

Coordinating Instruments for Multi-Messenger Astrophysics

**Daisy Wang
Ye Htet
Marion Sudvarg
Roger Chamberlain
Jeremy Buhler
James H. Buckley**

Daisy Wang, Ye Htet, Marion Sudvarg, Roger Chamberlain, Jeremy Buhler, and James Buckley, "Coordinating Instruments for Multi-Messenger Astrophysics," in *Proc. of 22nd ACM International Conference on Computing Frontiers Workshops and Special Sessions*, May 2025, pp. 213-218.

DOI: 10.1145/3706594.3727948

Invited paper at the Special Session on Computer Architectures in Space (CompSpace).

Dept. of Computer Science and Engineering
and
Dept. of Physics
Washington University in St. Louis

Coordinating Instruments for Multi-Messenger Astrophysics

Invited Paper

Daisy Wang

Washington University in St. Louis
Saint Louis, Missouri, USA
w.yanwang@wustl.edu

Ye Htet

Washington University in St. Louis
Saint Louis, Missouri, USA
htet.ye@wustl.edu

Marion Sudvarg

Washington University in St. Louis
Saint Louis, Missouri, USA
msudvarg@wustl.edu

Roger Chamberlain

Washington University in St. Louis
Saint Louis, Missouri, USA
roger@wustl.edu

Jeremy Buhler

Washington University in St. Louis
Saint Louis, Missouri, USA
jbuhler@wustl.edu

James Buckley

Washington University in St. Louis
Saint Louis, Missouri, USA
buckley@wustl.edu

Abstract

In multi-messenger astrophysics, signals of multiple types (e.g., gravitational waves, neutrinos, electromagnetic waves) are combined in an effort to learn more about the observed phenomena of interest. The Advanced Particle-astrophysics Telescope (APT) is a mission concept for a space-borne instrument that detects gamma-ray bursts (GRBs) omnidirectionally, facilitating multi-messenger observations by identifying and localizing celestial events of interest. Here, we describe the on-instrument computations for APT and its Antarctic Demonstrator (ADAPT) as well as techniques for follow-up observations of transient events.

CCS Concepts

• **Applied computing** → **Astronomy; Physics; • Computer systems organization** → **Embedded software; Embedded hardware.**

Keywords

Gamma-ray astronomy, multi-messenger astrophysics

ACM Reference Format:

Daisy Wang, Ye Htet, Marion Sudvarg, Roger Chamberlain, Jeremy Buhler, and James Buckley. 2025. Coordinating Instruments for Multi-Messenger Astrophysics: Invited Paper. In *22nd ACM International Conference on Computing Frontiers (CF Companion '25)*, May 28–30, 2025, Cagliari, Italy. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3706594.3727948>

1 Introduction

The astrophysics community has a strong interest in observing transient celestial events using multiple imaging modalities. This *multi-messenger* approach may include, e.g., gravitational waves, electromagnetic waves, neutrinos, and cosmic rays [15, 16]. Because transients can represent short-lived phenomena [1], fast detection

and localization is key to supporting cooperative multi-modal observation. This work describes cooperative multi-wavelength electromagnetic observation: localizing a transient from its gamma-ray emissions, with follow-up performed by optical telescopes. However, the computational techniques we employ generalize to multi-messenger observation.

The Advanced Particle-astrophysics Telescope (APT) [7] mission concept is a proposed gamma-ray and cosmic-ray observatory that will orbit the Sun-Earth L_2 Lagrange Point, which avoids obstruction by the earth and ensures a nearly omnidirectional, 4π -steradian field of view (FoV). Its goals include prompt detection and localization of gamma-ray bursts (GRBs), which are early indicators of, e.g., neutron star and black hole mergers, blazar and magnetar flares, and supernovae. APT's localizations will permit follow-up observation of such events by optical telescopes, which typically have quite narrow FoVs. It is predicted to localize GRBs with better than 1° accuracy and computational latency under 1 s [11]. Our recent work towards APT focuses on improving the accuracy of event localization in the presence of noise and uncertainty.

Other full-sky instruments have larger localization uncertainties, e.g., $>20^\circ$ [3] for the gravity-wave detectors LIGO and VIRGO, and $\sim 5\text{--}10^\circ$ [12] for APT's upcoming Antarctic Demonstrator (ADAPT). To convey this uncertainty, these instruments must supplement point localization with more detailed *likelihood maps* that provide a spatial probability distribution over possible locations. The most likely regions of the map can then be searched by fast-slewing, wide-FoV ($\sim 2\text{--}10^\circ$) optical telescopes to localize a source for observation by more sensitive, narrow-FoV ($\leq 1^\circ$) instruments.

Several factors contribute to the delay between detection and secondary observations, including time to localize/map the transient, communication latency to follow-up instruments, and the time for these instruments to physically search the sky as directed by a likelihood map. Real-time computation of localizations and maps is particularly challenging for space-based instruments like APT, for which limited communication bandwidth and high speed-of-light latency to earth (APT will orbit 5 light-seconds away) make real-time processing on the ground impractical. As a result, the computational pipeline that transforms raw sensor data into localizations must execute on the instrument (in space) [7] and must therefore meet significant size, weight, and power (SWaP) constraints.

This work describes four important computational elements associated with APT and their ongoing development for ADAPT:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CF Companion '25, Cagliari, Italy

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1393-4/25/05

<https://doi.org/10.1145/3706594.3727948>

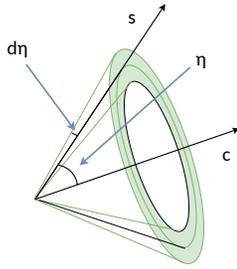


Figure 1: A gamma ray’s Compton ring is defined by the vector c between its first two interactions and cosine η (with uncertainty $d\eta$) of its first scattering angle (from [12]).

(1) improved trajectory reconstruction for individual gamma rays, (2) use of machine learning in point localization, (3) work in progress to generate likelihood maps, and (4) instrument control approaches to implement the follow-up observations.

2 Trajectory Reconstruction

APT and ADAPT infer a GRB’s direction by combining the trajectories of individual gamma-ray photons that interact within them. This section gives an overview of the instrument designs and the methods we are developing to reconstruct gamma ray trajectories from raw sensor data produced by their detectors.

The APT and ADAPT Detectors. APT and ADAPT are constructed with layers of scintillating tiles that emit visible light when incoming gamma-ray photons scatter within them. This light is captured by perpendicular arrays of wavelength shifting (WLS) optical fibers that line the top and bottom surfaces of the tiles, then measured by silicon photomultipliers (SiPMs) placed at the end of each fiber. This overlay of 1-dimensional fiber arrays into a 2-dimensional mesh, along with the relative position of the tile, allows us to resolve the 3-dimensional position $\mathbf{r} = (x, y, z)$ of each gamma-ray/scintillator interaction. Additional SiPMs, placed around the edges of the tile layers, improve light collection and provide an estimate of the energy E deposited with each interaction.

Both instruments will fly with onboard computational hardware, including preamplifier and waveform digitizer ASICs to shape, sample, and digitize analog signals from the SiPMs. FPGAs process the ASIC data, reducing it to spatial coordinates and energy measurements. A processor builds the final list of (\mathbf{r}_i, E_i) values for the interactions associated with each gamma ray, then uses these to perform *Compton reconstruction*, constraining the gamma-ray photon’s initial trajectory to a ring of the form illustrated in Figure 1.

FPGA Pipeline Prototype. ADAPT’s digitizer ASICs continuously sample the output voltages from the SiPM preamplifiers and store the values in a ring buffer with an analog memory depth of ≥ 256 entries. When a gamma ray is detected, all ASICs are triggered simultaneously to digitize and read out these values. Given the speed at which the gamma-ray photon moves within the detector, all of its interactions are captured in a single readout, and cannot be temporally disambiguated. We refer to the collection of data from a single gamma ray’s interactions as an “event.”

ADAPT will produce around 200 KB–1 MB of raw data for a single event (APT will produce $\sim 100\times$ more), and during a burst,

several thousand events trigger every second (APT will trigger $\sim 10\text{--}100\times$ faster). To sustain these high data rates, a hierarchy of FPGAs reduces the data for later processor-based handling. In previous work, we described an initial prototype FPGA pipeline specified with high-level synthesis (HLS) [21].

Our initial prototype first performs **pedestal subtraction**, removing the capacitive charge pedestal inherent to the ASIC’s analog memory cells from the digitized readouts to yield the true sampled signal values. Each memory cell has a unique pedestal value.

Next, to infer the number of photons captured by each fiber or edge detector’s SiPMs, a **signal integration** stage sums over digitized output values. ADAPT’s digitizers sample 2.5 ns windows; capturing a significant portion of the scintillation time therefore requires integrating over tens or hundreds of samples. Our implementation computes multiple integrals with configurable bounds to capture different components of the waveform. **Zero-suppression** then sets integral values under a configurable threshold to zero; sufficiently low values may be caused by various noise effects in the circuit, and are assumed not to correspond to “real” photons.

To infer interaction coordinates from the array of integrated signal values, **island detection** and **centroiding** take the intensity-weighted mean of WLS fiber positions over islands of adjacent non-zero channels. In our original prototype, island detection merely counted islands, rather than identifying their boundaries. We assumed this was acceptable because ADAPT does not perform trajectory reconstruction for events with more than two interactions in a single layer, as these are more likely caused by background radiation in the earth’s upper atmosphere [10]. Under this simplifying assumption, centroiding can take the weighted mean across *all* channels. If island detection identifies only a single island, the result is correct, since all channels outside the island have zero signal and therefore do not contribute. Otherwise, the result is discarded.

In [21], we describe several HLS-based optimization techniques for this pipeline, including the use of pragmas specified to the AMD XILINX Vitis HLS platform, wide vector operations, parallelization via loop unrolling, and a dataflow pipelined architecture. We achieved a throughput of $> 2 \times 10^5$ events per second in simulation, even with a conservative 100 MHz system clock.

Recent FPGA Developments. Our prototype pipeline had limitations that we have since worked to address. First, our prototype was designed to handle a *single* event, data from which was assumed to already reside in block memory (BRAM) on the FPGA. Under this assumption, registers were allocated for important header information, e.g., the offset between logical and physical addresses. However, in a truly pipelined design, new events may arrive while old ones are being processed, resulting in race conditions on these values. To address this, we implemented a finite state machine (FSM) **packet handler** to read out the serial interface from the digitizer ASICs. The FSM packages relevant header information into a data structure that is passed among each pipeline stage via a sequence of connecting FIFOs. It also vectorizes the received sample values across 16 adjacent readout channels, pushing a vector for each logical sample number into a FIFO for pipelined processing in the subsequent stages, beginning with pedestal subtraction.

The packet handler architecture provides additional flexibility, allowing us to easily swap or handle multiple types of waveform

digitizers. ADAPT is expected to demonstrate the capabilities of two digitizer ASICs: the ALPHA [13], developed by collaborators at the University of Hawai‘i at Mānoa, and the HDSoc from Nalu Scientific, both based on the TARGET ASIC [5]. Although similar in design, the ALPHA and HDSoc use different packet header formats and send channel/sample data in different orders. By designing a unified vectorization scheme and metadata format for the downstream pipeline, we merely need to swap out packet handler implementations to support both ASICs.

The output of the signal integration stage is a digital value reflecting the total charge captured during the integration window. To convert this value to the corresponding number of photons captured, we have implemented a **gain correction** stage in the pipeline. This multiplies the integrated value by a per-channel fractional gain to obtain the photon count. It then subtracts the expected *dark count* for the duration of the integration window. Due to thermally-generated electrons a SiPM will spontaneously emit an output current equivalent to the one generated by an incident photon; these should be subtracted to obtain the real photon signal. The dark count rate is a property of the SiPM, its input voltage, and its temperature; it can be estimated [19] or calibrated per channel.

We have also combined **island detection and centroiding** into one functional module that supports multiple islands. It makes a single scan over the fiber channel array, using registers to track whether it is currently in an island, and whether it has just entered or left an island. While in an island, registers maintain the island’s beginning channel position, width, and total intensity, while a multiply-accumulate circuit tracks the intensity-weighted position. When leaving an island, a single divide computes the weighted mean to get the interaction position, and the centroid data is pushed into a FIFO to be transmitted via Ethernet to the processor.

Compton Reconstruction. From the WLS centroids and edge detector signal intensities, we build a list of interactions, or *hits*, (r_i, E_i) for each individual event. From these, we reconstruct the gamma-ray photon’s trajectory to constrain the burst’s direction in a process referred to as *Compton reconstruction* [6].

For a gamma-ray photon that scatters following an interaction with an electron, the Compton law gives the relationship between the cosine η of the photon’s scattering angle and its energy before and after the interaction. Given the vector c between its two interactions and this η value, we can constrain the gamma ray’s source direction s to a circle projected on the unit sphere, as illustrated in Figure 1. Spatial and energy measurement errors spread the circle into a ring, or *annulus*; by propagation of error, we can estimate the uncertainty $d\eta$ in its radius [18].

Reconstruction is challenging because the list of hits is temporally unordered due to the gamma ray’s speed-of-light travel within the instrument. We therefore use the methods described in [6] to infer the most likely ordering of hits. We developed an accelerated branch-and-bound tree search algorithm [18, 19] that achieves a throughput of around 3×10^5 events per second when utilizing all 4 cores of ADAPT’s flight computer [20].

3 Point-Source Localization

Localization aims to determine the most likely source direction for a GRB using the Compton rings inferred by reconstruction.

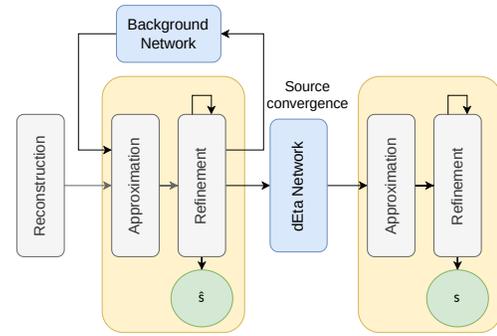


Figure 2: ADAPT GRB localization pipeline (from [12]).

Approach. Localization fixes a GRB’s source direction s by “intersecting” multiple Compton rings. In principle, three rings suffice to fix s ; however, we must contend with both the uncertainties $d\eta$ of each ring and the fact that many observed rings ($\geq 50\%$ for ADAPT) arise from unrelated, diffuse *background radiation*.

As described in [11], ADAPT’s localization operates in two stages. The first stage, *approximation*, considers source directions that lie on one of a random sample of reconstructed rings and selects a direction s_0 that maximizes the joint likelihood of the sample. The second stage, *refinement*, uses an iterative least-squares approach to adjust s_0 until it converges to a maximum-likelihood estimate s given all the rings. To improve the robustness of refinement, each iteration of least-squares fitting excludes as background any rings that do not pass near the current estimate s_j of the source direction.

Use of Machine Learning. Our recent work [12] supplements ADAPT’s localization pipeline with machine learning inference designed to address background noise and uncertainty estimation. We introduce two multilayer perceptron models, the *background network* and the *dEta network*. The first model classifies a Compton ring as originating from either the GRB or the background, allowing background suppression; the second more accurately estimates $d\eta$ for the surviving Compton rings. Each model takes as input the energy and position estimates of interactions that gave rise to the ring, along with an estimate of the GRB’s polar angle with respect to the detector z -axis.

Providing polar angle as an input proved essential to the accuracy of our ML models; however, it required altering the flow of localization as shown in Figure 2. We now iterate between the basic localization computation — which produces an estimated source direction \hat{s} that includes its polar angle — and relabeling the input rings using the background network, discarding those that are labeled as background. Once the estimate \hat{s} converges, or after a predetermined number (currently 5) of iterations, we re-estimate $d\eta$ for all surviving rings and obtain s with a final run of the core algorithm. This iterative approach both increases overall predictive accuracy and allows us to trade off between accuracy and efficiency.

Validation. Figure 3 summarizes experiments from [12] that measure the impact of our machine-learning additions to localization. These results measure the accuracy with which ADAPT can localize simulated 1-second GRBs with brightness 1 MeV/cm^2 . Using ML consistently improved localization accuracy, both in the common

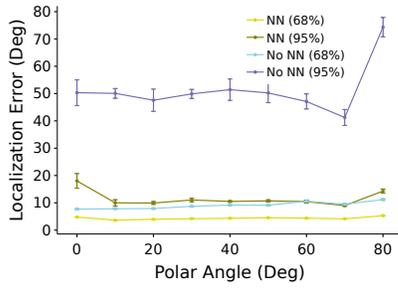


Figure 3: Localization accuracy vs. polar angle for old (“No NN”) and new (“NN”) pipeline (from [12]).

case (68%ile error) and especially for outliers (95%ile error). We predict that across all polar angles, ADAPT can localize GRBs with fluence $\geq 1 \text{ MeV/cm}^2$ to within 6° of error at least 68% of the time.

We also measured the computational cost of localization on a single-board computer with a quad-core, 1.92 GHz Intel Atom E3845 processor that will fly on ADAPT. Reconstruction and localization for a 1 MeV/cm^2 , normally-incident burst required around 220 ms. The time spent applying the neural networks was roughly comparable to the time spent in approximation and refinement.

We are currently investigating how to implement the architecture of Figure 2 as a streaming computation. A streaming approach could produce a series of increasingly accurate localizations over time, continuously communicating estimates to a follow-up optical instrument as it decides how to best to search for the GRB.

4 Likelihood Mapping

The techniques of Section 3 identify one likely location for a GRB in the sky but do not communicate the likelihoods of alternate locations. Observing partners demand this additional information to help them decide whether to commit their narrow-FoV telescopes to search for an optical counterpart to a GRB, and if so, how to prioritize possible locations during the search. The information is communicated in the form of a *likelihood map*, as illustrated in Figure 4, that scores each possible sky location as to the likelihood that it contains the GRB. Here, we describe mapping methods we are developing for ADAPT and APT.

Approach. Our mapping computation follows that of the *cosipy* library [14] released as part of the planned Compton Spectrometer and Imager (COSI) mission, which has been used for multiple prior

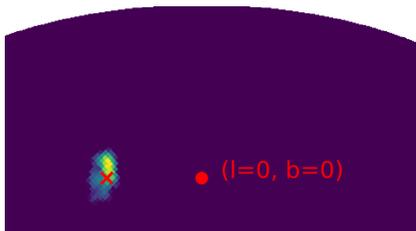


Figure 4: Partial likelihood map for GRB 140329295 from Fermi GBM catalog, showing 99% containment region. Lighter-colored pixels are more likely to contain the GRB source. The red cross denotes the actual source.

gamma-ray instruments. Mapping, like point-source localization, begins with a set D of Compton rings. As in Figure 1, each ring in D has a center vector and radius; it also has an associated E_m , the measured energy of the photon. These three parameters define the *Compton data space* (CDS) of possible rings. The rings in D are assumed to arise either from a source that appears for time Δt at location \mathbf{s} or from background radiation.

The source and background are respectively characterized by an *instrument response* $R(\mathbf{s}, E_i)$ and a *background model* B . R and B each describe the expected number of rings in a given volume within the CDS observed during time interval Δt . R assumes that rings arise from photons of energy E_i arriving from a GRB of unit intensity in source direction \mathbf{s} , while B assumes that they arise from the background. R and B are derived empirically from extensive simulations that capture biases in observed ring parameters due to the telescope’s geometry (e.g., whether shielding blocks photons from some directions), detector electronics, and analysis software, as well as biases inherent to the physics of Compton scattering.

Suppose for simplicity that all GRB photons have fixed energy E_i . To generate a likelihood map, we compare for each source direction \mathbf{s} the hypothesis $H_1^{\mathbf{s}}$ that some portion of D arose from a GRB point source at \mathbf{s} , versus the null hypothesis H_0 that D arose from background alone. For $H_1^{\mathbf{s}}$, the expected number of rings produced in a given volume of CDS is determined by $R(\mathbf{s}, E_i) \cdot \rho + B$, where ρ , the actual intensity of the GRB, is unknown *a priori* and so must be fit to maximize the likelihood. For both $H_1^{\mathbf{s}}$ and H_0 , the observed number of rings within a given volume of CDS is assumed to be Poisson with that volume’s expectation. The map score for direction \mathbf{s} is the log-likelihood ratio of $H_1^{\mathbf{s}}$ versus H_0 given D .

For computational convenience, the CDS, and hence R and B , is discretized along its axes to form voxels; in particular, spatial directions \mathbf{s} and \mathbf{c} are discretized into pixels using HEALPix [9], an equal-area tiling of the sphere. D is discretized into counts in each CDS voxel. Moreover, the energies E_i of incident gamma rays are not in fact fixed but rather follow a distribution, the GRB’s *energy spectrum*. This spectrum must be known or estimated so that we may average over it to obtain the response R for a given source direction. Finally, although a log-likelihood ratio can be computed for any source direction, we report only the *containment region* of all \mathbf{s} for which $H_1^{\mathbf{s}}$ lies within some confidence interval (typically 90 or 99%) of the maximum observed likelihood.

Computational Cost. A key question is whether likelihood map generation, like point-source localization, can be done in real time for short-duration GRBs, so that ADAPT and APT can coordinate with follow-up telescopes seeking optical counterparts that could fade within seconds. The computation is straightforward to parallelize across source directions, and the response R and background B are kept as arrays in DRAM for speed of access. We also implemented multiresolution mapping, in which a map is produced at low resolution for the whole sky and then refined only in areas with likelihood scores high enough to plausibly contain the GRB.

We tested our implementation on 17 simulated transients from the 3rd COSI Data Challenge [8], using COSI’s instrument response R and inferring the background B from three months of simulated observations in low-earth orbit. We generated likelihood maps with 12,288 HEALPix pixels ($\sim 2^\circ$ resolution), limiting output to the

90% containment region for the source. On eight cores of an Intel Xeon Gold 6240 processor, map generation consistently completed in under 700 ms. Further improvements, even for lower-power processors, will arise from porting our implementation to C++ and exploiting a GPU to accelerate linear-algebraic operations.

Challenges. Robust likelihood mapping for ADAPT and APT requires two further advances: real-time inference of GRB spectra, and efficient representation and inference of the response R .

Thus far, we have relied on the ground-truth spectrum provided by COSI to infer a burst’s distribution over E_i . For ADAPT and APT, we must infer this spectrum from observation. Cosipy’s spectral estimation fits a model, such as a power law or a Band spectrum [4], to maximize the likelihood of the observed Compton rings. This approach requires nonlinear optimization and also assumes the GRB’s source location is known, resulting in a circular dependence with map generation. We will investigate simplified real-time fitting approaches that do not assume a known source location.

The size of the instrument response R — several gigabytes for even a low-resolution model — places large demands on memory and data bandwidth. More seriously, computing R for a new telescope such as ADAPT/APT requires casting trillions of photons at a model of the detector using the GEANT4 particle physics simulator [2] to learn the distribution of Compton rings across CDS bins. Ongoing research includes compact machine-learning models that can approximate R . We will also investigate efficient active-learning approaches to infer R from fewer simulated photons.

5 Follow-up Observations

Given a likelihood map of potential source locations, a narrow-FoV telescope must decide how to explore this map to maximize its chance of seeing the source’s optical counterpart. Here, we formulate this problem as a combinatorial optimization. The map is divided into *tiles*, each representing a distinct observing field for the telescope. The goal is to pick a subset of these tiles and an order in which to visit them that maximizes the chance of seeing the source. The solution is constrained by a time *budget* determined by the time until the source’s counterpart is likely to fade below the limit of detection. Both time spent dwelling (observing) at a tile and time spent slewing (moving the telescope between tiles) count towards this budget.

Problem Formulation. We model our problem as a variant of the *Budgeted Prize-Collecting Traveling Salesman Problem (BPC-TSP)* [17]. In this variant, the observation path is open, i.e., it need not return to the starting point, and the cost of traveling between nodes (tiles) includes both slew and dwell time.

Formally, $G = (V, E)$ denotes a complete directed graph, where each node represents a tile. Node v has a value $L(v)$, the likelihood that the source will be found at this tile, and edge $(u, v) \in E$ has cost c_{uv} that is a sum of two parts: $\text{SlewTime}(u, v)$, the time to slew from u to v , and $\text{DwellTime}(v)$, the time needed to observe at tile v to see if the source is present. The graph is symmetric if dwell times are uniform and asymmetric otherwise. We add a dummy node v_0 with $L(v_0) = 0$ to represent the telescope’s initial position. The goal is to select a path P in G to maximize its value $\sum_{v \in P} L(v)$; however, P must satisfy $\sum_{(u,v) \in P} c_{uv} \leq T$, where T is the budget.

Heuristic Algorithms. Because BPC-TSP is an NP-optimization problem, we tried several heuristics to solve it. A simple greedy heuristic is as follows: if the path has reached tile u , select the unvisited tile v that maximizes the ratio $L(v)/c_{uv}$. This method reduces to the well-known greedy knapsack heuristic if all c_{uv} are uniform. Other, potentially more powerful heuristics explore a space of possible paths. These include simulated annealing, genetic, and ant-colony approaches that construct, modify, and/or mix together paths to seek locally optimal solutions.

We also formulated our problem as an integer linear program (ILP), using polynomial-size encoding as for TSP. We then applied branch-and-bound optimization using the Gurobi ILP solver. For instances with few tiles or a limited budget T , Gurobi returns an optimal solution; for larger instances, we constrained it to run with a 0.1% optimality gap and a time limit of 1800 s.

Experimental Setup. To compare the different solution strategies, we tested them on the likelihood map of Figure 4 for GBM140329295. The map gives the 99% containment region for the GRB, an ellipse of roughly 446 deg^2 centered near 0° latitude covering 133 tiles.

The dwell time was assumed to be the same for all tiles. Slew time between tiles depends strongly on the characteristics of the telescope. We considered two models of telescope movement, each of which describes a class of existing telescopes. In the first model (“geodesic movement”), the time to slew between directions \mathbf{s}_i and \mathbf{s}_j is proportional to the angular distance $\arccos(\mathbf{s}_i, \mathbf{s}_j)$ between them. In the second (“two-axis movement”), the telescope’s altitude and azimuth are controlled by two different motors, and the time to slew from \mathbf{s}_i to \mathbf{s}_j is proportional to the maximum of their difference in altitude and their difference in azimuth. In each model, the constant of proportionality is the telescope’s *slew rate*, which can range from less than a degree up to tens of degrees per second.

When choosing parameters for testing, we simplified our parameter space to two quantities:

- **Traversal Budget \mathcal{B}** = Budget \times SlewRate — the total distance the telescope can traverse;
- **Dwell Effort \mathcal{D}** = DwellTime \times SlewRate — the effective observation time.

Given these two quantities, the slew rate uniformly scales node values, edge costs, and time budget T and so is not expected to change the outcome. We considered nine scenarios $(\mathcal{B}, \mathcal{D})$:

S1: (1600, 192), S2: (1600, 128), S3: (1600, 64),
 S4: (3200, 192), S5: (3200, 128), S6: (3200, 64),
 S7: (4800, 192), S8: (4800, 128), S9: (4800, 64).

Larger \mathcal{B} permits more exploration, while a larger ratio of \mathcal{B} to \mathcal{D} reduces dwell times and so permits more total slew time.

Strategies were evaluated based on the fraction of total map likelihood contained in tiles visited within the time budget, and on time needed to compute the path on a 2.3 GHz Intel Xeon Gold 5118 processor.

Results. Figure 5 illustrates the fraction of total map likelihood collected by the path produced by each strategy. As expected, scenarios that restrict the overall traversal budget, or which require more time spent dwelling at each tile, more severely limit the likelihood that can be collected.

For the map tested here, the different strategies collected roughly the same fraction of total likelihood for a given budget, assuming

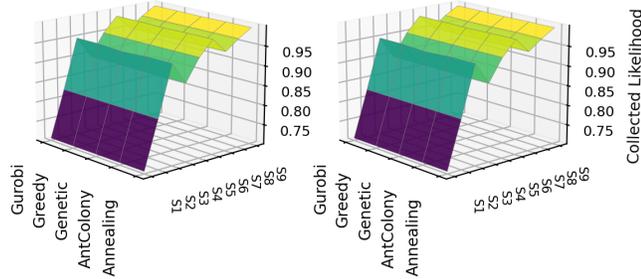


Figure 5: Fraction of map likelihood collected across scenarios and methods. Left surface: geodesic movement; right surface: two-axis movement.

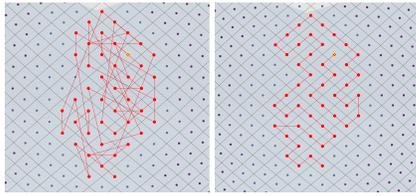


Figure 6: Paths found by Greedy vs. Gurobi for test burst.

either geodesic or two-axis movement. The greedy method was consistently slightly worse than the others, though not dramatically so. We note, however, that while our example models the common case of well-localized GRBs at moderate inclination to the zenith, we have encountered other scenarios with different behavior. In particular, (1) with two-axis movement, bursts centered near the zenith are challenging for simpler heuristics because small changes in azimuth have high movement cost, and (2) if the containment region is not compact but consists of multiple, widely-separated patches in the sky, simpler heuristics are again disadvantaged.

In terms of running time, a greedy solution was typically found within 10^{-4} s. All other approaches required substantially longer — more than 0.1 s for simulated annealing and 1–10 s for other heuristics. Given the large gap in running time between greedy and the other heuristics, there remains ample room to develop computationally efficient approximations or other combinatorial strategies for this problem that perform well even in more complex scenarios. As evidence that improvement on greedy is possible, Figure 6 compares the paths computed by greedy vs. ILP for the case of geodesic movement. We note that the greedy path contains many long edges, including edges that cross other edges, while the ILP-based path (which recovers slightly more likelihood) does not. We conjecture that spanning-tree approaches similar to those used to approximate TSP might yield better solutions in practice.

6 Conclusions and Future Work

This paper describes our ongoing efforts toward enhancing coordinated multi-wavelength observations of astrophysical transients such as gamma-ray bursts (GRBs). The Advanced Particle Astrophysics Telescope (APT) is a mission concept for a space-borne gamma-ray instrument, and we anticipate the launch of its Antarctic Demonstrator (ADAPT) on a high-altitude balloon at the end of 2026. Toward prompt optical observations of GRBs, APT and ADAPT will detect, reconstruct, and localize GRBs to a point source

or a likelihood map using onboard computation. Optical telescopes may then search for the GRB, guided by the likelihood map and the search strategies we have presented. Ongoing efforts will improve the algorithms discussed in this paper, better model the instrument response and background, explore tradeoffs between computational speed and accuracy, and evaluate a broader class of likelihood maps.

Acknowledgments

The authors would like to acknowledge the entire APT Collaboration (see <https://adapt.physics.wustl.edu/>). Support was provided by NASA award 80NSSC21K1741, the McDonnell Center for the Space Sciences, the Peggy and Steve Fossett Foundation, and a Washington University OVCR seed grant.

References

- [1] Rasha Abbasi et al. 2024. All-sky Search for Transient Astrophysical Neutrino Emission with 10 Years of IceCube Cascade Events. *ApJ* 967, 1 (2024), 48.
- [2] S. Agostinelli, J. Allison, K. Amako, et al. 2003. Geant4 — a simulation toolkit. *Nucl. Instrum. Methods Phys. Res. A* 506, 3 (2003), 250–303.
- [3] Rodolfo Artola et al. 2020. TOROS optical follow-up of the advanced LIGO-VIRGO O2 second observational campaign. *Monthly Notices of the Royal Astronomical Society* 493, 2 (2020), 2207–2214.
- [4] D. Band et al. 1993. BATSE Observations of Gamma-Ray Burst Spectra. I. Spectral Diversity. *ApJ* 413 (Aug. 1993), 281.
- [5] K Bechtol et al. 2012. TARGET: A multi-channel digitizer chip for very-high-energy gamma-ray telescopes. *Astropart. Phys.* 36, 1 (2012), 156–165.
- [6] SE Boggs and P Jean. 2000. Event reconstruction in high resolution Compton telescopes. *A&AS* 145, 2 (2000), 311–321.
- [7] James H. Buckley, Jeremy Buhler, and Roger D. Chamberlain. 2024. The Advanced Particle Astrophysics Telescope (APT): Computation in Space. In *Proc. of 21st Int'l Conference on Computing Frontiers Workshops and Special Sessions*. 122–127.
- [8] The Compton Spectrometer and Imager (COSI) Collaboration. 2025. *COSI Data Challenges*. doi:10.5281/zenodo.15126188
- [9] Krzysztof M Gorski et al. 2005. HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *ApJ* 622, 2 (2005), 759.
- [10] Ye Htet et al. 2023. Localization of Gamma-ray Bursts in a Balloon-Borne Telescope. In *Proc. of Workshops of the Int'l Conf. on High Performance Computing, Network, Storage, and Analysis*. 395–398.
- [11] Ye Htet et al. 2023. Prompt and Accurate GRB Source Localization Aboard the Advanced Particle Astrophysics Telescope (APT) and its Antarctic Demonstrator (ADAPT). In *Proc. of 38th Int'l Cosmic Ray Conf.*, Vol. 444. 956:1–956:9.
- [12] Ye Htet et al. 2024. Machine Learning Aboard the ADAPT Gamma-Ray Telescope. In *Proc. of Workshops of the Int'l Conf. on High Performance Computing, Network, Storage, and Analysis*. 4–10.
- [13] M. Kuwahara and G. S. Varner. 2023. Design and Development of Advanced Low Power Hybrid Acquisition (ALPHA) ASIC for Antarctic Demonstrator for the Advanced Particle Astrophysics Telescope (ADAPT). In *Proc. of Nucl. Science Symp., Medical Imaging Conf. and Int'l Symp. on Room-Temp. Semicond. Detectors*.
- [14] Israel Martinez et al. 2023. The cosipy library: COSI's high-level analysis software. In *Proc. of 38th Int'l Cosmic Ray Conf.*, Vol. 444. 858:1–858:8.
- [15] Péter Mészáros, Derek B Fox, Chad Hanna, and Kohta Murase. 2019. Multimessenger astrophysics. *Nature Reviews Physics* 1, 10 (2019), 585–599.
- [16] Kohta Murase and Imre Bartos. 2019. High-energy multimessenger transient astrophysics. *Annu. Rev. Nucl. Sci.* 69, 1 (2019), 477–506.
- [17] Alice Paul et al. 2020. Budgeted prize-collecting traveling salesman and minimum spanning tree problems. *Mathematics of Operations Research* 45, 2 (2020), 576–590.
- [18] Marion Sudvarg et al. 2021. A Fast GRB Source Localization Pipeline for the Advanced Particle Astrophysics Telescope. In *Proc. of 37th Int'l Cosmic Ray Conf.*, Vol. 395. 588:1–588:9.
- [19] Marion Sudvarg et al. 2023. Front-End Computational Modeling and Design for the Antarctic Demonstrator for the Advanced Particle Astrophysics Telescope. In *Proc. of 38th Int'l Cosmic Ray Conf.*, Vol. 444. 764:1–764:9.
- [20] Marion Sudvarg et al. 2023. Parameterized Workload Adaptation for Fork-Join Tasks with Dynamic Workloads and Deadlines. In *Proc. of 29th Int'l Conf. on Embedded and Real-Time Computing Systems and Applications*. 232–242.
- [21] Marion Sudvarg et al. 2024. HLS taking flight: Toward using high-level synthesis techniques in a space-borne instrument. In *Proc. of 21st ACM Int'l Conf. on Computing Frontiers*. 115–125.