Upgrades to the Event Simulation and Reconstruction for the Fermi Large Area Telescope

Leon Rochester and Tracy Usher (SLAC), Robert P. Johnson and Bill Atwood (SCIPP/UCSC)
on behalf of the Fermi Large Area Telescope Collaboration

Abstract

The pre-launch event simulation and reconstruction performed beyond expectation for real data, essentially without modification, and made possible the immediate start of science analysis. But the on-orbit data exhibit unanticipated features that necessitate upgrades to both the simulation (essentially complete) and the reconstruction (ongoing). The major new effect encountered on orbit is the presence of "ghosts," that is, remnant detector response to particles passing through the detector before the particle that triggered the event. These ghosts appear primarily in the form of extra tracks and/or energy deposits. As part of this upgrade, we plan to enhance our ability to discriminate against background particles by introducing additional analysis during the reconstruction phase. We present a description of the effect of ghosts, and of the work needed to deal with them, done and planned, as well as some other ideas for improving the reconstruction.

Our encounter with ghosts

When data started coming down from the Fermi Large Angle Telescope (LAT), we discovered that a few percent of the events were of a type that we hadn't anticipated. These events contained more than one track, clearly separated from each other. Here's an example of such an event:

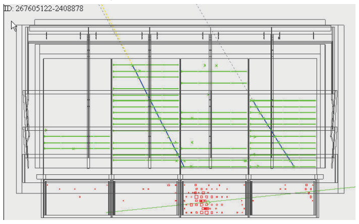

The green horizontal line segments and **X**s show the hit "x" and "y" strips on the silicon detectors in the tracker (TKR), and the blue lines show the found tracks. The "best" track is indicated by the dotted yellow line projecting upward from the head of the track. The red squares at the bottom indicate energy deposited in the calorimeter (CAL); the green box shows the centroid of the deposited energy, and the nearly horizontal green line shows the inferred axis of the shower.

In its present form, this event is essentially useless. First, we don't know what to make of the two tracks; could this be a shower of some sort? Second, the current energy calculation lumps all the deposited energy into a single cluster, whose direction is obviously incorrect in this case, as is, to a lesser extent, its centroid.

The origin of these extra tracks turns out to be the finite response time for each of the subsystems in the LAT. For example, the signal in the TKR of a hit from a cosmic ray track will be present for at least ten μsec after it passes through the detector, and longer for higher-Z particles. Such latent hits are nicely captured by the periodic trigger (PT) in the LAT. This trigger is generated internally at a uniform rate of two per second, completely independent of the actual state of the detector, and thus gives us a picture of the background environment on orbit.

The extra signals are tracks that crossed the detector at some time before the trigger, and whose detector response has decayed by the time of the trigger. This is borne out by the timing information recorded with the event. We call these tracks "ghosts."

Most of the PT events are essentially empty. But here's a sample of PT events with significant ghosts:

The first has a signal in the CAL, but none in the TKR; the second, vice versa; and the bottom two have signals in both the CAL and TKR, as well as reconstructed tracks. In principle, there can also be remnant signals in the anti-coincidence detector (ACD), but generally these have a shorter time constant and decay quickly; hence, they are seldom seen.

The charged-particle background rate (R_B) in the detector depends mainly on the local magnetic field. Here, it is shown as a function of earth coordinates.

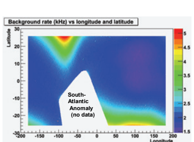

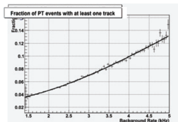

Using the PTs we can measure the potential contamination of real events due to ghosts as a function of the measured R_B.

Upgrading the simulation

Because a ghost can corrupt a gamma event so that the current analysis no longer recognizes it, the presence of ghosts affects the efficiency of detection of gammas, in a way that depends on the charged-particle trigger rate and the energy of the triggering event. We hope to ultimately recover many of these events. But as a first step, we decided to try to account for these ghosts in our Monte Carlo (MC) simulation. To this end, we modified the simulation so that we could overlay a real PT event on top of each simulated γ event. The resulting events should be very similar to our actual gamma events.

This indeed turned out to be the case. At the right, we show the effective area of these "haunted" MC γ-rays relative to that for unmodified ones, as a function of the energy of the gamma.

As a check of the simulation, we can divide our data into samples with different R_Bs and compare our measurements of the γ-ray flux of the Vela pulsar for each R_B bin.

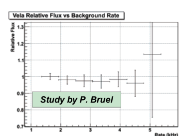

Each point in the plot above shows the result of the flux calculation using the effective area for that R_B bin, as determined by the simulation. The results are found to be independent of the background rate of the sample. Similar results are obtained for the other parameters of the fit, for example, the power-law index.

Finding CAL clusters

The first step, essential for developing a more effective event reconstruction, is to identify isolated clusters of energy deposition in the CAL. This seems straightforward, but turns out to be challenging.

One complication comes from the substantial gaps between the CAL modules in adjacent towers. We have developed a strategy to link up pieces of the cluster that cross towers.

Another issue is the treatment of single isolated crystals. For now we count each of these as a cluster, but we may eventually learn how to attach them to a nearby extended cluster.

Finally, we will need to re-tune the reconstruction of the energy associated with each event, since removing isolated crystals from the overall cluster will involve an apparent decrease of measured energy.

The displays below show an example of clustering, especially across tower boundaries.

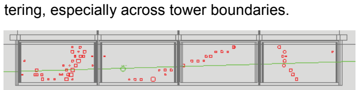

Clustering of a complex event using the current single-cluster algorithm. The green line shows the axis of the resulting "cluster."

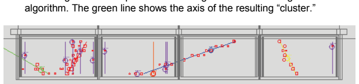

The same event with the newly developed multi-cluster algorithm. The resulting cluster axes are color-coded roughly in order of decreasing energy: red, orange, yellow, green, blue and violet. Some of the single-crystal clusters probably belong to nearby clusters.

This first attempt has yielded promising results, but we will no doubt have to refine it as we better understand the constraints of the problem.

Identifying ghost tracks

Tagging ghost tracks with trigger info

It turns out that many of the TKR hits produced by ghost tracks can be easily identified. This is because they are out of time and fail to trigger at the level of individual TKR planes or full towers. The triggering efficiency of normal tracker hits is well above 99%; thus a track whose component hits should have generated a trigger, but didn't, can be tagged as a ghost.

Here's an event with a normal track and one that consists of ghost hits tagged by trigger information.

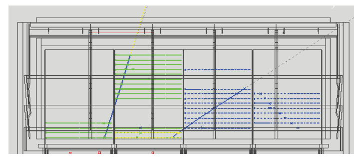

The tagged ghost hits are colored blue. The hits colored yellow are ones that happen to produce plane triggers (and thus are not tagged), but which end up on the tagged track. Such hits are very likely to come from the ghost particle.

Also note the track on the right, which was not found by our current code. We hope to improve on this situation in the next round!

Using the new information

Here is the same event shown at the beginning of the poster, but with the added information coming from the upgrades to the TKR and CAL analyses:

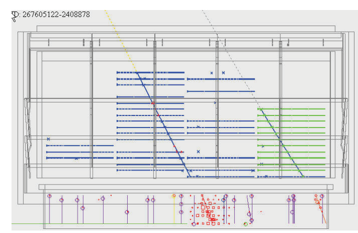

We see that the track on the left, which was originally chosen as the best track, is in fact a ghost, and that the energy of the most energetic CAL cluster in the event (1726 MeV) is actually associated with this ghost track. The cluster associated with the "real" track, on the right, has an energy of 129 MeV. Note that the axes of both of the clusters line up with their corresponding TKR tracks.

Based on the TKR time-over-threshold signals and the energy deposit in the CAL, the real track appears to be an upward-going proton that stops in the middle of the tracker. The ghost seems to be the remnant of a heavy ion.

Although unraveling this particular event was relatively straightforward, the typical haunted event will be more complicated.

Putting it all together

The remaining piece of the puzzle will be the development of strategies to associate tracks and clusters and to characterize the particles in complex events. This will also include accounting for the events which are still too complicated to analyze.

This will require redesigning the current reconstruction algorithm to be more flexible, allowing for iteration and multiple paths.

While the specific implementation of our reconstruction algorithm doesn't yet have these features, the framework itself is designed to allow this flexibility.

Other upgrades

The original aim of our reconstruction strategy was to find gammas, not background. But as we understand more about our actual problem, we've come to understand that we must be able to do both.

Identifying more background tracks

Since the hit efficiency of the TKR is so high, it's very unlikely that more than one or two hits will be missing on a normal track. So our track-finding is tuned to reject tracks with multiple gaps. But in many cases, the hits on a ghost track have decayed to the point where a substantial number of them are missing, even though it's obvious by eye that they belong to a track. The upper-right event in the PT sampler shows such a track. For the reasons cited above, these defective tracks are usually not found. The hits themselves are often tagged by their failure to trigger, but recognizing them as tracks may help us to eliminate any CAL cluster that they point to.

ACD-seeded track-finding

As noted earlier, we've seen events with seemingly findable TKR tracks that are not actually found. An important class of such tracks consists of those that point to a struck ACD tile. These are missed background tracks, which may have actually triggered the event.

The event display below shows one such missed track. Since such particles can be a source of background, we need to tune our software to better identify and deal with them.

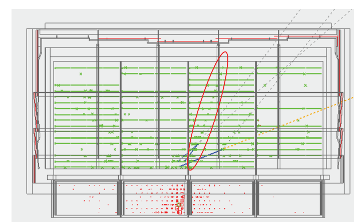

"Vector-links" track-finding

Our current track-finding is based on local track following and tends to find many track stubs in EM showers. Only the one or two tracks at the head of the shower contribute any information about the direction of the gamma; the remaining tracks are essentially useless, except to indicate the presence of a shower.

We are exploring a complementary approach, in which we try to link up local vectors of hits into a global structure. Two examples are shown below.

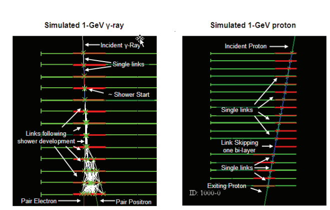

Initially, we are using this approach to find individual tracks, in order to compare it to the current one. It appears to be roughly comparable, even at this early stage. Eventually we hope to derive the kinematic information about the event directly from the global structure and eventually to characterize the "gamma"-ness of the event.