `G4OpticalPhotons`. Due to rounding errors, the trapezoidal prisms generated using the light guide dimensions in Elton's work aren't planar (some sides of the guides have four vertices that do not share a common plane). Geant4 is more strict than Geant3 when building trapezoidal prisms, so a procedure was needed for fine-tuning the dimensions of the light guides to make the surfaces sufficiently planar. This produce is discussed in Appendix A. The efficiencies we found with Geant4 were lower than what Elton had found for the three layers we tested:




Figure 2: Geant4 rendering of the simulation setup used in Elton’s Geant3 studies featuring a g6 light guide and a SiPM separated by a 0.5 mm air gap.

- Layer 1: 65.14 (G4) compared to 74.8 (Elton’s G3)
- Layer 6: 56.19 (G4) compared to 64.3 (Elton’s G3)
- Layer 10: 41.75 (G4) compared to 47.7 (Elton’s G3)

As a check, we re-built the simulation again in Geant3, primarily to check that we were using the correct optical properties and geometries that were used in Elton’s simulations, and we indeed saw similar efficiencies to Elton’s old numbers. As such, we concluded that there was a difference in how the G3 and G4 simulations were handling the physics.

After corresponding with Richard Jones, the author of the optical photon propagation code in Geant3, we surmised that the photon polarization was probably not being set properly in the Geant3 simulation. To test this, we constructed, in both G3 and G4, a simple simulation to test the implementation of the Fresnel equations. These equations allow for calculations of the percentage of incident photons that are transmitted and reflected at optical boundaries and are dependent on the indices of refraction of the materials involved, the angle of incidence of the photon, and the polarization of the photon, and are built-in to both G3 and G4. We threw photons at a fixed angle at an acrylic-air interface and counted the number of photons passing into the air volume in G3 and G4 using the same optical photon generation techniques as used in the light guide simulations. In G4, the Fresnel equations predicted perfectly what we saw in the simulations at each incident angle we tested, but in G3, we saw full transmission regardless of the incident angle. After some more digging into the G3 optical photon code and with some more assistance from Richard, we found that there was no mechanism coded by which a thrown optical photon could be assigned a polarization and no default polarization would be assumed, so the Fresnel equations would be ignored. The polarization could be assigned only when generating an optical photon as a result of something like Čerenkov radiation. Richard provided a workaround by throwing a geantino (particle type 48 in Geant3) rather than an optical photon, then immediately stopping that geantino and generating the optical photon using the geantino’s position and trajectory. This generation call also allows for a polarization assignment. For these simple tests, we assigned a fixed polarization and observed the behaviour we expected from the Fresnel equations.

With the discrepancy between G3 and G4 GlueX light guide efficiencies understood, we concluded that the G4 simulation (located online at <https://github.com/beattite/BIC-Light-Guide-Simulations>) was working as intended, providing accurate results that could be extended to BIC light guide geometry and length studies.

2 BIC Light Guide Studies

2.1 Sector Layout and Light Guide Shapes

In order to fully cover the fiber matrix of a sector (called a 'module' in GlueX) with light guides, each layer consists of 5 guides with trapezoidal bases. The angle of the trapezoid side of each guide is equal to the angle of the sector side. Since there are 48 sectors, each sector subtends a polar angle of 7.5° , so the angle between the inner radial edge of the sector and its side is 93.75° . The five guides in each layer are arranged alternating in vertical orientation so the guides fill the space on the sector face, as in Figure 3a. The lengths of the top and bottom of the input trapezoidal face are calculated from the total width of the sector at the top and bottom of each layer. In this layout, each light guide in a given layer of a sector lies the same distance away from the inner radial edge of the sector. This layout differs from GlueX's light guide layout, which has 4 trapezoidal light guides per layer. Each GlueX guide has the same orientation and shape (the angle between the bottom edge of the guide and its side is 90.9375°), arranged in an arc such that the centers of each light guide in a given layer share a radial coordinate, as in Figure 3b.




Figure 3: Light guide layout scheme for a) ePIC's BIC sector and b) GlueX's BCAL module. Diagrams are to-scale relative to each other. Layer 7 light guides are shown in each diagram.

The detectors which will be used for each channel of the BIC are the Hamamatsu S13 or S14 SiPM arrays (4×4 grids of $3 \text{ mm} \times 3 \text{ mm}$ SiPMs) with $50 \mu\text{m}$ pixels [3][4], the physical dimensions of which are $13 \text{ mm} \times 13 \text{ mm}$ (the dimensions of the active area of the SiPM array are $12.6 \text{ mm} \times 12.6 \text{ mm}$ with 0.2 mm dead crosses between individual SiPMs)¹. The purpose of the BIC light guide is to funnel light from the fiber matrix to the face of the SiPM array. Various light guide shapes can be used to do this, and part of these studies is to quantify the differences between a few of them in terms of efficiency and the degree of light mixing.

The most straightforward shape is what we call the 'nominal' shape (designation 'b'), for which the output face of the guide matches the physical dimensions of the SiPM array. This shape brings the trapezoidal input face to a square output face using two curved sides. Figure 4a shows a Geant4 render of such a light guide, with the top and bottom sides being planar and the left and right sides having a smooth curve which brings the 93.75° input trapezoid angle to the 90° output square angle. In Geant4, this shape is achieved through the use of the G4GenericTrap class, which automatically segments the volume into a series of planar trapezoidal prisms. In terms of production, these are likely the most difficult to handle, since the curved sides could make these guides more tricky to machine (and, if necessary, polish) compared to guides with all planar sides.

The simplest shape is what the 'trapezoidal' shape (designation 'bt'), for which the output face of the guide is trapezoidal with the same side angle as the input trapezoid (See Figure 4b). This is a fully planar trapezoidal prism, much like the GlueX light guides, which makes for easy machining and polishing. However, since the output face does not match the dimensions of the SiPM array, either light will be lost out of the top corners of light guide (in the case where the bottom edge of the output face matches the SiPM array size) or no light will be deposited in the bottom corners of the SiPM array (in the case where the top edge of the output face matches the SiPM array

¹Individual SiPM pixels are not simulated in these studies.




Figure 4: Geant4 renders of the Nominal (‘b’), trapezoidal (‘bt’), and hexagonal (‘bh’) BIC light guide shapes. The active SiPM area is shown in red.

size). There are many options here for the exact size of the output trapezoid, but for simplicity, only results where the bottom edge of the output trapezoid matches the SiPM size are presented. The top edge length is calculated based on the height of output face (which also matches the height of the SiPM array) and the input trapezoid angle.

Figure 4c shows a second option, termed ‘hexagonal’ (designation ‘bh’), for bringing the trapezoidal input face to a square output face. This shape replaces each of the curved sides in the nominal design with two triangular planar sides. In Geant4, this shape constructed as the boolean intersection between the trapezoidal shape and a square input, square output prism, with the input square dimension equal to the top length of the nominal base trapezoid and the output square dimension equal to the physical size of the SiPM array. This shape allows the light output to match the SiPM array dimensions, as with the nominal shape, but uses planar sides, which should make machining and polishing more consistent.

Table 1 gives the dimensions of the input and output faces of each tested shape and layer, with the designation again being given as [shape designation][layer number], with layers numbered radially out from the center of the detector. As with the GlueX light guides, any trapezoidal prisms (‘bt’ guides or intermediate prisms used in boolean operations) are run through the parameter adjustment procedure in Appendix A.

2.2 Material Optical Properties

For the BIC simulations, the acrylic light guides have the same absorption length, λ , as Elton’s studies, 240 cm, but the index of refraction, n , is variable depending on the wavelength of light. We generate photons with the same distribution as Elton’s study at a wavelength of 450 nm. At this wavelength, $n_{\text{acrylic}} = 1.499906 \approx 1.5$. We simulate events with a 0.5 mm air gap ($n \approx 1.00028$, $\lambda = 300$ m) and with a 1-mm-thick silicone optical cookie ($n = 1.43$, $\lambda = 240$ cm). The silicone cookie’s thickness and index of refraction are based on a standard product, ‘EJ-560,’ from Eljen Technology [5]. The SiPM is modeled after an S13 SiPM array[3], with a 0.1-mm-thick epoxy resin window ($n = 1.55$)². The absorption length of the window is set arbitrarily to $\lambda = 420$ cm, and the index of refraction of the silicon crystal itself is set to match that of the window. A ‘shield’ extending from the output face of the light guide and surrounding the SiPM is also implemented to stop photons from entering the SiPM crystal (sensitive volume) from anywhere other than through the output face of the light guide. This is implemented mainly for tests with short light guides, where photons can leave the side of the light guide with a steep enough angle to enter the side of the window, air gap, or optical cookie.

For some studies, θ_{max} for both double-clad (27.6°) and single-clad (20.4°) fibers were simulated. The assumed cladding and gap material will be specified in figures and/or body text for each study.

2.3 Efficiency and Light Mixing vs. Length

The efficiency calculated in these studies is defined as the percentage of thrown photons arriving at the face of the silicon crystal through the SiPM window (called ‘successful’ events). In this way, it’s a combined efficiency of the light guide, gap, and SiPM window. It does not account for the photon detection efficiency (PDE) of the SiPMs, cross-talk, or the dead crosses between SiPMs in the SiPM array.

Aside from efficiency, the light guides are required to ‘mix’ the light coming from the fibers so that, ideally, light entering the light guide at any point on the input face is equally likely to exit the output face anywhere. This is to help mitigate saturation effects. Because each pixel in a SiPM is sensitive only to whether a photon is present or not, it cannot distinguish between one photon and multiple photons arriving in the same event. If there is too

²An S14 SiPM array, which has a 0.15-mm-thick glass epoxy window ($n = 1.57$) [4], was also studied. The results using the S13 and S14 arrays were identical.

Table 1: Input and output face dimensions of BIC light guides of various shapes and layers. Trapezoid dimensions are given as ‘bottom length - top length \times height.’

Light Guide Designation	Input Dimensions (mm)	Output Dimensions (mm)
b1 & bh1	20.691-23.505 \times 21.463	13.00-13.00 \times 13.00
b2 & bh2	21.707-24.521 \times 21.463	13.00-13.00 \times 13.00
b3 & bh3	22.723-25.537 \times 21.463	13.00-13.00 \times 13.00
b4 & bh4	23.739-26.553 \times 21.463	13.00-13.00 \times 13.00
b5 & bh5	24.755-27.569 \times 21.463	13.00-13.00 \times 13.00
b6 & bh6	25.801-28.504 \times 20.625	13.00-13.00 \times 13.00
b7 & bh7	26.360-29.063 \times 20.625	13.00-13.00 \times 13.00
b8 & bh8	26.893-29.597 \times 20.625	13.00-13.00 \times 13.00
b9 & bh9	27.452-30.155 \times 20.625	13.00-13.00 \times 13.00
b10 & bh10	28.011-30.714 \times 20.625	13.00-13.00 \times 13.00
b11 & bh11	28.544-31.248 \times 20.625	13.00-13.00 \times 13.00
b12 & bh12	29.103-31.806 \times 20.625	13.00-13.00 \times 13.00
bt1	20.691-23.505 \times 21.463	13.00-14.7044 \times 13.00
bt2	21.707-24.521 \times 21.463	13.00-14.7044 \times 13.00
bt3	22.723-25.537 \times 21.463	13.00-14.7044 \times 13.00
bt4	23.739-26.553 \times 21.463	13.00-14.7044 \times 13.00
bt5	24.755-27.569 \times 21.463	13.00-14.7044 \times 13.00
bt6	25.801-28.504 \times 20.625	13.00-14.7037 \times 13.00
bt7	26.360-29.063 \times 20.625	13.00-14.7037 \times 13.00
bt8	26.893-29.597 \times 20.625	13.00-14.7037 \times 13.00
bt9	27.452-30.155 \times 20.625	13.00-14.7037 \times 13.00
bt10	28.011-30.714 \times 20.625	13.00-14.7037 \times 13.00
bt11	28.544-31.248 \times 20.625	13.00-14.7037 \times 13.00
bt12	29.103-31.806 \times 20.625	13.00-14.7037 \times 13.00

little mixing, an event which lights up only a portion of the input fibers will correspondingly light up only a portion of the SiPM array, focusing all the scintillation light from the event on a small number of pixels. In this case, the probability that we lose signal due to saturation is higher compared to the case where light is mixed perfectly. There are a few things we look at to determine the degree of light mixing. The first are ‘correlation plots.’ These are a set of four plots showing possible correlations between entrance positions (thrown positions, in our case) and exit positions (where photons leave the output face of the light guide) of photons which make it to the SiPM crystal.

Figure 5a shows an example of correlation plots under poor light mixing conditions. There are clear correlations between entrance and exit positions. In this specific example, the light guide is only 5 mm long, so most light thrown near the edges of the light guide (more extreme x and y values) are lost out the sides of the guide, while most light that makes it through to the SiPM crystal do so in a mostly-straight line without reflecting off the walls of the guide, causing a linear correlation between entrance x and exit x and between entrance y and exit y . In contrast, Figure 5b shows correlation plots with good light mixing, with no obvious correlations between entrance and exit positions whatsoever. Here, for photons that reach the SiPM crystal, a given entrance position has no effect on the exit position, which ensures that the SiPM will uniformly illuminated even if only a small number of fibers on the input face of the guide are providing light. This example uses a light guide length of 100 mm, which, coupled with the angular distribution of the thrown photons, ensures that most of the photons internally reflect at least once along their path to the SiPM crystal. In addition, the extra length makes the angle of the light guide wall less steep relative to the photon’s travel direction, so less light is lost to refraction compared to a shorter guide.

The second thing we look at is a ‘detection plot.’ This is a plot showing the position successful photons enter into the SiPM crystal. As mentioned above, if we have good light mixing, photons should be evenly spread over the SiPM crystal regardless of the input distribution of the light. We look at the detection plot when using a ‘quadrant source,’ restricting the thrown photons to only one quadrant of the input face of the light guide. Figure 6a shows an example of a quadrant source detection plot under conditions of poor light mixing (a 5-mm-long light guide), with a clear enhancement in the upper-right corner of the plot and very few successful events in other regions, and Figure 6b shows an example of a detection plot with good light mixing (a 100-mm-long light guide), with a uniform distribution of photons over the SiPM crystal. These plots, and all others like them in these studies, are normalized between values of 0 and 1 unless otherwise labeled. The idea of looking at detection plots for light uniformity can be extended to other input distributions (single fibers, line sources representing a particle track, etc.), but only

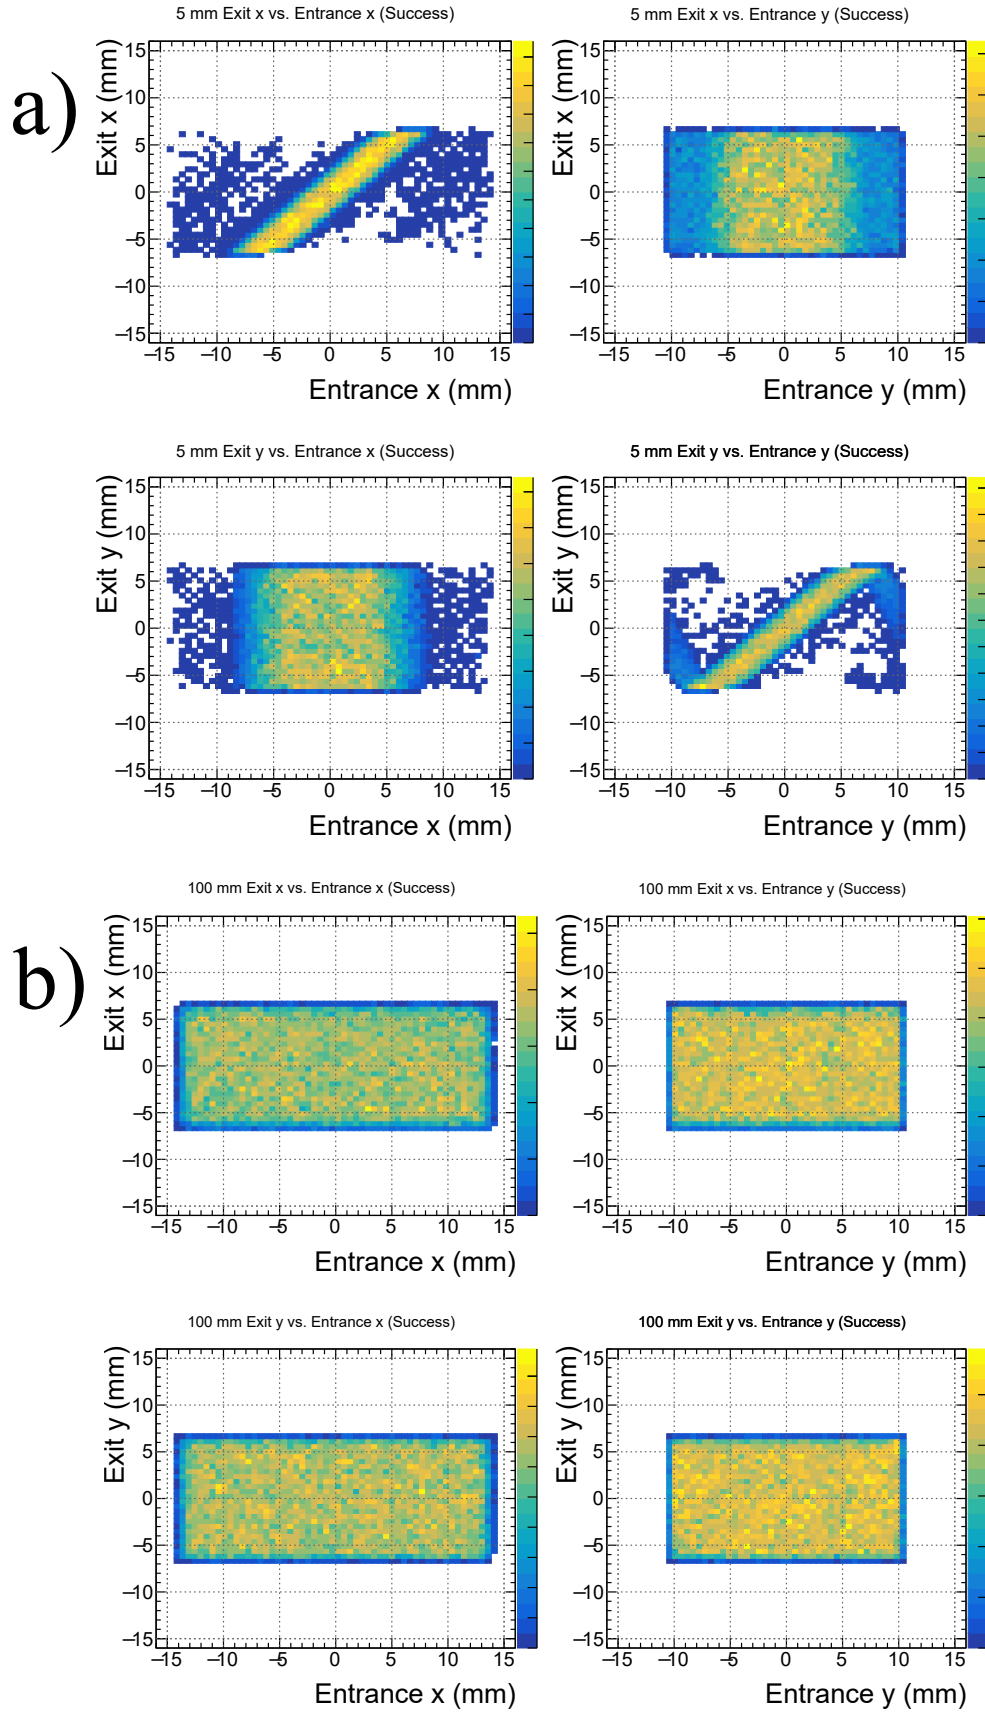


Figure 5: Example correlation plots of light guide entrance position and exit position of successful events for a) a 5-mm-long ‘b6’ guide (poor light mixing) and b) a 100-mm-long ‘b6’ guide (good light mixing). Plots normalized between values of 0 (blue) and 1 (yellow).

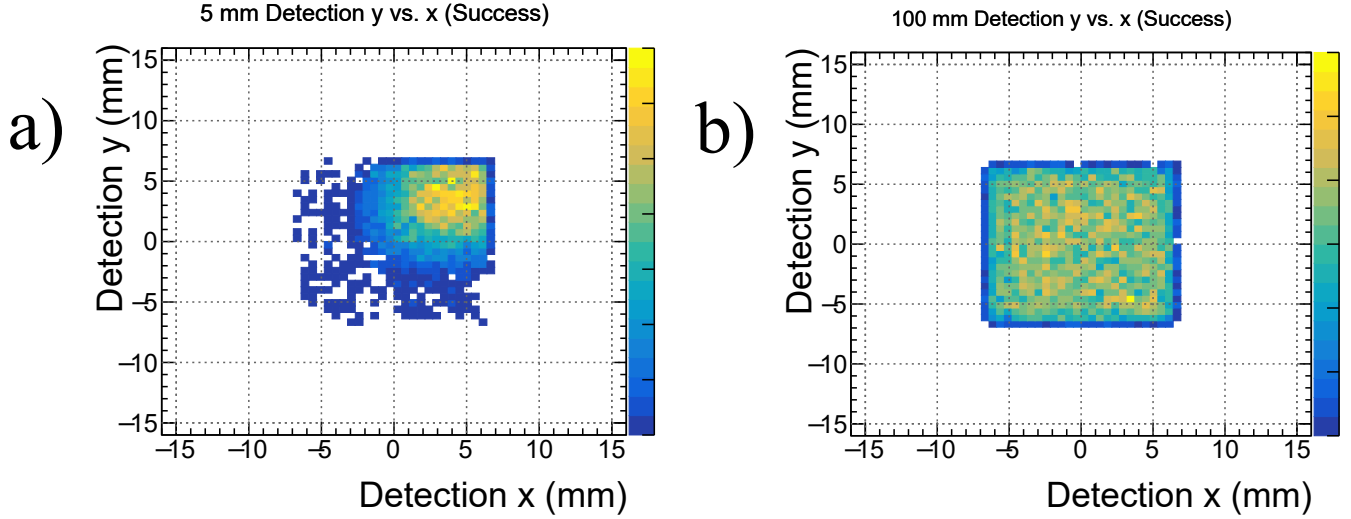


Figure 6: Example detection plots of SiPM array hit position of successful events using a quadrant source for a) a 5-mm-long ‘b6’ guide (poor light mixing) and b) a 100-mm-long ‘b6’ guide (good light mixing). Plots normalized between values of 0 (blue) and 1 (yellow).

quadrant sources have been considered in our studies thus far.

One way of quantifying the degree of light mixing is with a uniformity metric. In this work, we use the NAAD (normalized absolute average deviation) metric [6], given in Equation 1.

$$\text{NAAD} = 1 - \frac{1}{N\bar{Y}} \sum_{i=1}^N |Y_i - \bar{Y}| \quad (1)$$

Here, N is the number of pixels in the region of interest, Y_i is the value of bin i in the region of interest, and \bar{Y} is the average bin value of the region of interest. For our detection plots, we remove the two outermost rings of bins (the data covers 24 bins in x and y with our histogram binning scheme, and we take the innermost 22×22 bins as our region of interest). A NAAD value can be between 0 and 1, with a value of 1 meaning perfect uniformity and lower values correspond to less uniformity.

Each light guide shape and layer was tested at various lengths, with the air gap and with the silicone cookie, and assuming double-clad and single-clad fibers. The interesting studies and results are summarized below.

2.3.1 Air Gap vs. Optical Cookie

BIC plans to use optical cookies between the light guides and the SiPM arrays, unlike GlueX, which used an air gap. The difference in efficiency between the two strategies is demonstrated in Figure 7, which shows efficiency vs. layer number for various light guide lengths, ranging from 20 mm to 80 mm. With the same light guides, SiPM, and input light distribution, the 1-mm-thick optical cookie results in 50-60% higher efficiency than the 0.5 mm air gap. In both cases, the efficiency is not much effected by lengthening the guides past 40 mm or so, in agreement with Elton’s older studies which suggested the same thing.

Figure 8 shows the NAAD of the detection plots vs. layer number for the same light guide lengths, again for the air gap and the optical cookie. The optical cookie performs better in this metric for lengths of 40 mm and below, but the difference mostly disappears above 40 mm. For longer light guides, the air gap outperforms the optical cookie in terms of light mixing, if only by a percent or two, because light exiting the guide will refract at a larger angle in air than in the cookie, which helps spread the light onto the SiPM face.

The difference in light mixing between an air gap and the optical cookie is small for light guides of 45 mm length or more, but the efficiency difference is considerable. Studies still need to be done to see what sort of efficiency gain we see with real light guides and optical cookies, but for the rest of these simulation studies, an optical cookie will be used.

Figure 9 shows the detection plots (for a quadrant source) and a selection of the correlation plots (for a complete source) for the optical cookie simulations from this study. The rows give plots for some of the light guide lengths studied. The 20 mm length plots show clear features which flatten out gradually as the length increases. Based on these plots, along with the efficiencies and NAAD values, we conclude that light guides of 45 mm to 60 mm length with an optical cookie should provide adequate light mixing while maintaining good efficiency.

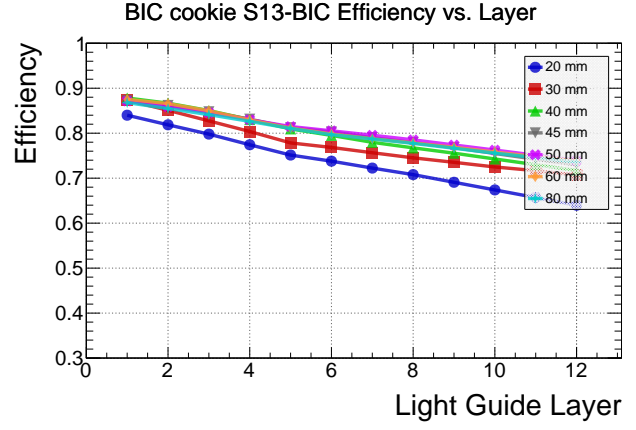
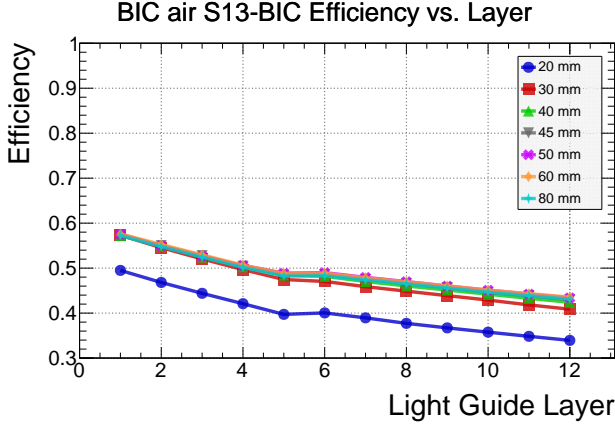


Figure 7: Efficiency vs. layer number for double-clad fibers and nominal BIC guides of various lengths with a 0.5 mm air gap (left) and a 1-mm-thick optical cookie (right).


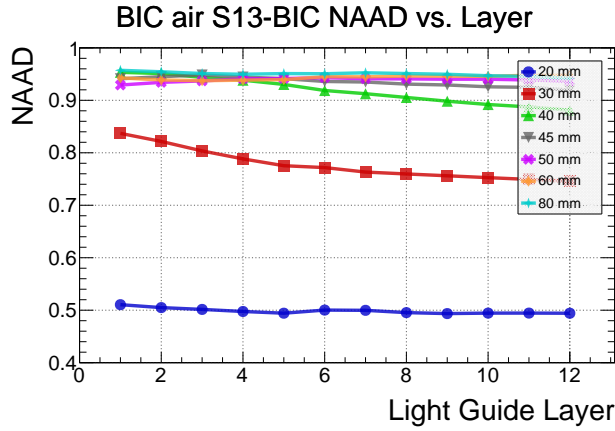


Figure 8: NAAD vs. layer number for double-clad fibers and nominal BIC guides of various lengths with a 0.5 mm air gap (left) and a 1-mm-thick optical cookie (right).


2.3.2 Single-Clad vs. Double-Clad

Fibers with only a single layer of cladding capture light by total internal reflection up to a smaller polar angle than double-clad fibers. As a result, single-clad fibers will transport less scintillation light to the detectors. For this study, only the effect of changing the maximum polar angle of the thrown photons is studied (we generate the same number of photons in each case). For single-clad fibers, we use $\theta_{\max} = 20.4^\circ$, again taken from the Kuraray fiber documentation [2].


Figure 10 shows the efficiency vs. layer number for various light guide lengths assuming double-clad and single-clad fibers. Because the thrown light is restricted to a tighter cone for single-clad fibers, a higher percentage of the light passes through the guide with fewer reflections off the walls. As a result, less light reaches the output face at extreme angles, so the efficiencies with single-clad fibers are higher than with double-clad fibers.

By the same token, light mixing suffers. Less reflections means the light more often passes straight through the guide, resulting in more obvious correlation between the input distribution and output distribution. The NAAD metric for each cladding scheme is shown in Figure 11, again vs. layer number and for various guide lengths. The difference in the flatness metric is slight, but the single-clad fibers do shift the NAAD down compared to the double-clad fibers. The effect on light mixing is more readily seen in this case with the detection (quadrant source) and correlation (complete source) plots, shown for various guide lengths in Figure 12. In comparison with the same plots for double-clad fibers (Figure 9), the features and spatial correlations are more pronounced and are easily visible in even the longer guides. More investigation is to be done into understanding the degree of light mixing required for the experiment.


20 mm:




30 mm:



40 mm:



50 mm:



60 mm:




Figure 9: Detection (left) and correlation (middle and right) plots for BIC guides of various lengths (rows) with double-clad fibers and a 1-mm-thick optical cookie. Plots normalized between values of 0 (blue) and 1 (yellow).





Figure 10: Efficiency vs. layer number for nominal BIC guides of various lengths with a 1-mm-thick optical cookie assuming double-clad fibers (left) and single-clad fibers (right).





Figure 11: NAAD vs. layer number for nominal BIC guides of various lengths with a 1-mm-thick optical cookie assuming double-clad fibers (left) and single-clad fibers (right).

2.3.3 Alternate Shapes

Two alterations to the nominal BIC light guide shape were tested: a trapezoidal guide and a hexagonal guide, as described in Section 2.1. Figure 13 shows the efficiency vs. layer number for 50-mm-long guides of these three shapes with double-clad and single-clad fibers. Apart from some minor differences in the trends between the shapes, the only major difference is the relatively low efficiency of the trapezoidal guide for lower layer numbers.


Similarly, the differences in the flatness metric between the shapes is minimal. Figure 14 shows the NAAD vs. layer number for the same 50-mm-long guides with double- and single-clad fibers. As an aside, Figures 13 and 14 give a more readable comparison between single- and double-clad fibers for one of the tested guide lengths than the overview figures from the previous section.

Based on these studies, we have settled on the ‘nominal’ BIC light guide shape and a light guide length of 50 mm. Additional Lab studies can be done when we have access to prototype light guides of various shapes and lengths.


2.4 BIC Mini Light Guides

Since space at the end of the sectors is short, we also investigated the effects of using smaller SiPMs and half-sized light guides. The idea is to use 6 mm × 6 mm SiPMs rather than the S13 or S14 arrays and replace each light guide with four smaller guides, which should preserve the approximate shape and reflection properties of the guides while potentially saving a few centimeters of space. Because of the light guide input trapezoid shape and layout scheme where the same shape of light guide is alternated in vertical orientation to fill a layer, an odd number of guides is required to fill the space in a layer of a sector, so rather than expanding from 5 guides per layer to 10, we use either 9 or 11 guides. In this study, 11 was chosen.


20 mm:




30 mm:



40 mm:



50 mm:



60 mm:




Figure 12: Detection (left) and correlation (middle and right) plots for BIC guides of various lengths (rows) with single-clad fibers and a 1-mm-thick optical cookie. Plots normalized between values of 0 (blue) and 1 (yellow).

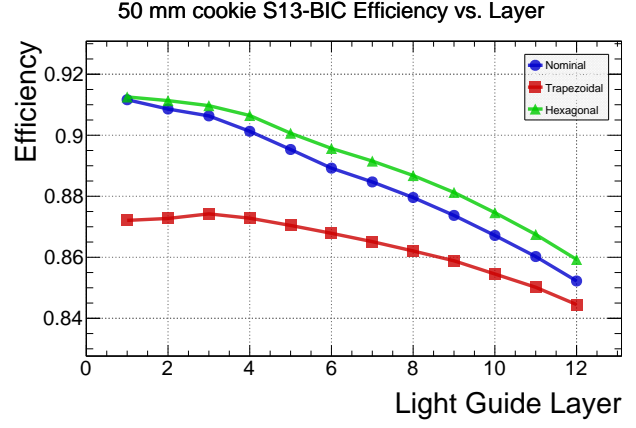



Figure 13: Efficiency vs. layer number for 50-mm-long BIC guides of various shapes with a 1-mm-thick optical cookie assuming double-clad fibers (left) and single-clad fibers (right).

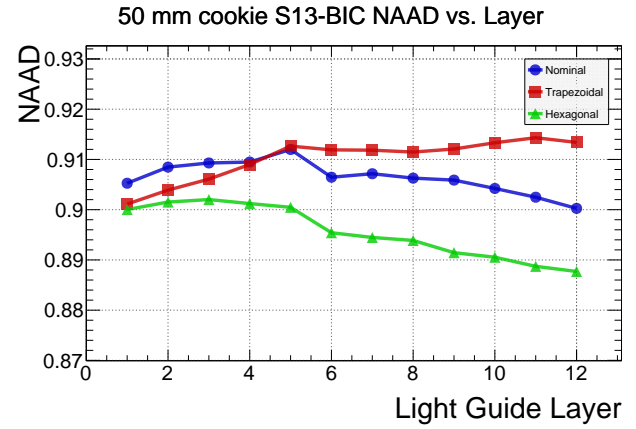



Figure 14: NAAD vs. layer number for 50-mm-long BIC guides of various shapes with a 1-mm-thick optical cookie assuming double-clad fibers (left) and single-clad fibers (right).

Using these mini guides, we would have now 24 layers per sector. In this study, we only constructed 12 layers, each based on the existing full-sized BIC layers. The heights of the mini layers are half the heights of the nominal layers, and the output face is a 6 mm × 6 mm square (rather than a 13 mm × 13 mm square). The top and bottom lengths, l_t and l_b , of the input trapezoid for each layer are calculated based on the widths of a sector along the top and bottom of the layer, w_t and w_b , as in Equation 2.

$$\begin{aligned} 6l_t + 5l_b &= w_t \\ 5l_t + 6l_b &= w_b \end{aligned} \quad (2)$$

w_b for a given mini layer is the same as w_b for the corresponding nominal layer with 5 full-sized guides and w_t is calculated based on w_b , the half-height of the nominal guides, and the trapezoid angle of the sector. With this procedure, we are testing every second mini layer starting with the innermost one. Table 2 gives the dimensions of the input and output faces of each of the 12 mini guides tested.

Two SiPM models are used for this study: a Hamamatsu S13 SiPM [7], which has a 0.4-mm-thick silicone resin window with an index of refraction of 1.41, and a Hamamatsu S14 SiPM [4], which has a 0.15-mm-thick glass epoxy window with an index of refraction of 1.57. The same absorption length of $\lambda = 420$ cm was used for the window. This short study was done only with a silicon optical cookie and with double-clad fibers.

Figure 15 shows the efficiency and NAAD vs. layer number for various mini light guide lengths ranging from 10 mm to 40 mm using the S13 SiPM. Past 25 mm in length, the efficiency and NAAD stop improving. Figure 16 shows the same plots using the S14 SiPM, which provides slight improvements to both the efficiency and NAAD compared to the S13 SiPM. Compared to the nominal guides, however, the NAAD and efficiency are quite a bit lower. Figure 17 shows better the quality of light mixing from these guides; little correlation is visible at 25 mm and

Table 2: Input and output face dimensions of every second layer of BIC mini light guides assuming 11 light guides per layer.

Light Guide Designation	Input Dimensions (mm)	Output Dimensions (mm)
bm1	9.4049-10.8121 x 10.7315	6.00-6.00 x 6.00
bm2	9.8667-11.2739 x 10.7315	6.00-6.00 x 6.00
bm3	10.3285-11.7358 x 10.7315	6.00-6.00 x 6.00
bm4	10.7903-12.1976 x 10.7315	6.00-6.00 x 6.00
bm5	11.2521-12.6594 x 10.7315	6.00-6.00 x 6.00
bm6	11.7379-13.0791 x 10.3125	6.00-6.00 x 6.00
bm7	11.9820-13.3332 x 10.3125	6.00-6.00 x 6.00
bm8	12.2240-13.5762 x 10.3125	6.00-6.00 x 6.00
bm9	12.4784-13.8295 x 10.3125	6.00-6.00 x 6.00
bm10	12.7325-14.0836 x 10.3125	6.00-6.00 x 6.00
bm11	12.9745-14.3266 x 10.3125	6.00-6.00 x 6.00
bm12	13.2288-14.5800 x 10.3125	6.00-6.00 x 6.00

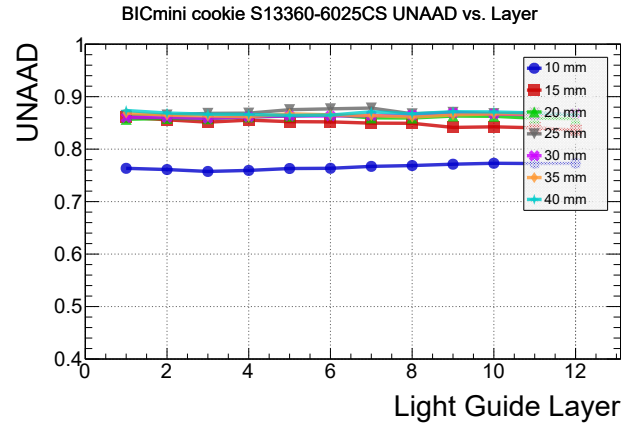
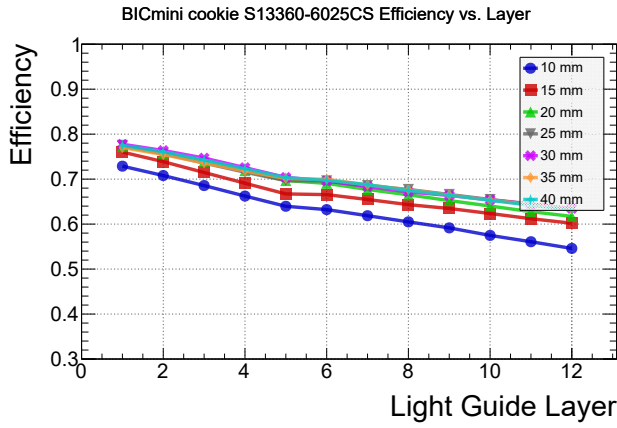


Figure 15: Efficiency vs. layer number (left) and NAAD vs. layer number (right) for double-clad fibers and BIC mini guides of various lengths with a 1-mm-thick optical cookie and an S13 SiPM.

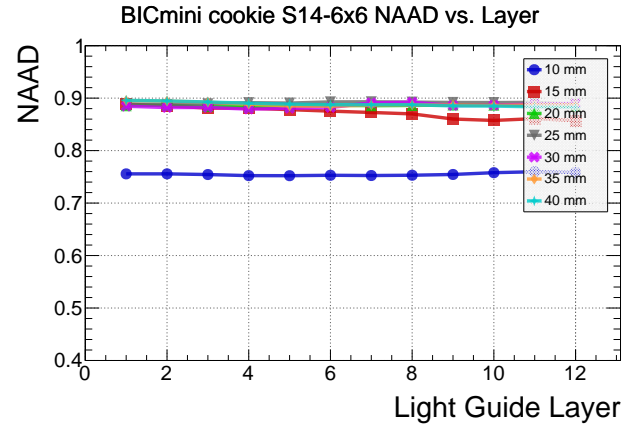
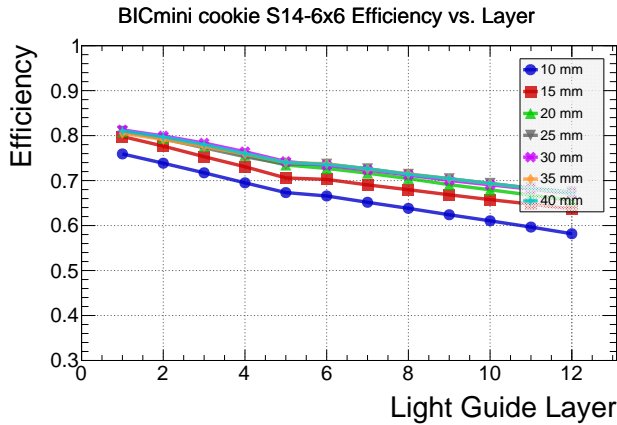




Figure 16: Efficiency vs. layer number (left) and NAAD vs. layer number (right) for double-clad fibers and BIC mini guides of various lengths with a 1-mm-thick optical cookie and an S14 SiPM.


10 mm:




15 mm:



20 mm:



25 mm:



30 mm:




Figure 17: Detection (left) and correlation (middle and right) plots for BIC mini guides of various lengths (rows) with double-clad fibers and a 1-mm-thick optical cookie and an S13 SiPM. Plots normalized between values of 0 (blue) and 1 (yellow).

above in length, and the detection plots for these longer guides look reasonably flat, though it is more difficult to discern if there are significant enhancements in the detection plots with the reduced number of histogram bins. The drop in efficiency near the edges also has more impact on the overall efficiency of these guides. With smaller guides and SiPMs, there is more edge area and less inner area, so the efficiency drop near the edges has a greater effect relative to guides and SiPMs with larger dimensions. The efficiency of these mini guides is about 10% lower than the nominal guides with S14 SiPMs. This, coupled with the factor of 4.4 increase in the total number of readout channels (or some channel summing scheme), makes this option undesirable.

3 Complementary Study: SiPM linearity

Using the data generated in the previous studies, we also look at the effects of SiPM linearity. As mentioned above, non-linearity effects arise when multiple photons strike the same SiPM pixel, and during an event with thousands of photons, this may happen enough that the signal coming out of the SiPM no longer matches the number of photons incident on the SiPM. Because the chances of overlapping photons increases with the number of incident photons, however, this effect may be corrected for in data by estimating the average non-linearity in events of varying sizes.

The data from the previous studies' simulations are stored in ROOT trees which we fill into 2D histograms with bins that are the same size as the physical pixels on the SiPM arrays. The particular SiPMs we are considering for the BIC have a $50\text{ }\mu\text{m}$ pitch. Since the SiPM array is $12.6\text{ mm} \times 12.6\text{ mm}$ in area, the histogram is 252 bins in each dimensions for a total of 63504 bins. Further, any photons landing in the dead crosses between SiPMs in the array are ignored, so we're left with a map of how many photons strike which SiPM pixels in an event with a variable number of photons. Multiple trials are done with each number of thrown photons to get averages for the number of photons that reach the SiPM, the number of pixels fired, and the linearity, defined as the number of pixels fired divided by the number of photoelectrons (NPe). The NPe is the number of photons which strike a SiPM pixel and generate an avalanche. Figure 18 shows two example events, one with 10000 thrown photons and one with 40000 thrown photons, from one of the b5 quadrant source simulations.

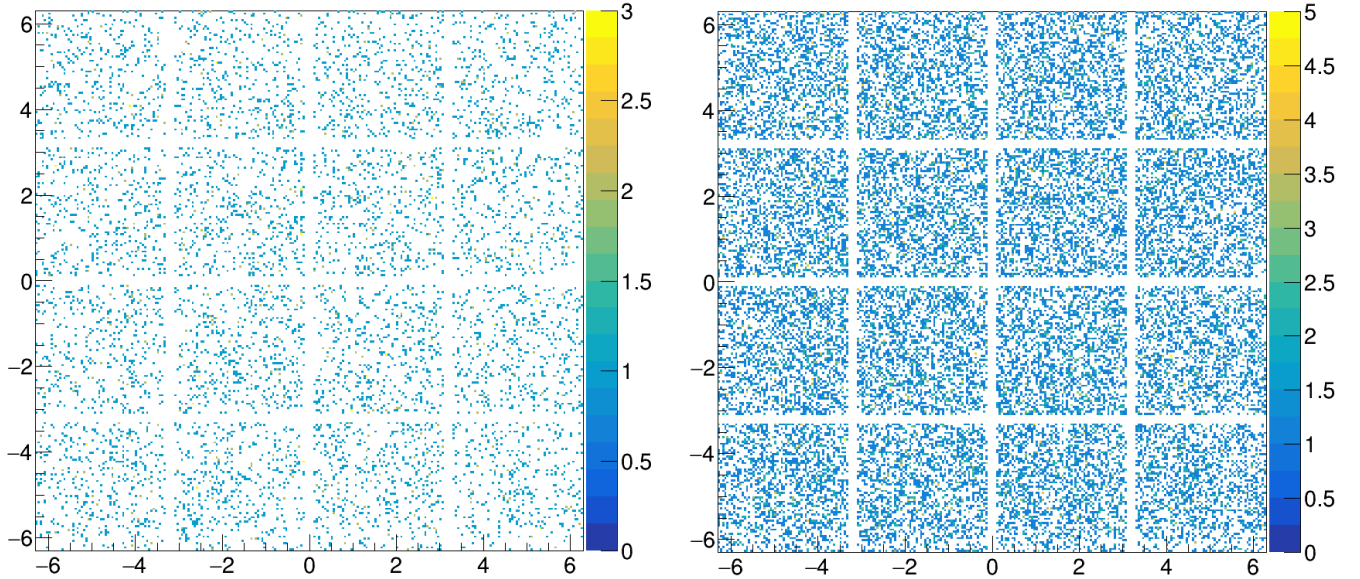


Figure 18: Detection plots, binned such that each histogram bin corresponds to one SiPM pixel $50\text{ }\mu\text{m}$ in pitch, from quadrant source simulations of a 50-mm-long, layer 5 nominal BIC light guide with a 1-mm-thick optical cookie and single-clad fibers showing events with 10000 thrown photons (left) 40000 thrown photons (right).

In order to accurately calculate the linearity of the SiPMs, we also need to simulate cross-talk, which occurs when an avalanche caused by a real photon in one SiPM pixel causes an avalanche in an adjacent pixel spontaneously. The chance of cross-talk quoted by Hamamatsu for the BIC SiPMs is about 7% [4], meaning that, on average, a photon striking a SiPM pixel of an S14 SiPM array which causes an avalanche has a 7% chance of incurring an additional avalanche in a neighbouring pixel, which is indistinguishable from the effects of a second photon hitting that neighbouring pixel. There are a number of mechanisms by which cross-talk can occur, but for this study, we simplify the simulation by assuming that each hit has a 7% chance of incurring a cross-talk hit and that each neighbouring pixel is equally likely to receive the cross-talk avalanche. In addition, each cross-talk hit also has a

7% chance of incurring another cross-talk hit. No special attention is paid to pixels on the edges or corners of SiPMs, meaning that hits on the edges or corners are equally likely to induce cross-talk hits even though they have less neighbouring pixels. Though unrealistic, this is the most simple way of implementing the quoted 7% cross-talk chance.

To generate one trial, we first choose a number of thrown photons, then we histogram those photons as described above to find the locations where photons strike active SiPM pixels. The number of entries in this histogram is the NPe, the number of photoelectrons caused by incident photons (in the case of 100% PDE, this is the same as the number of thrown photons which hit an active SiPM pixel). Then, for each bin with at least one hit, we generate a random number between 0 and 1. If that number is less than or equal to the cross-talk chance of 7%, we generate a cross-talk hit in a neighbouring pixel. The pixel is chosen again by generating a random number, this time between 0 and 0.8, with each range of 0.1 assigned to one of the eight surrounding pixels. If the chosen pixel lies outside of the SiPM array or lies in one of the dead crosses between SiPMs, a new random number is generated until we find a valid pixel. One count is then added to a second histogram with the same binning in the chosen pixel. This process is repeated until each occupied pixel in the original histogram has been checked for cross-talk hits. This second histogram, then, represents avalanches that would be caused by first-order cross-talk hits. If this second histogram is not empty, the procedure is repeated using this second histogram to generate second-order cross-talk hits. This is repeated up to fifth-order cross-talk hits. These histograms are then summed together to form the final detection plot for the event consisting of our chosen number of thrown photons. Finally, we count the number of non-empty bins in the histogram, which gives our number of pixels fired, and the linearity is calculated using the pixels fired and the NPe.




Figure 19: Average linearity and pixels fired vs. number of thrown photons for 50-mm-long, layer 5 nominal BIC light guides with a 1-mm-thick optical cookie assuming single-clad fibers.

For this study, we looked at trials with thrown photon numbers between 100 and 100000. We conducted 100 trials for each number of thrown photons, and each trial consists of an independent set of simulated photons. The NPe, pixels fired, and linearity are averaged over the 100 trials and the errors on these values are taken as the sample standard deviations. Figure 19 shows the average linearity and pixels fired vs. the number of thrown photons. Note that, because of the way the linearity is defined here, the linearity exceeds 1 at low number of thrown photons. This is because approximately 7% of the photons which generate avalanches also generate cross-talk hits, so the linearity only gets below 1 when the effects of saturation catch up with the effects of cross-talk. This feature could be removed by redefining the linearity to account for (or not include) cross-talk, but since the effects of cross-talk and non-linearity are not separable in real data, we'll keep this definition for now.

Figure 20 shows the same average linearity and pixels fired, but this time vs. the average NPe (incident photons which caused avalanches, referred to as 'detected photons'). Figure 21 shows the average linearity and pixels fired vs. the average number of incident photons, which is simply the average NPe divided by the PDE of the detectors (for this study, we used a PDE of 50%, which is typical for S14 SiPM arrays).




Figure 20: Average linearity and pixels fired vs. average NPe for 50-mm-long, layer 5 nominal BIC light guides with a 1-mm-thick optical cookie assuming single-clad fibers.




Figure 21: Average linearity and pixels fired vs. average number of incident photons for 50-mm-long, layer 5 nominal BIC light guides with a 1-mm-thick optical cookie assuming single-clad fibers and a 50% PDE.

The signal for an event out of a SiPM is proportional to the number of pixels fired, so, for a real data event, these types of simulations could be used to estimate the number of photons striking the SiPM and the effects of cross-talk and non-linearity for the purposes of correcting the signals for high-energy events.

References

- 303 [1] E.S. Smith. “Acceptance Studies for the Production Light Guides of the BCAL”. https://halldweb.jlab.org/DocDB/0017/001784/003/bcal_prod_guides.pdf. GlueX-doc-1784. 2011.
- 304
- 305 [2] Kuraray. “Plastic Scintillating Fibers”. https://www.kuraray.com/uploads/5a717515df6f5/PR0150_psf01.pdf. Product brochure. 2014.
- 306
- 307 [3] Hamamatsu. “MPPC (Multi-Pixel Photon Counter) arrays - S13361-3050 series”. https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s13361-3050_series_kapd1054e.pdf. 2024.
- 308
- 309
- 310 [4] Hamamatsu. “MPPC (Multi-Pixel Photon Counter) - S14160/S14161 series”. https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s14160_s14161_series_kapd1064e.pdf. 2024.
- 311
- 312
- 313 [5] Eljen Technology. “Silicone Rubber Optical Surface - EJ-560”. https://eljentechnology.com/images/products/data_sheets/EJ-560.pdf. 2021.
- 314
- 315 [6] Frank L. Goerner et al. “A comparison of five standard methods for evaluating image intensity uniformity in partially parallel imaging MRI”. In: *Medical Physics* 40.8 (2013), p. 082302. DOI: <https://doi.org/10.1118/1.4816306>.
- 316
- 317
- 318 [7] Hamamatsu. “MPPC (Multi-Pixel Photon Counter) - S13360 series”. https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s13360_series_kapd1052e.pdf. 2024.
- 319
- 320

A Trapezoidal Prism Parameter Adjustment

In Geant4, trapezoidal prisms are built using the ‘G4Trap’ class, which takes 11 arguments, where dimensions are taken as half-lengths:

0. dz (half-length)
1. θ
2. ϕ
3. dy1 (half-height of side 1)
4. dx1 (half-bottom of side 1)
5. dx2 (half-top of side 1)
6. α_1
7. dy2 (half-height of side 2)
8. dx3 (half-bottom of side 2)
9. dx4 (half-top of side 2)
10. α_2

$\alpha_{1(2)}$ is the angle between the y-axis and the line defined by the origin and the center of side 1 (2). In our simulations, θ , ϕ , α_1 , and α_2 are set to 0. The routine below, provided by Geant4 forum user ‘evc’ (<https://geant4-forum.web.cern.ch/t/g4trap-coplanarity-problem-with-geant4-version-10-4/1831>), takes the 11 G4Trap parameters as input, calculates the necessary adjustments to parameters 4 and 8, then outputs the parameters to be used to construct the G4Trap object.

```
void DetectorConstruction::TrapParamAdjustment(const double (&in)[11], double (&out)[11])
{
    // get input
    double dy1 = in[3], dx1 = in[4], dx2 = in[5], alp1 = in[6];
    double dy2 = in[7], dx3 = in[8], dx4 = in[9], alp2 = in[10];

    // adjust alpha
    alp1 = alp2 = std::atan((std::tan(alp1) + std::tan(alp2))/2.);

    // adjust dx
    double k1 = (dx1 - dx2)/(dy1 + dy1);
    double k2 = (dx3 - dx4)/(dy2 + dy2);
    double k = (k1 + k2)/2.;
    double m1 = (dx1 + dx2)/2.;
    double m2 = (dx3 + dx4)/2.;
    dx1 = m1 + k*dy1; dx2 = m1 - k*dy1;
    dx3 = m2 + k*dy2; dx4 = m2 - k*dy2;

    // set output
    out[0] = in[0]; out[1] = in[1]; out[2] = in[2];
    out[3] = in[3]; out[4] = dx1; out[5] = dx2; out[6] = alp1;
    out[7] = in[7]; out[8] = dx3; out[9] = dx4; out[10] = alp2;
}
```