

DETECTION OF HIGH ENERGY PARTICLES & JETS FROM
HIGGS DECAYS AT THE INTERNATIONAL LINEAR COLLIDER

by

BAYLEY BURKE

A THESIS

Presented to the Department of Physics and the Robert D. Clark
Honors College in partial fulfillment of the requirements for the
degree of Bachelor of Science

April 2023

Table of Contents

The Standard Model and the International Linear Collider	5
An Introduction to Monolithic Active Pixel Sensors	11
Pixel Hit Clustering	15
Methodology	17
Results & Analysis	17
Jet Finding & Analysis	25
Methodology	25
Results & Analysis	28
Summary & Outlook	36
References	38

List of Figures

Figure 1. The Standard Model	6
Figure 2. The International Linear Collider	8
Figure 3: The SiD Detector	11
Figure 4: Particle-Sub-Detector Interactions	12
Figure 5: TDR vs. MAPS Showers	13
Figure 6: Scaled Energy Resolution of MIPs vs. All Hits	15
Figure 7: Example Cluster	16
Figure 8: MIP Resolution Replication	18
Figure 9: Example Bugged Cluster	18
Figure 10: Distribution of Cluster Count Per Layer	20
Figure 11: Cluster Size & MIP Distribution	21
Figure 12: Cluster MIPs vs. Size	22
Figure 13: Cluster MIPs for Each Size	22
Figure 14: Average Cluster MIPs vs. Size Profile Plot	24
Figure 15: Example Two-Jet Event	27
Figure 16: $\log(R_{jj})$ Distribution	29
Figure 17: Jet Energy vs. $\log(R_{jj})$	30
Figure 18: Jet Scaled Energy Resolution	31
Figure 19: Jet Scaled Energy Resolution vs. $\log(R_{jj})$	32
Figure 20: $\log(R_{pp})$ vs. $\log(\text{Photon Scaled Energy Resolution})$	33
Figure 21: Scaled Photon Resolution vs. $\log(R_{jp})$	34
Figure 22: Scaled Photon Resolution vs. $\log(R_{jp})$ for $<3, 3-5, 5\text{GeV}$ Photons	35

The Standard Model and the International Linear Collider

The Standard Model (SM) is one of the most successful and accurate physical theories ever devised. In it, there exist just twenty-five fundamental particles (plus twelve antiparticles): twelve fermions (and their antiparticles) that make up the visible matter in the universe and thirteen bosons, which mediate the fundamental forces (except gravity, which the SM does not explain). The gluon (of which there are eight flavors, distinguished by their differing color quantum numbers) mediates the strong force, the force that holds atoms together. Two W bosons (W^+ and W^-) together with the Z boson mediate the weak force, which is responsible for phenomena like radioactive decay. Then there is the photon, which composes light and mediates the electromagnetic force. This is responsible for almost all day-to-day forces we experience, like the force keeping you from falling through the floor, and the forces between electrically charged materials and magnets. Finally, there is the Higgs boson, which is closely associated with the mechanism that is responsible for giving the other fundamental particles mass.

Standard Model of Elementary Particles




Figure 1: The Standard Model

https://upload.wikimedia.org/wikipedia/commons/thumb/0/00/Standard_Model_of_Elementary_Particles.svg/1200px-Standard_Model_of_Elementary_Particles.svg.png

A helpful reference diagram of the particles of the Standard Model. This chart is roughly analogous to the more familiar Periodic Table, but details components more fundamental than the elements.

Despite the extraordinary success of the SM, we understand that it is an incomplete description of nature due to its inability to explain phenomena like gravity or dark matter. There exist several potential additions to the model and parameters of it which are not known which may be able to account for these observations, but it remains to be seen if these describe our universe. The most promising path toward determining if this is the case lies with the Higgs boson. Theorized since the 1960's, it was not discovered until 2012 at the Large Hadron Collider (LHC), located at the European Organization for Nuclear Research (CERN). As such, a number of its properties are still not known to the remarkable degree that those of many of the other SM

particles are, leaving ample room for them to deviate from the predictions of the SM. If such deviation is found, this could indicate the presence of new processes, particles, or other phenomena (new physics). For example, one of the most famous problems vexing modern physicists is dark matter, a mysterious substance that does not interact with the electromagnetic force, and is therefore invisible, but makes up a staggering approximately 80% of the matter in the observable universe. If we find that the Higgs decays to invisible particles more often than the SM predicts, then this may mean that dark matter particles are being produced, which could then be further probed. Another example is the asymmetry between antimatter and normal matter. Antimatter can be created in particle colliders, but is largely absent in the universe as a whole. The SM predicts that matter and antimatter should have existed in equal quantities in the early universe, so there should exist some explanation for there being so staggeringly much more matter today, 13.8 billion years later. If we observe some asymmetry in the way that the Higgs decays, this might point the way to a modification to the SM which could account for the lack of antimatter. In order to probe these questions, a new particle collider is being planned to add to the discoveries made by the LHC: the International Linear Collider (ILC) [1] [2].




Figure 2: The International Linear Collider [2]

A diagram of the ILC from the Technical Design Report (TDR). The full collider is 31 km in length, to potentially be upgraded to higher energy and 45 km in length. For perspective, each dash in the dashed line below the collider is a full soccer field. The blue lines show the path of the electron beam, and the green path shows that of the positron (anti-electron) beam, which is created from the electron beam in the gray cylinder on the left.

The ILC differs from the LHC in a few critical ways. First, it is a linear collider rather than circular. This choice was made to avoid the effects of synchrotron radiation, a phenomenon which causes accelerating charged particles to radiate away their energy. To keep charged particles moving in a circle, they must constantly be accelerated to change the direction they are moving, so circular colliders must fight against this effect on top of ramping up the particles' energies. Synchrotron radiation also saps far more energy from low-mass particles than it does from high-mass particles, so heavier particles feel much less of its effects. This is precisely the reason that synchrotron radiation is manageable in the case of the LHC; it collides protons, which are relatively heavy. The ILC instead collides electrons and positrons (the antiparticle of electrons), which are much lighter than protons and thus much more susceptible to synchrotron radiation. This makes a linear collider much more attractive in order to avoid the much stronger

energy drain. The main trade-off is that the beams only get one chance to collide rather than being able to make multiple laps around the collider. This, combined with electrons and positrons having much smaller cross-sections, means that the beams must be squeezed down to an extremely small size in order to ensure that the particles interact. This is a significant challenge, but ultimately worth it to reap the benefits of electron-positron collisions.

As opposed to electrons and positrons, protons are not fundamental particles, instead being composed of three quarks bound together using gluons by the strong force. When colliding at the energies produced in the LHC, they break apart and are able to form new particles within the debris. However, because there are six quarks plus the gluons that bind them together going into the collision, that debris can become very messy, producing large numbers of particle tracks which need to be sifted through. Electron-positron colliders reduce this problem drastically by using matter and antimatter; upon colliding, instead of breaking apart, electrons and positrons undergo a process called annihilation, which simply frees up their energy to cleanly create new particles, leaving no debris.

After the electron and positron collide, the energy gained from the accelerator may create a Z boson and a Higgs boson. This process is called Higgsstrahlung and will be the primary production method of Higgs bosons at the first phase of the ILC (with center-of-mass energy of 250GeV). The clean, controlled nature of this process makes the ILC an ideal environment for producing large numbers of Higgs bosons with low background noise, acting as a ‘Higgs factory’ [3]. These particles are incredibly short lived, almost immediately (on the order of 10^{-22} s [4][5]) decaying into other particles which fly into the detectors, often in concentrated streams called jets.

By detecting these precisely, we can trace their paths backward and reconstruct the original Z and Higgs they came from, allowing us to probe the Higgs' properties. This detection is planned to be done by two main detectors, called ILD and SiD, the latter of which is the focus of this work.

An Introduction to Monolithic Active Pixel Sensors

The SiD detector [6], depicted in Figure 3, consists of a number of sub-detectors, each of which contribute to the full detector's ability to measure particles. My research focuses on SiD's electromagnetic calorimeter (ECal), a device which measures the energy of incoming electrons, positrons, and photons as precisely as possible (a diagram showing where different particles tend to interact is shown in Figure 4). The ECal is constructed of thirty alternating tungsten and silicon layers, with the last ten tungsten layers being twice as thick as the first twenty to offset cost. Its shape is that of a dodecagonal prism, with the long section with a dodecagonal cross-section referred to as the "barrel" and the remaining two faces referred to as the "endcaps." When a photon, electron, or positron, enters the detector, it hits the densely-packed (high-Z) atoms in the tungsten, creating a shower of additional electrons and photons that develops through the layers, depositing energy in the silicon which can be measured.




Figure 3: The SiD Detector [7]

A 3D model of the SiD detector. Different subsystems are modeled with different colors. The red subsystem in the center is the Tracker, the pink subsystem is the Hadronic Calorimeter (HCal), and the green subsystem sandwiched between them is the ECal.




Figure 4: Particle-Sub-Detector Interactions

<https://indico.cern.ch/event/471037/contributions/1980825/attachments/1225875/1794509/ATLAS-Introduction.pdf>

A diagram showing which parts of a detector several different particles generally interact with. Notably, this is not a diagram of SiD, but of the ATLAS detector at the LHC, though these interactions are the same for SiD and the structure is highly similar.

In the baseline design outlined in the ILC Technical Design Report (TDR) [2][6], the silicon is divided into 13 mm^2 hexagonal analog pixels capable of measuring the amount of energy deposited in each. We can achieve good performance with this design, but incoming photons, electrons, and positrons may enter the detector very close together, and the lack of granularity can make it difficult to disentangle them. This limits the spatial precision of the reconstruction process, and thus the detector's ability to make precision measurements of the Higgs.

To address this, a novel design has been proposed, based on monolithic active pixel sensors (MAPS). In this new design, the silicon is divided into a grid of $25\mu\text{m} \times 100\mu\text{m}$

rectangular digital pixels which each record a hit if a particle passes through them, but not the particle’s energy. Fortunately, by carefully counting the number of hits, we can still extract the energy of an incoming particle despite not being able to measure it as directly. Additionally, because the pixels are much smaller, we gain a much higher resolution measurement of the location of each particle in the detector. This is illustrated in Figure 5 below, which depicts displays from each detector for the same event in which a neutral pion decays to two photons, each of which enter the detector and create the showers shown. MAPS provides a much finer spatial resolution, allowing us to pinpoint the center of each shower more precisely and separate the two showers from each other more easily.

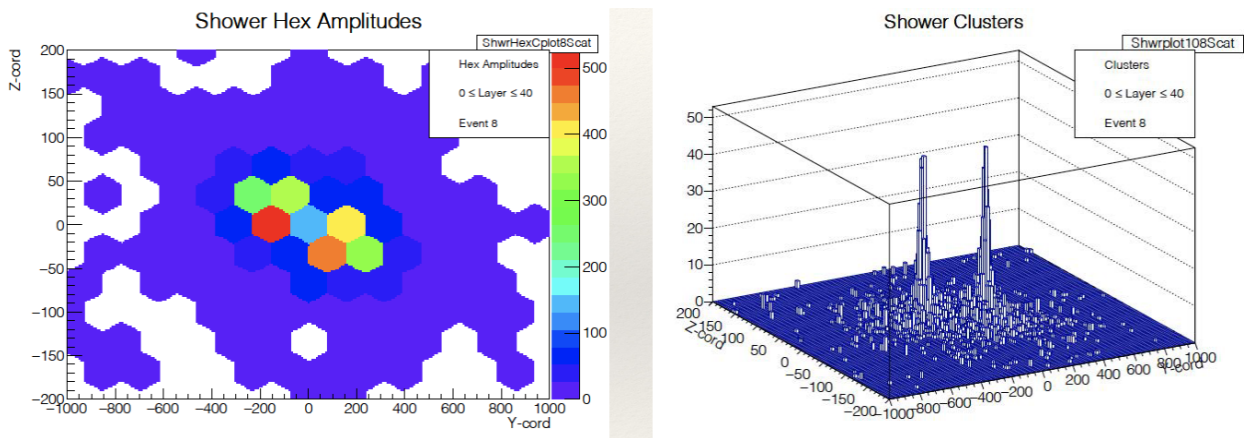


Figure 5: TDR vs. MAPS Showers [8]

Displays of the TDR detector (left) and in MAPS (right) for an event in which a neutral pion decays to two photons.

This measurement may then be paired with data from the other components of the detector to identify the particle and reconstruct the Z boson, if present, from all its decay products. In the event that the reconstruction does yield a Z-boson, we can use it, along with our knowledge of the total energy in the system, to find the Higgs. We can examine the reconstructed Z boson and find its “recoil mass:” the effective mass of particles other than the Z boson which are present in the event. If this recoil mass matches that of the Higgs, then we know that there is

a Higgs in the event before examining its decay products. This is useful because if the Higgs can decay into invisible dark matter particles, we would never be able to observe them directly. Using recoil mass, we can still identify the Higgs in these events and determine how often it decays to invisible particles. The Standard Model predicts this will happen only about 0.1% of the time [5], so if we find this to be happening at an elevated rate, that would be an indication of new physics.

This reconstruction process consists of a number of steps, first gathering the raw hits into clusters, then reconstructing particles from these clusters, next finding jets of particles, then from these jets finally reconstructing the Z, and from it the Higgs. Here, I will focus first on the performance of the initial clustering stage in MAPS, before exploring how MAPS can contribute to improving the jet reconstruction stage.

Pixel Hit Clustering

The loss of a direct analog energy readout from each pixel is the most significant disadvantage of MAPS compared to the TDR baseline detector design. At first glance, it seems that this should lead to a reduction in the detector's energy resolution (how precisely it can measure energy), but with careful analysis this may be recovered.

We can measure the energy of a particle by counting the number of Minimum Ionizing Particles (MIPs), which are charged particles above a minimum energy threshold (here 0.1 MeV). By counting how many are present in each layer of the detector, we can track the development of the shower produced by an incoming particle, and from this extract its energy. However, MAPS cannot determine which hits (pixels which have been activated by the impact of some particle) were caused by a MIP and which were caused by some other particle, nor can it tell if one MIP activates multiple pixels. Thus, we must use approximate methods to try to count MIPs. One way to do this is to simply count up all the hits, though it is clear that this leaves significant room for improvement, as shown in Figure 6 [8].

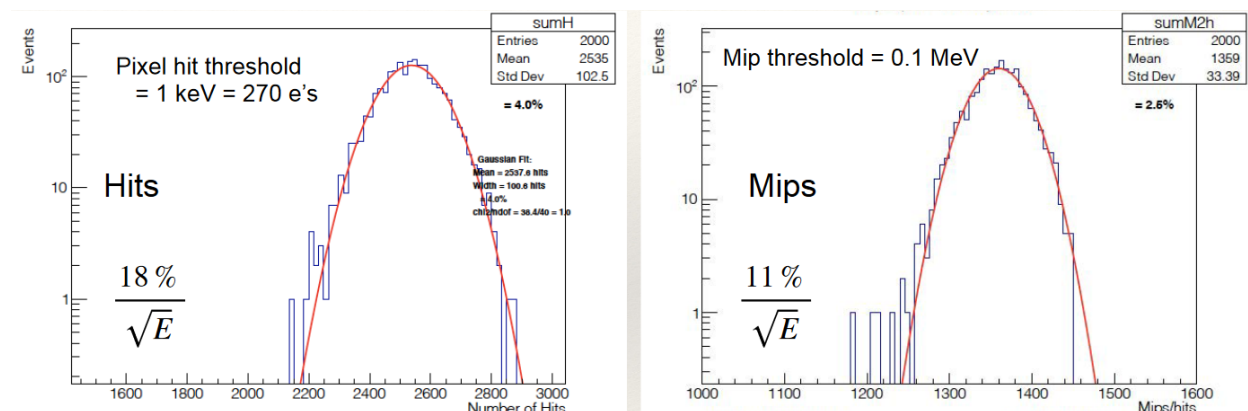


Figure 6: Scaled Energy Resolution of MIPs vs. All Hits [8]

Two histograms, the left showing the scaled energy resolution measurement from examining all hits, and the right showing the measurement of scaled energy resolution from examining MIPs.

Nonetheless, this provides a good starting point to which improvements may be made. In this work, the hits are grouped into clusters. MIPs often enter the silicon at an angle, activating one or more adjacent pixels by entering one pixel and traveling through the silicon to activate the others. Figure 7 is a clear example of this phenomenon, and of how clustering can help to resolve it. The cluster in the center of the display is created by a single MIP, entering the silicon layer in the black pixel at a shallow angle and traveling through the red pixels, activating all of them. Were we to simply count all the hits in this image, this particle would essentially be counted eleven times, since it activated eleven pixels. Clustering solves this issue, because once these hits are clustered together, the MIP is only counted once. This method has been shown to recover the energy resolution of the TDR design and even the potential to exceed it [7][11], which is expanded upon in this work.

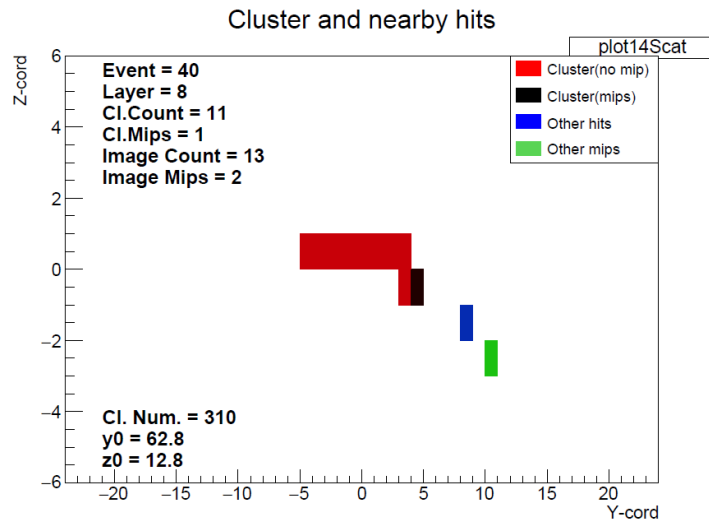


Figure 7: Example Cluster

An example cluster formed by merging adjacent hits. The cluster itself consists of a single hit where the MIP that created it entered the detector, in black, and ten adjacent hits which were created by the MIP traveling within the silicon after entering (not given MIP designation because they were not caused by a MIP entering the silicon). Nearby hits which are not part of this cluster are also shown in green (MIP) and blue (non-MIP), each of which are in their own, separate, single-hit clusters.

Methodology

Given that the ILC is still in its research and development phase, there exists no complete physical detector to provide data to form clusters with. As such, this work is based entirely on Monte Carlo simulations built using Geant4[9]. These allow for a detector to be built, tested, and tweaked virtually rather than using physical prototypes. Furthermore, because we are in control of the simulation, we may extract not only the data from the virtual detector, but the originally generated particles as well. This gives us two lists of data: the Monte Carlo Truth (MCT) data, a list of all particles generated by the simulation and all their properties; and the hit data, the data output by the detector once the generated particles are sent into it. Having access to the MCT data allows us to run Python analyses on the hit data using pyROOT[10] and compare the results to “truth,” giving diagnostics on how accurate these analyses are.

The clustering portion of this work is largely based on prior efforts by Professor James Brau, who wrote the framework for this analysis. To expand on this, I built my own clustering algorithm to run alongside the old one, comparing their outputs and revealing bugs in both.

Results & Analysis

Using these methods, previous results were able to be replicated, including the measurement of scaled energy resolution from Figure 6, shown below in Figure 8. Additionally, the robustness of the clustering algorithm was improved. These enhancements mostly consisted of removal of edge-case bugs, such as those shown in Figure 9.

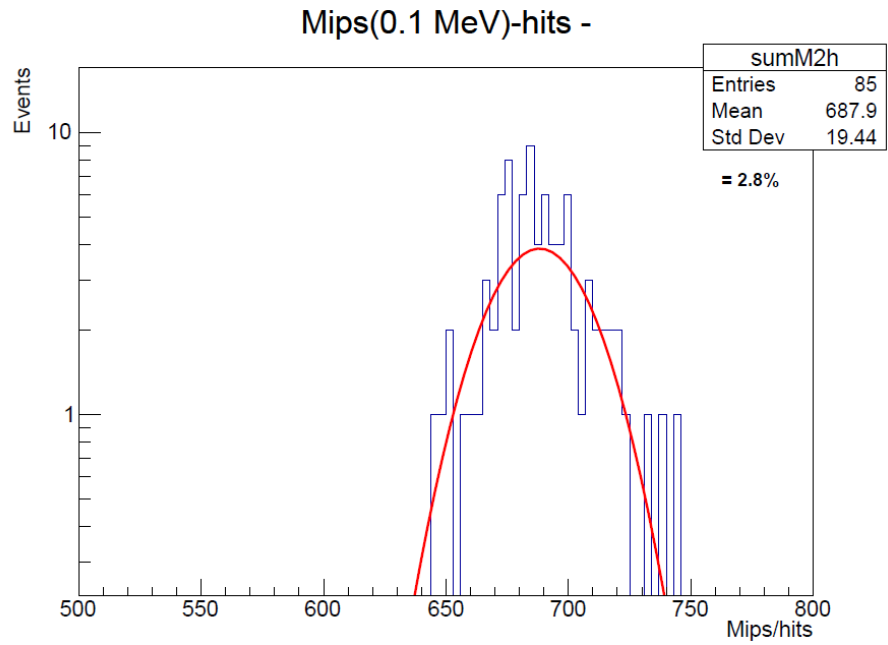


Figure 8: MIP Resolution Replication

A replication of the MIP energy resolution measurement shown in Figure 6.

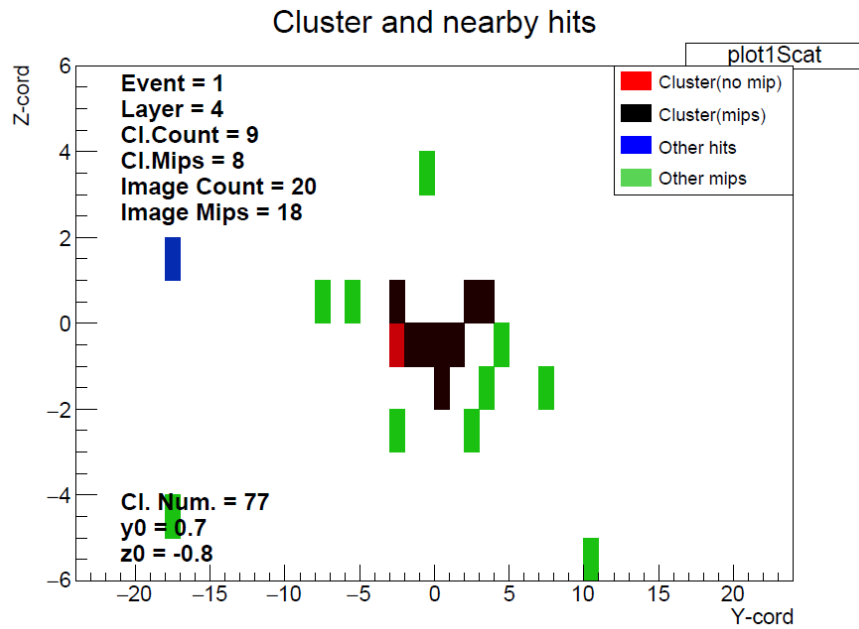


Figure 9: Example Bugged Cluster

An example of a bugged cluster which has now been remedied. Three hits are not successfully joined to the cluster.

With this in hand, we are able to examine a number of different parameters to characterize the detector and the clusters we create from its measurements, with the goal of finding ways to further improve our ability to count MIPs. First, we examine the number of clusters per layer in Figure 10, which shows us how the showers generated by incoming particles tend to behave as they develop through the calorimeter. This distribution is given for clusters with at least one hit caused by a MIP entering the silicon (MIP clusters, in red) and clusters with no such hits (non-MIP clusters, in blue).

We can see that the shape of each distribution is nearly identical, and the peak of each distribution is around layer ten, the midpoint of the thin layers. This is dependent on the energies of the particles entering the detector; low energy particles have less energy to impart to the shower they generate, so it does not penetrate as far into the detector as a shower from a high-energy particle, which can provide the energy to penetrate much further. Thus, if higher-energy particles had MIP clusters represent a larger portion of their produced clusters, we would expect the distribution of MIP clusters to be shifted deeper into the calorimeter than that of non-MIP clusters. Given that the distributions here are identical, we can safely say that this is not the case, and rule out any simple methods of differentiating MIP clusters from non-MIP clusters based on the layer they lie in.

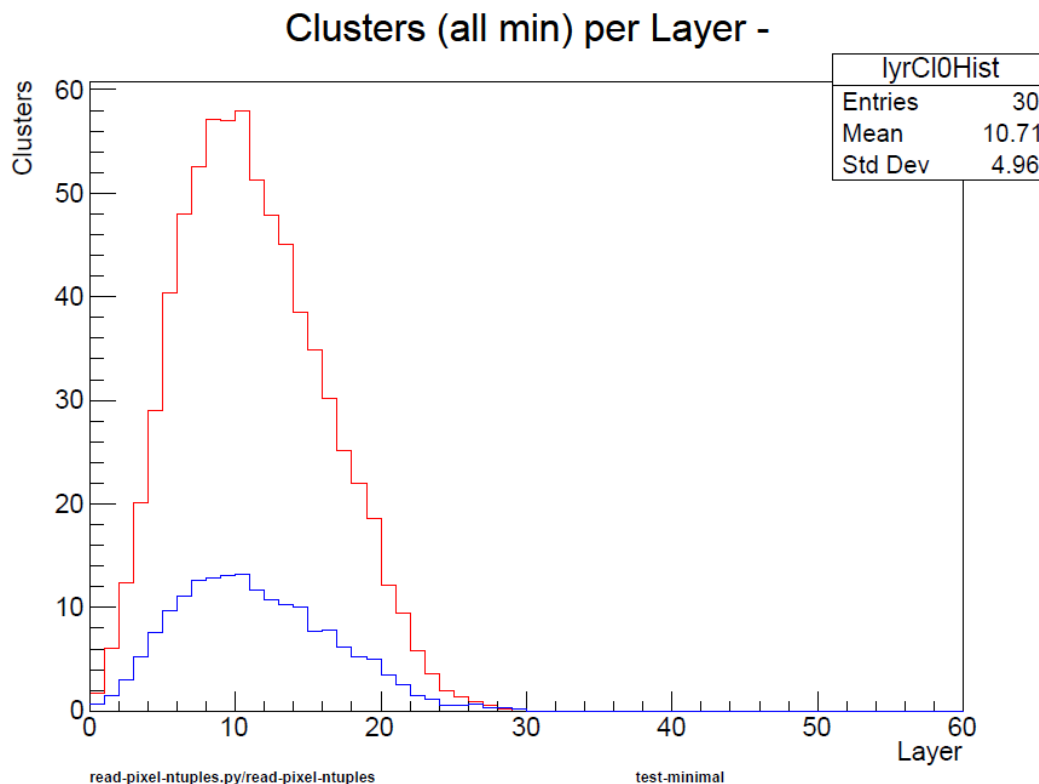


Figure 10: Distribution of Cluster Count Per Layer

A distribution of the number of clusters found in each layer of the calorimeter. The final ten layers are twice as thick as the first twenty. Clusters containing at least one hit caused by a MIP entering the silicon are in red, and clusters with no such hits are in blue.

We may also examine the properties of each cluster, such as size and the number MIPs contributing to creating each cluster, and how these relate. It is immediately apparent from Figure 11 that events are dominated by clusters containing only one or two hits, with the remainder making up less than a percent of the total. The number of MIPs follows a similar trend, though this is an incomplete picture without information about how these relate to cluster size, as the two are inherently linked. However, we can at least still see that the largest clusters tend not to be mostly made up of MIPs, since there are no clusters with >15 MIPs, and the distribution is overall skewed left.

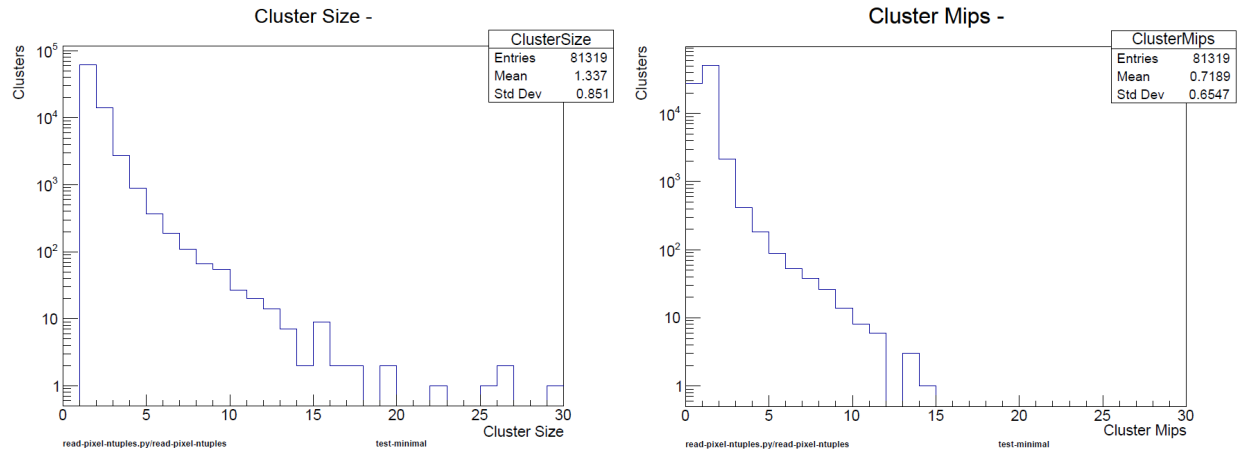


Figure 11: Cluster Size & MIP Distribution

Distributions of the number of clusters of each size (left) and the number of MIPs in each cluster (right).

This is further clarified by Figures 12 and 13; most clusters consist of a single hit made by a MIP, followed closely by clusters formed from a single hit not made by a MIP, then by two-hit clusters with one, zero, and two MIPs. As cluster size (the number of hits a cluster contains) continues to increase from that point, the number of entries becomes relatively insignificant, but clusters with zero or one MIP are most common. Figure 13 makes clear how dominated the analysis is by low multiplicity clusters.

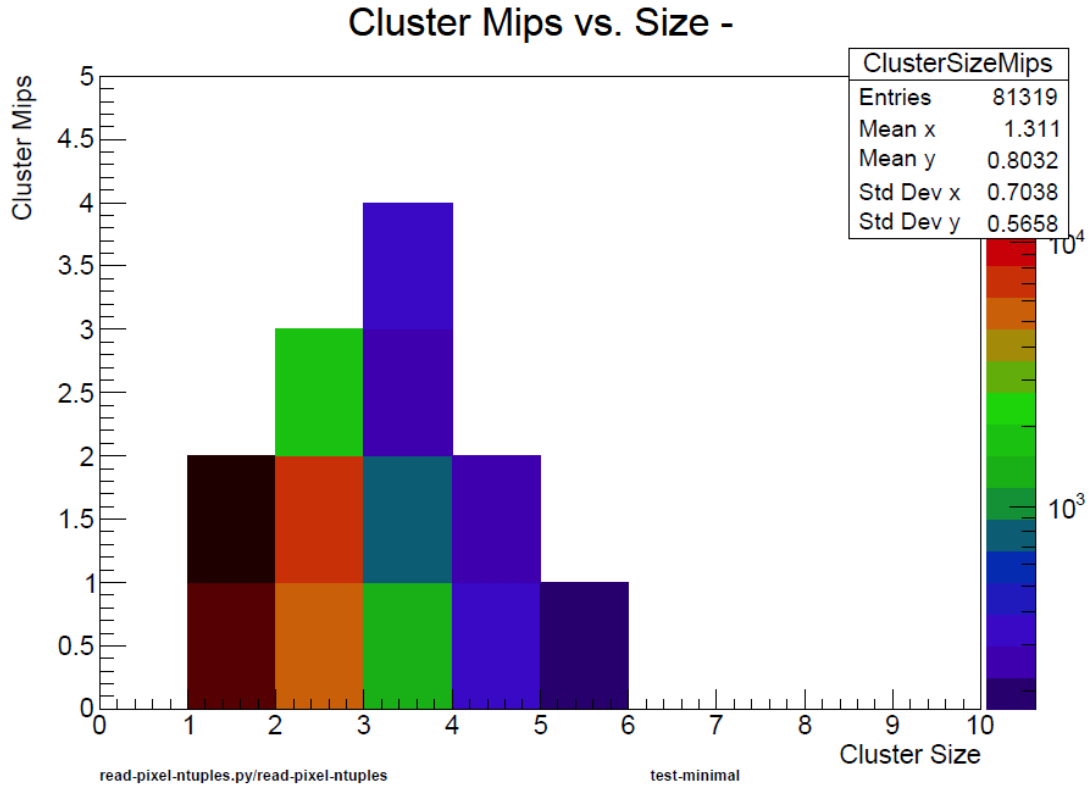


Figure 12: Cluster MIPs vs. Size

A 2D-Histogram of the number of MIPs in each cluster vs. the size of each cluster. A lower bound cutoff is imposed, not plotting a bin if the number of entries in it falls below 10^2 .

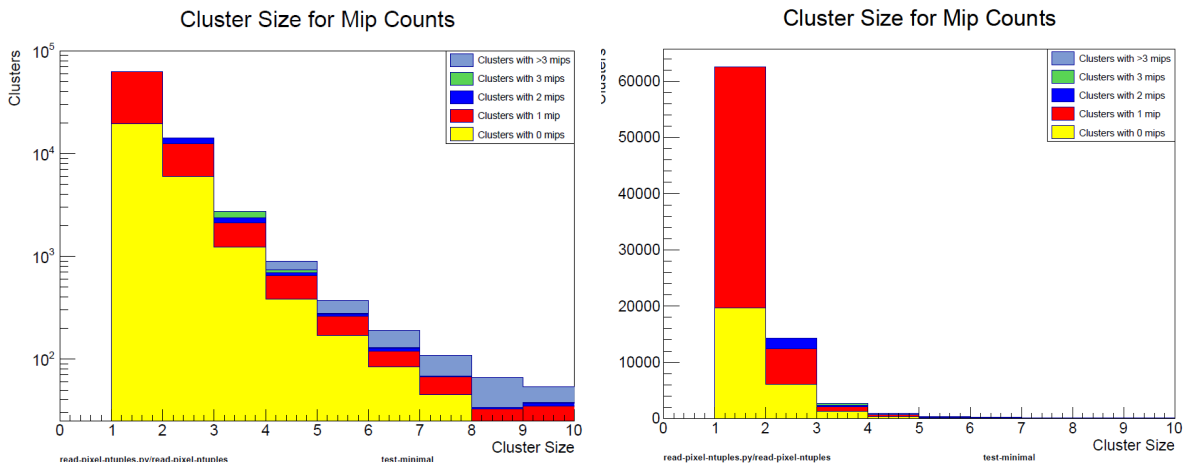


Figure 13: Cluster MIPs for Each Size

Stacked histograms showing the number of clusters containing different numbers of MIPs (in different colors) up to size 10. Logarithmic scale is on the left; linear scale is on the right.

Figure 14 presents the average number of MIPs in each size cluster, which illustrates how well counting clusters approximates counting MIPs. This is plotted in a profile plot, where for each bin a horizontal line is plotted to denote the mean of the entries in the bin, and a vertical line is plotted denoting their standard deviation (RMS). For example, here size nine clusters have approximately three MIPs on average, with an RMS of approximately 0.5. Bins past size ten begin to lose significance due to low statistics. Size one and two clusters have the same average number of MIPs per cluster at approximately 0.7. This means that purely counting clusters overcounts MIPs for these sizes. Fortunately, this can easily be accounted for by weighting the count such that size one and two clusters are only counted as approximately 0.7 of a cluster instead of 1. We can also make similar adjustments to the other sizes of cluster to get a full set of weights for each cluster size with sufficient statistics. This technique does require knowledge from the MCT which the real detector would not have access to, but this is not a problem. We can calibrate these values based on the simulation first, then use those to weight the real data.

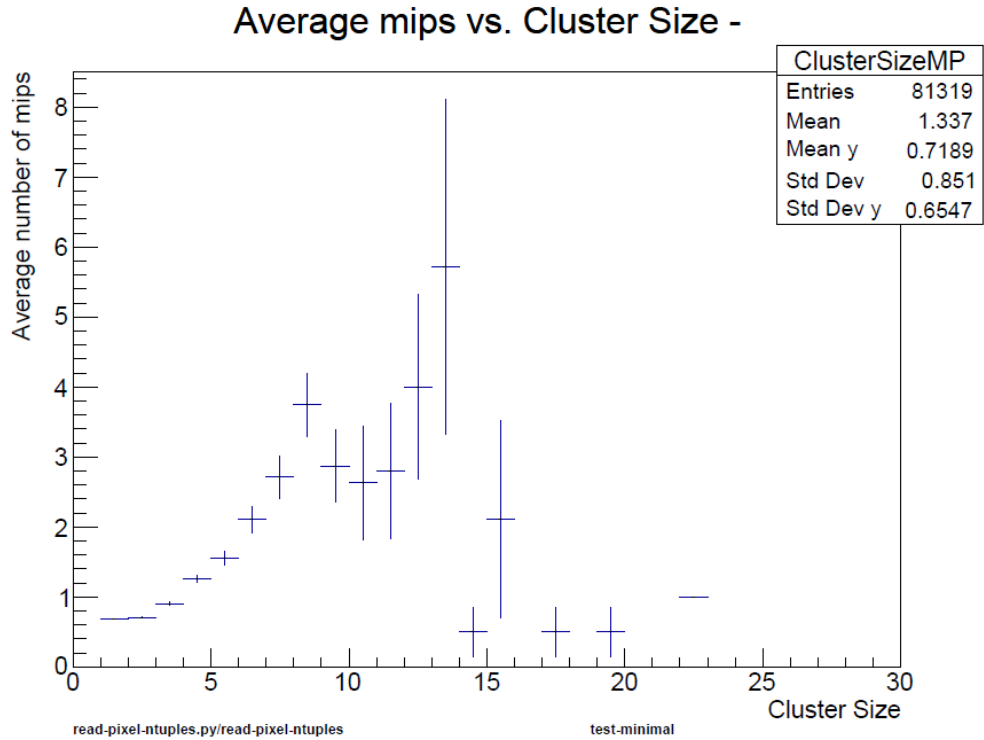


Figure 14: Average Cluster MIPs vs. Size Profile Plot

A profile plot of the average number of MIPs in each cluster for each cluster size. The horizontal line on each bin marks the mean of the entries in that bin, and the vertical lines extending from it mark the standard deviation (RMS) of the distribution.

These data show that the results of prior work have been replicated and give ample diagnostic data on the behavior of clusters. We can see that clustering provides a better method for counting MIPs than simply counting hits, and that it is able to account for important effects like size two clusters having the same number of MIPs as size 1 clusters. We can also see that the analysis is dominated by low multiplicity clusters. This limits opportunities for clever treatment of more complex cases to make a large impact on the overall results, but does simplify the analysis. With this understood, we may move to a different part of the reconstruction process in order to understand where this approach might stand to improve it.

Jet Finding & Analysis

From clustering, we move up the chain of steps in the event reconstruction process to jet finding. This takes the collection of particles identified in the detector and groups them into jets, which are then used to reconstruct the decay products of the Z . In this work, we examine Higgsstrahlung events with a Z and a Higgs in the final state, with the Higgs decaying invisibly. These events were generated by Chris Potter, with further details in the next subsection.

The precision of this process depends directly on how well we can determine the properties of jets. In particular, the energy resolution of each jet is critical. If the jet energy resolution is poor, then so too will be the resolution of the reconstructed Z . Here, MAPS has the opportunity to leverage its high granularity to improve over the TDR design. Jets are composed of many particles, with the density of particles increasing toward the center of the jet. If this density becomes sufficiently high, the resolution of the particles in this inner region begins to degrade, since they become packed tightly enough together that the detector begins having difficulty distinguishing them. Thus, it merges, misidentifies, or misattributes energy to them, introducing inaccuracy in the reconstruction of the particles and ultimately the jet. The smaller pixel size of MAPS allows the particles to be much closer together before this happens, which should improve their reconstruction.

Methodology

As in clustering, much of the process of jet analysis consists of comparing the data output by the simulated detector to the data generated by the simulation itself. There are two main differences, however. First, the detector being simulated is not MAPS, as the number of pixels would make this computationally untenable, but the TDR detector. Instead of directly examining MAPS, we can instead examine in detail the TDR design, identifying exactly where and how its

performance begins to degrade, and in which of these areas MAPS stands to make an improvement. Second, the most basic unit we will be examining is not single hits, but particles. For the data coming directly from the simulation, again called Monte Carlo truth (MCT), this is easy to retrieve, but this takes some work for the detector data. Fortunately, this work has already been done for us, encapsulated in the PandoraPFA algorithm [12]. This algorithm pieces together the information in each sub-detector, using clusters and charged tracks to reconstruct particles as a list of Particle Flow Output (PFO) data (so named because PandoraPFA is a “particle flow” algorithm). These data are a list of reconstructed particles analogous to the MCT data and represent the most complete picture we have of what the detector can see. As such, we can determine how well the detector is able to measure different particles by directly comparing the PFO and MCT data. For instance, neutrinos may be created in events, but these interact with other matter extremely rarely (in order to have just a 50% chance of stopping a neutrino, one would need to place a light year of solid lead in front of it). As such, the MCT contains all of the information for these particles, but the PFO never contains them, since they are nearly invisible to the detector.

The process of jet finding is similarly performed using PyJet, a Python-based implementation of FastJet [13], using an anti-kt clustering algorithm with $R=1$. It is important to note that FastJet is primarily intended for use in hadron colliders like the LHC, which perform their clustering in a different coordinate system than in electron-positron colliders, so the jets formed by it are slightly distorted, though not in any way that would compromise the parameters we examine. Jets are formed inclusively (i.e. using the “inclusive_jets” function in PyJet), then a cut is made requiring jet energy $>5\text{GeV}$ to ensure we are only examining true jets. Additionally, we examine only two-jet events where both jets have a combined energy of $>100\text{GeV}$ and lie

well within the barrel ($|\cos(\theta)| > 0.5$, θ being the angle from the beam axis) (an example event is shown in Figure 15). These requirements are assessed based on the MCT data alone. This removes the vast majority of events, but ensures that we have removed all possible factors which could impact the measurement aside from the performance of the detector itself. The barrel requirement ensures that there are no geometric effects from the corners where the barrel and endcaps meet; the two-jet requirement simplifies the underlying physics; the combined energy requirement ensures that the jets are well-behaved and constitute the majority of the energy of the Z (which has energy approximately equal to 110GeV). Finally, only events where the Z decays to a quark-antiquark pair ($Z \rightarrow qq$) are examined in order to select events with less complex physics.

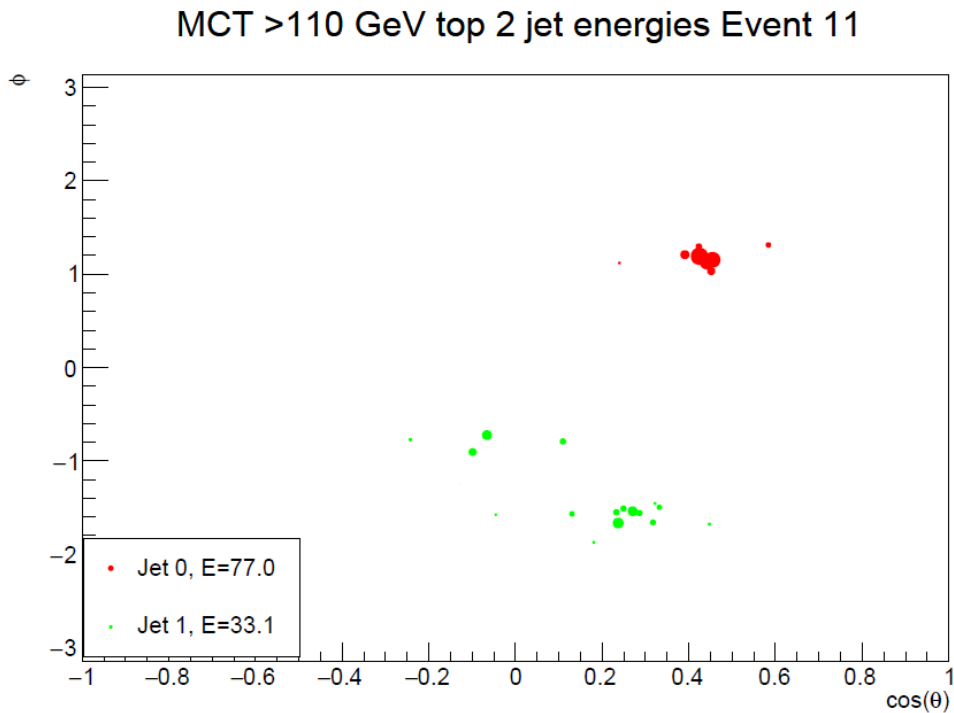


Figure 15: Example Two-Jet Event

A display of an example event with the above cuts applied. Each dot is a particle with color denoting jet membership and size proportional to energy.

PyJet is used on both the PFO and MCT datasets to produce two corresponding sets of jets, which can then be examined and compared in concert with the original lists of particles, all of which is again done using pyROOT[10] and Python. To compare these sets of data, we must also determine a method for matching MCT and PFO objects (particles or jets) to each other. For this purpose, we use a parameter $R_{ab} = 1 - \cos(\theta_{ab}) = 1 - (\mathbf{p}_{xa}\mathbf{p}_{xb} + \mathbf{p}_{ya}\mathbf{p}_{yb} + \mathbf{p}_{za}\mathbf{p}_{zb})/(\mathbf{p}_a\mathbf{p}_b)$, where a, b indicate the objects being compared, p_{ka} is the k-component of the momentum of object a, and p_a is the total momentum of object a. For example, R_{jj} compares two jets, whereas R_{jp} compares a particle to the center of the jet that contains it. R_{ab} ranges from 0 to 2, with small R_{ab} indicating the two objects are close together (meaning that the momenta are in very similar directions), and large R_{ab} indicating they are far apart.

Results & Analysis

We can use R_{jj} to examine how well the detector is able to reconstruct jets by comparing MCT and PFO jets. We can take each MCT jet in descending order of energy (so that the more dominant jet is favored in tiebreaker cases) and match it to the nearest (lowest R_{jj}) PFO jet. We can double-check that this does not match to the wrong jet by examining the distribution of R_{jj} , given by Figure 16, and observing that it has a single peak. If this did not select the correct jet, then there should be two peaks: one at the average R_{jj} between an MCT jet and its PFO counterpart and one at the average R_{jj} between the two jets. We can also use this to characterize the performance of the detector. The distribution peaks at roughly $R_{jj} = 10^{-4} = .8$ degrees, indicating good spatial reconstruction of the PFO.

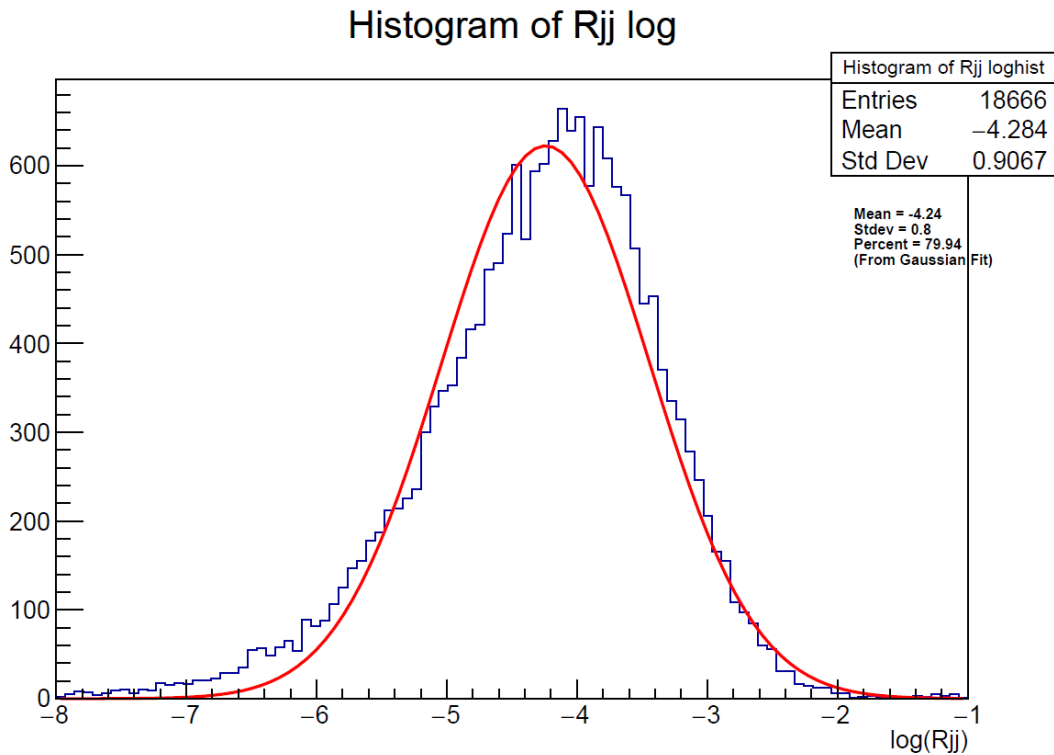


Figure 16: $\log(R_{ij})$ Distribution

A histogram of R_{ij} between each MCT jet and its matched PFO jet. The distribution is slightly non-gaussian with a left-sided tail.

We can see from Figure 17 how this matched R_{ij} varies with jet energy. It shows that higher energy MCT jets have their PFO partner reconstructed closer to them than low energy jets. This makes sense, since we know that the detector can reconstruct high energy jets better.

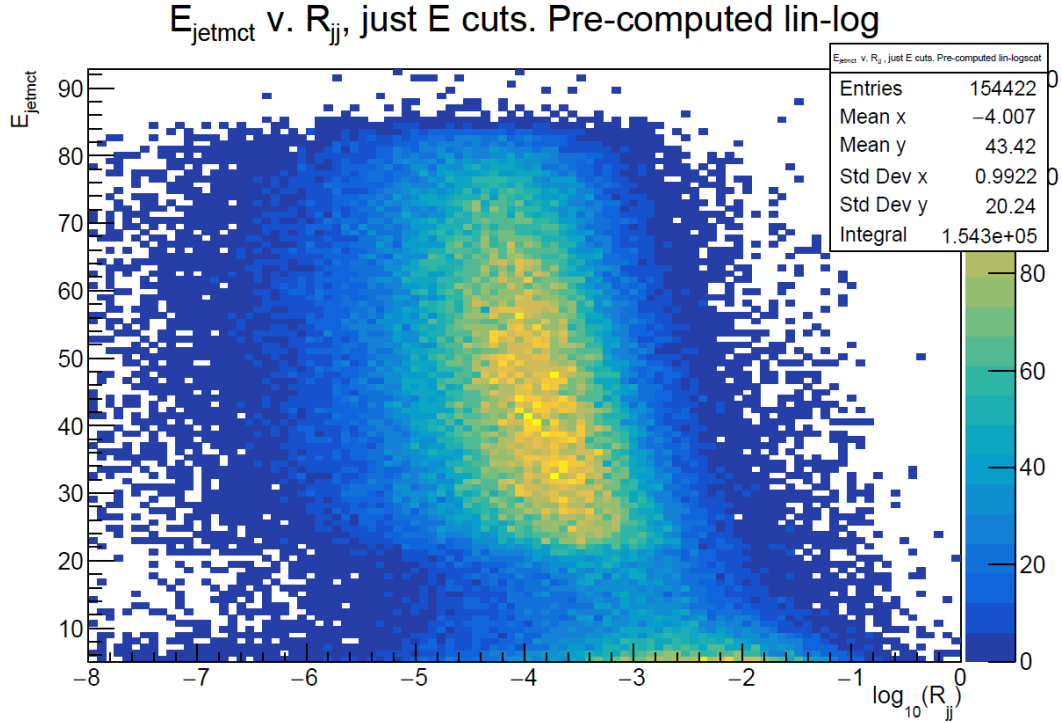


Figure 17: Jet Energy vs. $\log(R_{\text{jj}})$

A 2D histogram of the energy of each MCT jet vs. $\log(R_{\text{jj}})$ between it and the PFO jet matched to it.

More important than R_{jj} is energy resolution, which tells us how precisely the detector is able to measure the energy of a jet. There are two ways to examine this: scaled, or unscaled. Unscaled energy resolution is defined as the standard deviation of the distribution of $(E_{\text{MCT}} - E_{\text{PFO}})/E_{\text{MCT}}$. This examines how well the detector can measure the jet's energy by taking the difference between the actual and measured energies, then scaling it by the energy to account for how large this change is in relation to the jet itself (an energy difference of 5GeV is much worse for a 20GeV jet than an 80GeV jet). However, because we know that the detector can measure higher energy jets better than low energy jets, this still has a dependence on jet energy. This is removed, at least in principle, by examining the scaled resolution, defined as the standard deviation of the distribution of $(E_{\text{MCT}} - E_{\text{PFO}})/(E_{\text{MCT}})^{1/2}$, which will generally be preferred here.

In Figure 18, we see an overall scaled jet energy resolution of approximately $35\%/(E_{MCT})^{1/2}$, which is in-line with what we should expect.

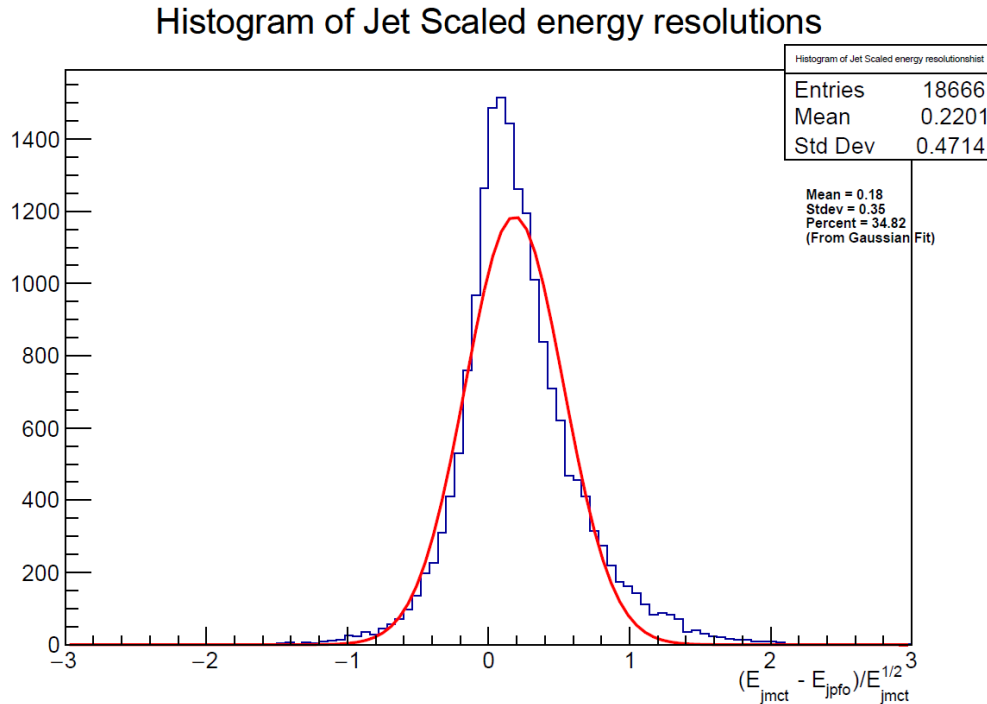


Figure 18: Jet Scaled Energy Resolution

A histogram of the scaled energy resolution of all jets. There is a slight bias in the positive direction, indicating that the PFO tends to underestimate jet energy.

Additionally, we can see how this scales with R_{jj} in Figure 19 (wherein each jet is placed in a bin corresponding to a range of R_{jj} values as in a histogram, then the mean and standard deviation of the scaled energy resolution of the jets in each bin is plotted as a horizontal line and vertical lines respectively). As expected, when R_{jj} is small, the scaled resolution is good, and when R_{jj} is large, the scaled resolution is poor. When the detector poorly reconstructs a jet's position in space, it also poorly reconstructs its energy, and vice versa. Notably, there is also a slight positive bias in the means of the bins, indicating that the PFO tends to systematically underestimate the jet energy; similar behavior is observed in Figures 21 and 22.

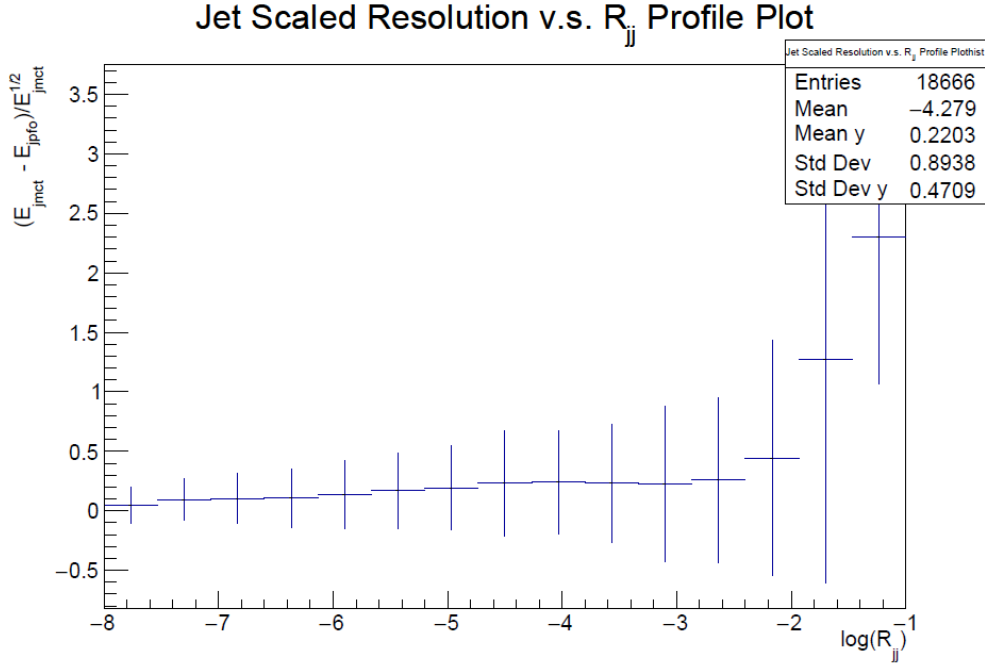


Figure 19: Jet Scaled Energy Resolution vs. $\log(R_{jj})$

A profile plot of the scaled energy resolution of jets for each R_{jj} . The horizontal lines indicate the mean scaled energy resolution of each bin, and the vertical lines denote the standard deviation (RMS) of the distribution. The two rightmost bins have little statistical significance due to low entries.

Now that we understand the behavior of the jets, we can examine their structure. Within a given jet, we can examine R_{jp} , which is calculated in the same way as R_{jj} , but instead of comparing two jets it compares the central axis of the jet to one of its constituent particles. We also need to determine the energy resolution of each particle, which is done similarly to comparing jets. We define the scaled energy resolution in the same way as before, using the standard deviation of the distribution of $(E_{MCT} - E_{PFO})/(E_{MCT})^{1/2}$, but now for all particles in a jet instead of all jets in an event. For each MCT particle, we examine R_{pp} (also defined in the same way as R_{jj} but comparing particles instead of jets) and match it to the closest PFO particle. R_{pp} is also required to be $> 10^{-5}$ in order to select the best quality matches. This selection is based on Figure 20, which shows how scaled energy resolution varies with R_{pp} for photons. Then, once

these matched pairs are formed, we can calculate the energy resolution and R_{jp} for each particle and see how they relate, as shown in Figure 21, again for photons. The decision to examine photons in Figures 20-22 is made because other particles' measurements rely on information from subdetectors other than the Ecal. Thus, we do not wish to examine them in order to measure the performance of the Ecal in relative isolation.

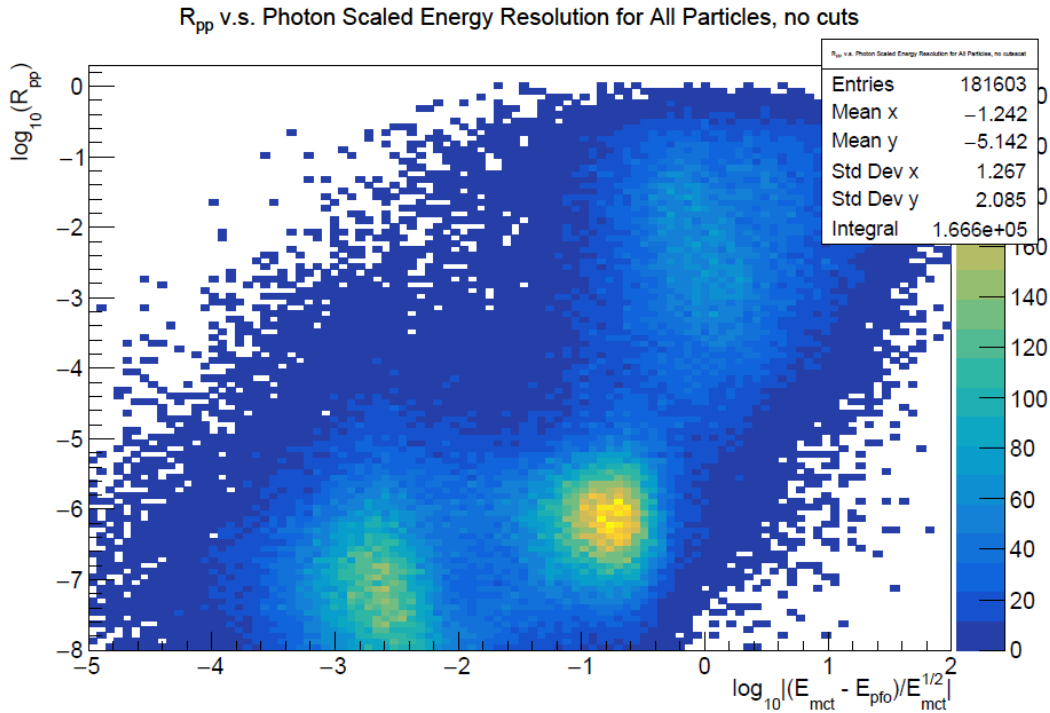


Figure 20: $\log(R_{pp})$ vs. $\log(\text{Photon Scaled Energy Resolution})$

A 2D histogram of R_{pp} vs. scaled energy resolution for photons, calculated before any of the cuts mentioned in the beginning of this section aside from $Z \rightarrow q\bar{q}$. This informs the cut of $R_{pp} > 10^{-5}$, removing the worst-resolution concentration of photons.

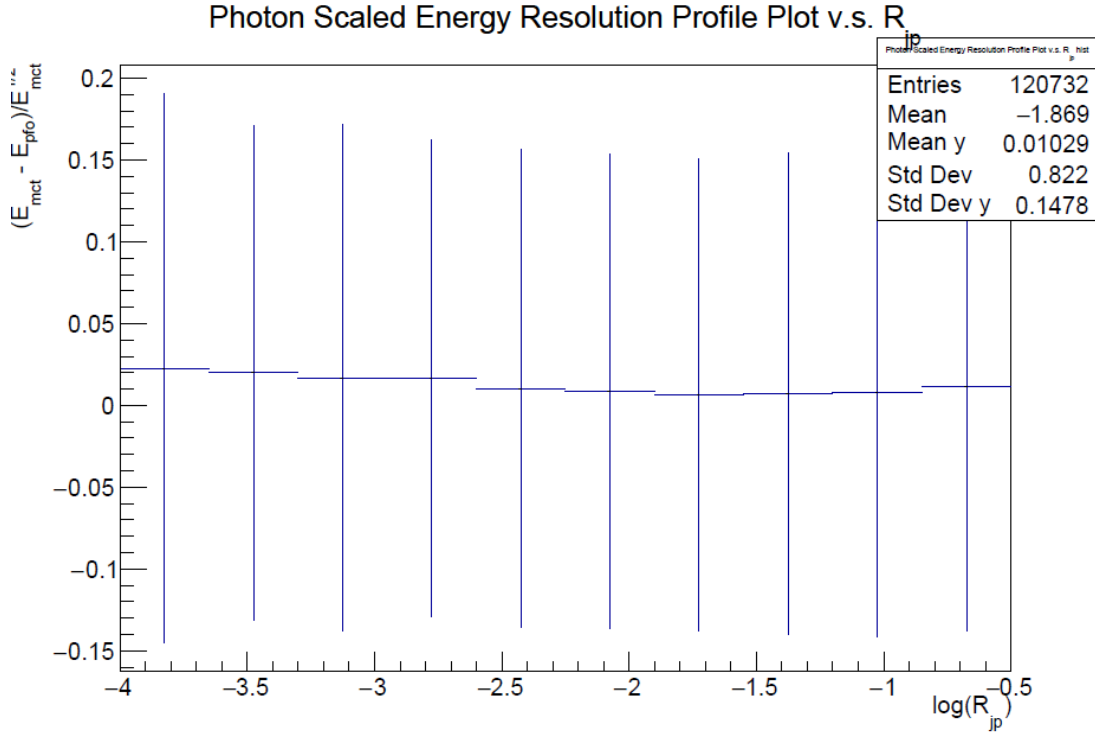


Figure 21: Scaled Photon Resolution vs. $\log(R_{jp})$

A profile plot of the scaled energy resolution of particles for each R_{jp} . The horizontal lines indicate the mean scaled energy resolution of each bin, and the vertical lines denote the RMS of the distribution.

Though the effect is not overly pronounced, there is a clear trend of scaled resolution improving as R_{jp} increases, up to roughly $R_{jp} = 10^{-2} = 8$ degrees before leveling off. This means that photons in the center of a jet tend to be reconstructed slightly more poorly than those on the outskirts. The concentration of particles tends to increase as one approaches the center of a jet, so this can be interpreted as evidence that in the busy inner-jet environment the particles get close enough that the detector is not able to reconstruct them as well.

This effect is even more pronounced when the photons are split into >5 , $3-5$, and <3 GeV energy ranges, with the higher energies showing a more pronounced trend (Figure 22). These higher-energy photons comprise a higher proportion of the overall jet energy, so are more important than the lower-energy photons to measure accurately. Seeing that the TDR design's

reconstruction ability degrades for these higher-energy photons in busier environments is a strong point in favor of MAPS.

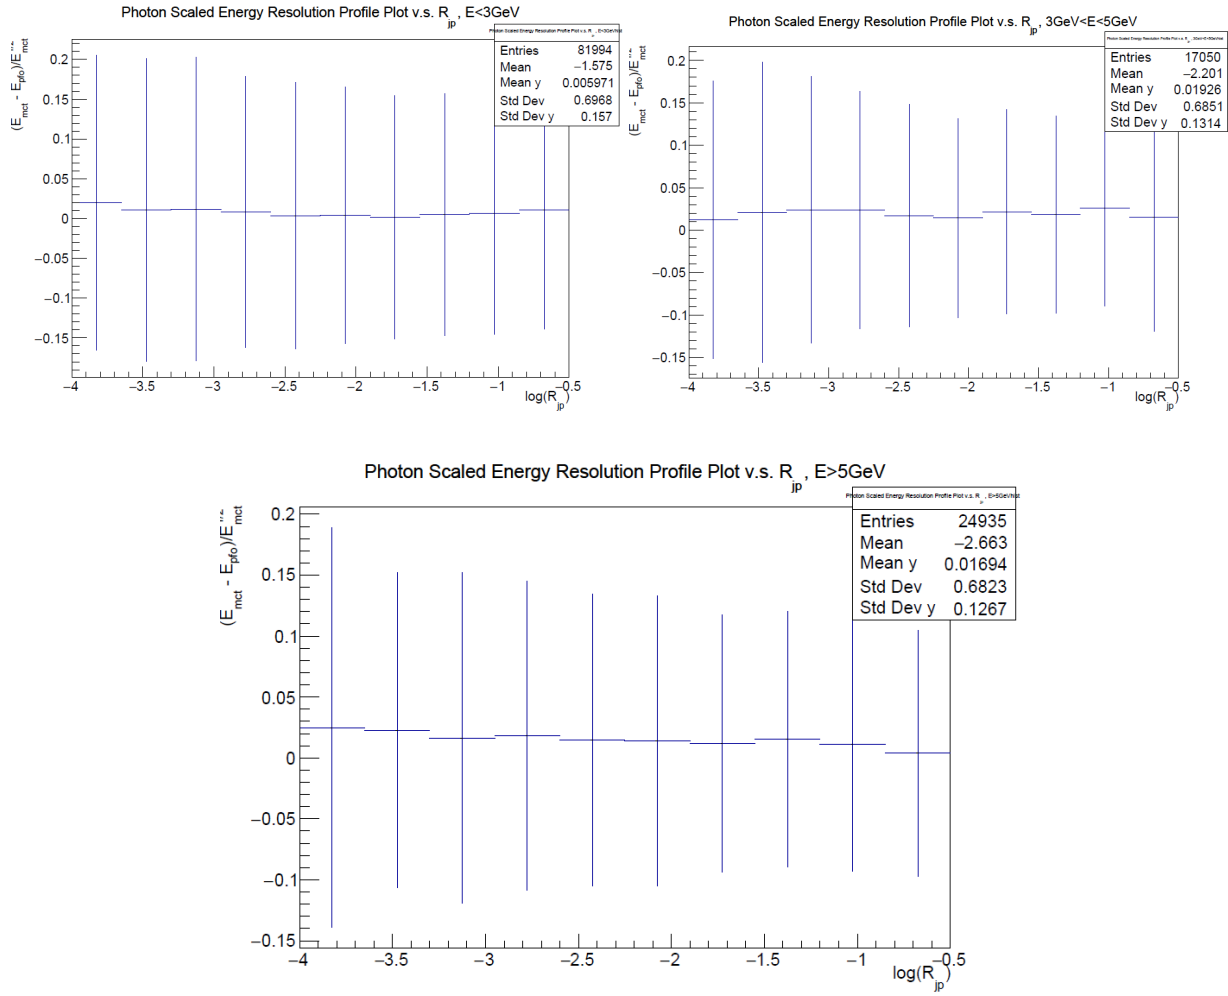


Figure 22: Scaled Photon Resolution vs. $\log(R_{jp})$ for <3, 3-5, 5GeV Photons

Profile plots of the scaled energy resolution of photons for each R_{jp} . Each plot is the same as in Figure 21 except split into three energy ranges.

Summary & Outlook

There is substantial evidence in favor of MAPS as a replacement of the TDR SiD Ecal design due to the potential for performance improvements provided by the higher-granularity pixels. By using clustering methods, MAPS is able to replicate the energy resolution of the TDR design, meaning that MAPS does not sacrifice the detector's ability to measure energy in return for its improved spatial resolution. Additionally, there is still room for improvement; though these clusters are dominated by low-complexity structures, the analysis may still be augmented further through a more sophisticated algorithm. This might include structure-based weighting or algorithms which can place non-adjacent hits into the same cluster or separate adjacent hits into different clusters. However, the complexity of the algorithm must be weighed against diminishing returns, given the low number of complex clusters for these methods to improve upon. Nonetheless, this represents an interesting line for future research to pursue, though for our purposes it is sufficient to confirm that the performance of MAPS at least matches the TDR.

Additionally, we observe a degradation in the TDR detector's ability to measure particles in the center of jets, where they tend to be most concentrated. The smaller pixel size of MAPS allows for easier spatial separation of particles within this environment, presenting an opportunity for improvement of the detector's measurement. Further research is ongoing to find more evidence to confirm this, looking into how the PFO forms each particle from clusters of energy and data from the tracker. Understanding these clusters (different from the clusters mentioned in the first half of this work) can allow us to understand in more detail how the PFO is performing its reconstruction and precisely where this begins to break down. The observed systematic underestimation of PFO energy has also been noted and will be fixed in the future.

However, given MAPS's demonstrated potential to match and exceed the TDR baseline, it is probable that future efforts will focus on optimizing the simulation for MAPS instead.

References

- [1] International Linear Collider, International Linear Collider. (n.d.). Retrieved from <https://linearcollider.org/>
- [2] P. Bambade et al., The International Linear Collider: A Global Project. 2019. doi:10.2172/1527401. Retrieved from <https://arxiv.org/pdf/1903.01629.pdf>
- [3] H. Baer et al., The International Linear Collider Technical Design Report - Volume 2: Physics. 2013. Retrieved from <https://arxiv.org/ftp/arxiv/papers/1306/1306.6352.pdf>
- [4] Particle Data Group, Particle Data Group. (n.d.). Retrieved from <https://pdg.lbl.gov/>
- [5] A. Steinhebel, Much Ado About Nothing: Searches for Higgs Boson Decays to Invisible Particles. June 2021. Retrieved from <https://scholarsbank.uoregon.edu/xmlui/handle/1794/26657>.
- [6] T. Behnke et al., The International Linear Collider Technical Design Report - Volume 4: Detectors. Jun 26, 2013. Retrieved from <https://arxiv.org/abs/1306.6329>
- [7] M. Breidenbach et al., Updating the SiD Detector concept. Oct 19, 2021. Retrieved from <http://export.arxiv.org/abs/2110.09965>
- [8] J. E. Brau., The SiD Digital ECal based on Monolithic Active Pixel Sensors: ILCX. Oct 27, 2021. Retrieved from <https://agenda.linearcollider.org/event/9211/contributions/49457/attachments/37426/58628/ILCX-brau.pdf>
- [9] Geant4. (n.d.). Retrieved from <https://geant4.web.cern.ch/about/>
- [10] ROOT Team, ROOT. 2023. Retrieved from <https://root.cern/>
- [11] J. E. Brau et al., The SiD Digital ECal Based on Monolithic Active Pixel Sensors. Instruments 2022, 6(4), 51 (2022). Retrieved from <https://doi.org/10.3390/instruments6040051>
- [12] M. A. Thompson, Particle Flow Calorimetry and the PandoraPFA Algorithm. Jul 21 2009. Retrieved from <https://arxiv.org/abs/0907.3577>
- [13] Scikit-HEP Project, PyJet. Jan 6 2023. Retrieved from <https://github.com/scikit-hep/pyjet>