ONLINE ICRC 2021
THE ASTROPARTICLE PHYSICS CONFERENCE
Berlin | Germany
37th International Cosmic Ray Conference
12–23 July 2021

A Fast GRB Source Localization Pipeline for the Advanced Particle-astrophysics Telescope

Marion Sudvarg[a], Jeremy Buhler[a], James Buckley[a], Wenlei Chen[a,b], Zachary Hughes[a], Emily Ramey[a,c], Michael Cherry[d], Samer Alnussirat[d], Ryan Larm[a], and Cristofer Berruz Chungata[a,e] for the ADAPT Collaboration

[a]Washington University in St. Louis, St. Louis, Missouri, United States [b]University of Minnesota, Minneapolis, Minnesota, United States
[c]University of California, Berkeley, California, United States [d]Louisiana State University, Baton Rouge, Louisiana, United States
[e]University of Bridgeport, Bridgeport, Connecticut, United States

Introduction

The Advanced Particle-astrophysics Telescope (APT) [1] is a space-based observatory, currently in development, to survey the entire sky for gamma-ray sources in the MeV to 100 GeV range. APT's goals include prompt detection of energetic transient events in the distant universe, such as gamma-ray bursts (GRBs), and rapid communication of these events to narrow-band instruments that can conduct follow-up observations in other spectral bands. Pursuant to this goal, we are developing analytical methods to perform real-time detection and localization of events, which will run on computing hardware on the orbiting APT platform.

In this work, we focus on detecting events for which most incident photons have energies in the Compton regime (511 keV to a few MeV). For such events, the photons of interest Compton-scatter one or more times within the detector, depositing energy with each scattering, until they are eventually photoabsorbed. All scatterings for one photon appear simultaneous at the time resolution of the detector. Our analytical tasks are twofold: first, to identify the first two scatterings for each photon in order of occurrence, which localizes the source's direction to an annulus centered on the line connecting these scatterings; and second, to combine the annuli from all detected photons to estimate the most likely direction in the sky for a distant point source (the GRB) that emitted them.

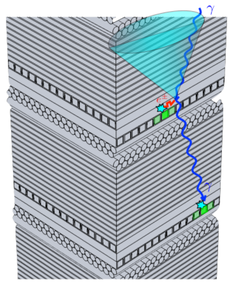

Figure 1: A cross-section of the APT instrument. A gamma-ray photon, γ, enters the instrument from the top, then Compton-scatters, before finally being photoabsorbed. The vector between the two interactions, and the scattering angle inferred from the photon's energy before and after the first interaction, define a circle that describes the set of possible source vectors of the photon.

Several computational challenges arise in building a robust pipeline of algorithms for event detection and localization. The pipeline must accurately localize even low-fluence events (at most a few thousand total incident photons) while being efficient enough to keep up with high-fluence events (around 10^5 photons/sec) involving many photons that might scatter 5 or more times within the detector. It must yield results quickly enough to permit rapid retargeting of narrow-band instruments to the event – ideally in well under a second. And it must deliver this performance using a low-power processor of the type feasible for a space-based platform. All of these criteria must be met while also dealing robustly with the measurement limitations of the APT detector.

This work describes a computational pipeline for Compton-regime reconstruction and event localization. We build on the basic approach of Boggs and Jean [2] to reconstruct photon trajectories within the instrument by minimizing disagreement between the reconstructed angle of each scattering and the energy it deposited. To eliminate redundant computation and ensure rapid analysis even of photons with multiple scatterings, we implement a tree search with pruning over possible photon trajectories to find one with the best agreement. Event localization from reconstructed photons then follows a maximum-likelihood approach, with random sampling of reconstructed photons to guess a plausible source direction, followed by iterative refinement. We demonstrate that our pipeline can reliably localize events to within 2.53 degrees (68% containment) for low-fluence events and 0.42 degrees for high-fluence events while delivering results within 200 ms even on a low-power ARM Cortex-A53 processor.

Acknowledgments

This work was made possible through current and ongoing support from NASA grant 80NSSC19K0625, NASA APRA award 20-APRA20-0148, and NSF award CNS-1763503. The collaboration also acknowledges generous ongoing support from the Washington University McDonnell Center for the Space Sciences and the Peggy and Steve Fossett Foundation. We are grateful to the Washington University technical staff including Richard Bose, Dana Braun and Garry Simburger who made invaluable contributions to the development and construction of the prototype detectors and the Antarctic flight of the APTlite instrument.

References

[1] J. Buckley, S.A. Nussirat, C. Altomare, J. Buhler, E. Burns, R. Chamberlain et al., *The Advanced Particle-astrophysics Telescope (APT) project status*, in *ICRC 2021*, Astroparticle Physics Conf., 2021.

[2] S. Boggs and P. Jean, *Event reconstruction in high resolution Compton telescopes*, *Astronomy and Astrophys. Supp. Series* **145** (2000) 311.

[3] O. Klein and Y. Nishina, *Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac*, *Zeitschrift für Physik* **52** (1929) 853.

[4] W. Gander, G.H. Golub and U.V. Matt, *A constrained eigenvalue problem*, *Linear Algebra Appl.* **114-115** (1989) 815.

[5] F. Tisseur and K. Meergergen, *The quadratic eigenvalue problem*, *SIAM Rev.* **43** (2001) 235.

[6] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce et al., *Geant4 — a simulation toolkit*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506** (2003) 250.

[7] W. Chen, J. Buckley, S.A. Nussirat, C. Altomare, J. Buhler, E. Burns et al., *The Advanced Particle-astrophysics Telescope: simulation of the instrument performance*, in *ICRC 2021*, Astroparticle Physics Conf., 2021.

[8] D. Band, J. Matteson, L. Ford, B. Schaefer, D. Palmer, B. Teegarden et al., *BATSE observations of gamma-ray burst spectra. I. spectral diversity*, *Astrophys. J.* **413** (1993) 281.

Compton Event Reconstruction

For each incident gamma-ray, we seek to infer the circle, illustrated at the top of Figure 1, that describes the set of possible source vectors of the photon. The instrument reports a sequence of N gamma-ray interactions for each photon, but the ordering is unknown. We must infer the most likely ordering according to the approach of Boggs and Jean [2]:

Find the ordering that minimizes the sum χ^2 of disagreements between the measured angle of each scattering and the angle implied by its deposited energy, according to the Compton law.

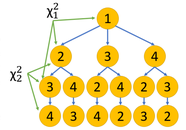

Figure 2: A subtree for an event with $N = 4$ representing all permutations beginning with interaction 1.

- The statistic χ^2 can be expressed as $\sum_{i=2}^{N-1} \chi_i^2$.
- Computing this sum for all $N!$ possible orderings can be quite slow for large N.
- To enable rapid reconstruction, we implement a *tree traversal* over possible orderings.
- A node at tree depth k corresponds to an ordering of the first k interactions.
- Each triple of interactions in a given sequence contributes to one term to the sum, as shown in Figure 2.
- For all children of a node of depth k, the first $k-2$ terms of χ^2 are the same and so need to be computed only once.
- We also precompute values for each possible triple.
- If the partial sum for a node at depth k exceeds either the current best χ^2 sum or the $p = 0.001$ significance level for a χ^2 statistic with $k-2$ degrees of freedom, we prune the tree at this node.

Localization

Given N annuli, each described by a triple $\langle \mathbf{c}_k, \phi_k, \sigma_k \rangle$, we seek to estimate the true source direction $\mathbf{s}$. First, we compute a rough approximation of the source direction. Second, we use an iterative least-squares refinement to derive a final estimate of $\mathbf{s}$.

Identifying an Approximate Source Direction

1. Select a single annulus i from the input at random.
2. Test a set of evenly spaced candidate source directions $\mathbf{s}_j$ that lie on the circle $\langle \mathbf{c}_i, \phi_i \rangle$, as illustrated in Figure 3.
3. Find the likelihood of each source direction with respect to all input annuli k according to:

$$\mathscr{L}\left(\mathbf{s}_j \mid \langle \mathbf{c}_k, \phi_k \rangle\right) \approx \frac{1}{\sigma_k \sqrt{2\pi}} e^{-(\phi_k - \beta_k)^2 / 2\sigma_k^2} \quad \text{where} \quad \beta_k = \arccos\left(\mathbf{s}_j \cdot \mathbf{c}_k\right)$$

4. Select the candidate direction $\mathbf{s}_j$ with the greatest joint log-likelihood as an approximate source direction.
5. Repeat steps 1 to 4 for 20 random input annuli.
6. Produce a combined estimate $\mathbf{s}_0$ as the average over these 20 estimates, each weighted by its likelihood.

Iterative Least-Squares Refinement of Source Direction

1. Begin with the estimate $\mathbf{s} = \mathbf{s}_0$, the vector produced by the approximation step.
2. For each annulus i, test whether the angle $\arccos(\mathbf{c}_i \cdot \mathbf{s})$ lies within $3\sigma_i$ of ϕ_i, as illustrated in Figure 4.
3. For those annuli that do, generate linear constraints $\mathbf{c}_i \cdot \mathbf{s} = \cos\phi_i$.
4. Require $\mathbf{s}$ to be a direction vector; this unit-norm constraint is quadratic in the coordinates of $\mathbf{s}$.
5. Reduce the problem to a quadratic eigenvalue problem as in [4].
   - The matrix has dimension proportional to the three coordinates of $\mathbf{s}$, independent of the number of photons.
   - Quadratic eigenvalue problems are computationally tractable for small matrices [5].
   - Forming the matrix for the problem has cost quadratic in the number of input photons.
6. Solve to get a refined estimate for $\mathbf{s}$.
7. Iterate 20 times, repeating steps 2 to 6. Each time pass in the refined estimate from the previous iteration; the final solution gives the estimated source direction of the GRB.

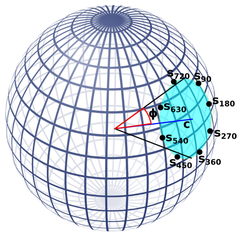

Figure 3: For initial approximation, candidate source directions are evenly-spaced around a randomly-selected annulus.

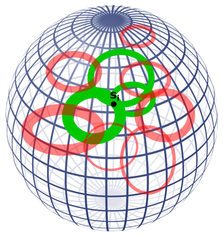

Figure 4: At refinement step i, only annuli (shown in green) within 3 standard deviations of the current estimated source vector $\mathbf{s}_i$ are used as constraints for the least-squares problem. Other annuli (shown in red) are ignored.

Evaluation Overview

- We simulated a typical long GRB using Geant4 [6] and APTsoft [7].
- Generated 10^6 photons
- Normally incident
- Distributed across a collimated beam with a cross-section of 18 m^2 to fully cover the APT detector
- Spectral-energy distribution modeled according to a Band function [8] with parameters:
  - $\alpha = 0.6$
  - $\beta = -2.5$
  - $E_{\text{peak}} = 1$ MeV
  - Energy range 300 keV to 10 MeV
- In this energy regime, Compton scattering effects dominate, though pair production does occasionally occur. We do not attempt to distinguish pair-production events in our evaluation of accuracy and execution time.

Execution Times

- We test the efficiency of our pipeline on a low-power embedded platform representative of what might fly on APT:
  - Platform: Raspberry Pi 3 Model B+
  - CPU: Cortex-A53 (ARMv8) quad-core, 1.4 GHz, 64-bit
  - Memory: 1 GB LPDDR2 DRAM
- We ran 200 trials each for:
  - Fluences from 0.01 to 0.05 MeV/cm^2 in increments of 0.01
  - Fluences from 0.1 to 1.0 MeV/cm^2 in increments of 0.05
- We used OpenMP to parallelize across cores.
- We introduced a faster, lower-precision approximation to the arccosine needed to obtain β_i that reduced its cost by an order of magnitude.

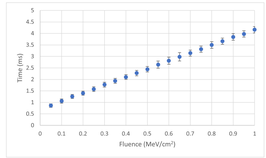

Reconstruction Times. Since each photon is reconstructed independently, reconstruction times scale linearly with fluence.

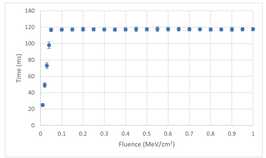

Initial Approximation. Execution time increases with fluence until the number of annuli exceeds $N = 1000$, after which we sample a constant-size subset.

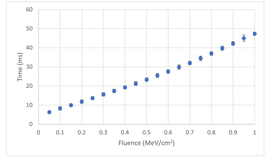

Iterative Refinement. Running times increase quadratically.

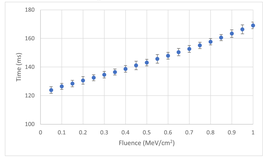

Total Pipeline Time. Time remains under 200 ms, with low variance, for high fluence.

Figure 5: Pipeline execution times for 200 trials per fluence. Error bars denote a single standard deviation about the mean.

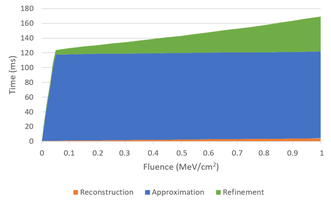

Figure 6: Mean execution times over 200 trials for each fluence tested, highlighting contributions from each pipeline stage.

Localization Accuracy

We measured localization accuracy across a range of fluences, running our pipeline 1000 times for each fluence, each time randomly sampling photons from the corpus of 10^6 events.

Fluence	Mean Error	Std Dev	68% Containment	95% Containment
0.03 MeV/cm^2	2.15	1.22	2.53	4.42
0.1 MeV/cm^2	1.21	0.64	1.45	2.32
0.3 MeV/cm^2	0.70	0.36	0.87	1.32
1.0 MeV/cm^2	0.35	0.20	0.42	0.72

Table 1: Localization error (in degrees) measured over 1000 trials.

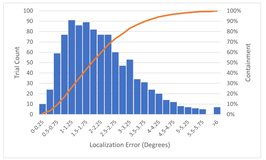

Fluence 0.03 MeV/cm^2

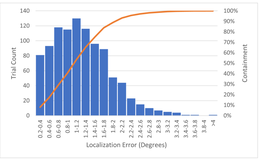

Fluence 0.1 MeV/cm^2

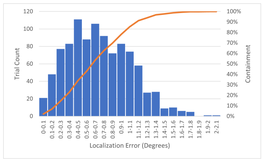

Fluence 0.3 MeV/cm^2

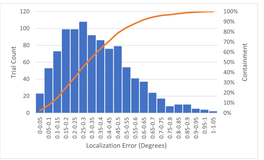

Fluence 1.0 MeV/cm^2

Figure 7: Pipeline accuracy distributions for 1000 trials per fluence.

Error Parameters

We estimate the spatial and energy error parameters of the detector, used in our calculation of the χ^2 value and the propagated error σ of each annulus used in localization, by comparing simulated true interaction details with output from our simulated digitizer. To achieve better fits, we eliminate pair-production events from our dataset.

Spatial Error

- We estimate the parameter $\Delta = \delta x = \delta y$.
- We assume Δ is approximately Gaussian.
- We compare the true simulated x- and y-axis coordinates for each interaction with the coordinates given by the simulated digitizer.
- We generated a histogram of error values, then fit a normal distribution to the plot.
- The standard deviation of the best fit, 0.7 mm, is assigned to Δ.

Energy Error

- We estimate the parameter δE.
- We assume δE is a function of energy: $\delta E = \sigma_E \sqrt{E}$.
- We compared the true simulated energy deposits E' with the energies E reported by the digitizer.
- We generated a histogram of values $\langle E - E' \rangle / \sqrt{E'}$, then fit a normal distribution to the plot.
- The standard deviation of the best fit, 0.095 MeV$^{1/2}$, is assigned to σ_E.

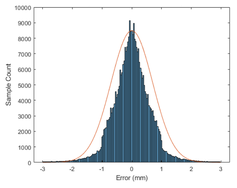

Figure 8: Histogram of measured coordinate noise, fit to a normal distribution with $\sigma = 0.7$ mm

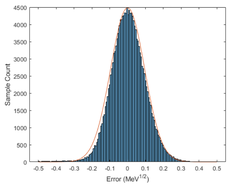

Figure 9: Histogram of measured energy noise, normalized to $\sqrt{E'}$, fit to a normal distribution with $\sigma = 0.095$ MeV$^{1/2}$