

Modeling and Optimizing Real-time Telescope Interaction for Multi-wavelength Observation of Gamma-ray Bursts

Ye Htet

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

Marion Sudvarg

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

Honghao Yang

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

Jeremy D. Buhler

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

Roger D. Chamberlain

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

James H. Buckley

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

ABSTRACT

Multi-wavelength observation of gamma-ray bursts (GRBs) requires real-time interaction among multiple telescopes. A gamma-ray telescope detects and localizes a GRB in the sky and must then communicate with an optical telescope to direct the latter toward the GRB as quickly as possible. We previously developed software for ADAPT, a suborbital gamma-ray telescope, to localize GRBs in real time, on a timescale shorter than that of the GRB itself. This work therefore studies progressive localization, in which ADAPT computes a series of increasingly accurate location estimates during a GRB to enable a partner instrument to more rapidly find it. We describe a modeling and optimization framework to decide when ADAPT should compute estimated GRB locations to minimize the time for the partner to find the GRB. Our framework can design progressive strategies that allow a partner telescope to find a GRB up to 42% faster than strategies using a single alert.

CCS CONCEPTS

• **Applied computing** → **Astronomy**; • **Computer systems organization** → **Sensors and actuators**; **Embedded software**.

KEYWORDS

multi-messenger astrophysics, interactive search, performance modeling, streaming computation, real-time systems

ACM Reference Format:

Ye Htet, Marion Sudvarg, Honghao Yang, Jeremy D. Buhler, Roger D. Chamberlain, and James H. Buckley. 2025. Modeling and Optimizing Real-time Telescope Interaction for Multi-wavelength Observation of Gamma-ray Bursts. In *Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25)*, November 16–21, 2025, St Louis, MO, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3731599.3767475>

1 INTRODUCTION

Multi-wavelength and multi-messenger observations of transient, high-energy astrophysical phenomena offer rich opportunities to

probe their underlying physics. Indeed, the “time-domain and multi-messenger program” was identified among the top scientific priorities for astrophysics in the U.S. National Academies’ Astro2020 Decadal Survey [9]. These observations involve cooperation among multiple telescopes that can observe a phenomenon almost simultaneously, either at different electromagnetic wavelengths or by detecting different types of signal from it (e.g., photons, neutrinos, or gravitational waves). However, effectively coordinating these instruments requires intensive computation that must occur on the short timescale dictated by the physical processes of interest. Deciding when and how to perform these computations is essential to maximize the scientific utility of coordinated observation.

In this paper, we consider the detection and localization of gamma-ray bursts (GRBs) by high-energy, omnidirectional instruments to enable follow-up observations by optical telescopes. In particular, we focus on the Antarctic Demonstrator for the Advanced Particle-physics Telescope (ADAPT), a suborbital gamma-ray telescope scheduled for a long-duration, high-altitude balloon flight from Antarctica. The prompt optical counterparts of a short- to mid-duration GRB might be visible for only seconds, so ADAPT must accurately localize the GRB in the sky and must issue urgent alerts to partner instruments to observe it within a short time.

In prior work on ADAPT, we developed a pipeline of algorithms [5, 14, 15] that runs aboard the telescope and can localize a GRB with an apparent energy density of 1 MeV/cm² to within 5–6 degrees in the sky in only a few hundred milliseconds after a one-second exposure. On-board computation avoids communication bandwidth constraints and latency associated with sending raw detector signals to earth for localization, thereby decreasing the time to produce an informative alert.

Because localization runs on a shorter timescale (hundreds of milliseconds) than many GRBs (seconds), ADAPT can compute a series of location estimates during a GRB, each using progressively more observations and hence yielding greater accuracy. These estimates can be immediately communicated to the optical partner as they are computed, improving the partner’s ability to find the GRB through more frequent interaction. Earlier localization estimates, while less accurate, may still enable the partner to begin turning (“slewing”) toward the source immediately, while later, more accurate estimates allow it to course-correct its movement. In the end, the partner can reach its observing target sooner than if it had waited for a single, final result from ADAPT. This paradigm leads



This work is licensed under a Creative Commons Attribution 4.0 International License. *SC Workshops '25, November 16–21, 2025, St Louis, MO, USA*

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

ACM ISBN 979-8-4007-1871-7/2025/11

<https://doi.org/10.1145/3731599.3767475>

to the following key question: *at what time(s) after the beginning of a GRB should location estimates be computed and communicated so as to enable the partner to reach its target soonest while respecting the resource constraints of ADAPT's limited on-board computing hardware?*

In this work, we develop a methodology to decide when ADAPT should compute and communicate progressive localization estimates to a partner instrument during an ongoing GRB. We first develop two types of model using extensive simulation of ADAPT: a *performance model* of the computational cost to localize a GRB, and an *accuracy model* to estimate the accuracy of GRB localization given some number of observations. We use these models to formulate an optimization task that chooses the times at which ADAPT should compute and communicate localization estimates to the partner. We show that progressive localization allows a partner that can slew at $50^\circ/\text{sec}$ to reach the final estimated location of a $1 \text{ MeV}/\text{cm}^2$ GRB up to 42% sooner compared to waiting to perform localization until after the GRB ends. Furthermore, our model-based strategy can improve the utility of interaction among telescopes while limiting ADAPT's computational resource usage.

2 BACKGROUND

Gamma-ray bursts (GRBs) result from cataclysmic events in the deep universe, such as neutron-star collisions and supernovae. Due to their immense distance from earth, they appear as point-source emissions at unpredictable locations in the sky. The gamma-ray signals from short- and medium-duration GRBs may appear for tens of seconds to less than a second. A GRB's brightness, or *fluence*, quantifies the total gamma-ray energy it deposits in a given area of the detector; ADAPT can detect and localize bursts of fluence at least $1 \text{ MeV}/\text{cm}^2$ [5].

2.1 Localizing GRBs in the Sky

The proposed Advanced Particle-astronomy Telescope (APT) and its Antarctic Demonstrator (ADAPT), like prior instruments such as Fermi GBM [8], observe gamma-ray transients whose photon energies lie in the Compton energy regime (0.1–10 MeV); these photons interact with ADAPT's detector primarily via Compton scattering [4]. ASIC and FPGA logic [11, 13] computes for each such interaction its 3D position $\mathbf{r} = (x, y, z)$ within the detector and its deposited energy E . For a single gamma-ray photon, the flight software gathers its list of interactions $(\mathbf{r}_i, E_i)$, collectively referred to as an *event*, and communicates this list to the analysis pipeline via an in-memory event queue.



Figure 1: ADAPT GRB localization pipeline.

As shown in Figure 1, the analysis pipeline consists of two stages: *reconstruction* and *localization*. Reconstruction converts the interaction list for each photon to a constraint on the source's location in the form of a *Compton ring* – a circle on the unit sphere on which the direction vector from ADAPT to the GRB source must lie. A photon's interactions are converted to a Compton ring by inferring

their temporal ordering from energetic and kinematic constraints via the Compton law [2]. In prior work, we developed an accelerated reconstruction algorithm [5, 10, 15] that can process $>10^5$ photons per second on ADAPT's flight computer [12].

Reconstruction continuously accumulates Compton rings in a *Compton ring queue*. At any time, the pipeline can perform *localization* by combining the constraints from all accumulated rings to produce an estimated direction vector to the source. Localization must contend both with uncertainty in the parameters of each inferred Compton ring and with spurious rings caused by atmospheric background radiation unrelated to the GRB. These background rings, which arise from gamma rays scattered by the atmosphere and also from energetic protons, neutrons, and electrons that trigger the detector's electronics (but cannot be distinguished from incident gamma rays by signal alone) can account for $\geq 50\%$, and perhaps as much as 80%, of all rings seen by ADAPT.

As described in [5], ADAPT's localization algorithm first constructs a coarse estimate of the source direction from a random sample of rings in the queue, then performs iterative least-squares refinement of this estimate using all available rings to obtain its final result. We use machine-learning models [6] to select Compton rings that are likely to be from the GRB rather than the background and to estimate the uncertainty in the radius of each ring.

2.2 Multi-wavelength Observation

GRBs may exhibit *optical counterparts* – visible-light signals that appear at the GRB's sky location within at most tens of seconds of the gamma-ray signal. Capturing light from the optical counterpart provides information about the physical phenomena that gave rise to the GRB. Capturing earlier light from the counterpart is more scientifically valuable, so a gamma-ray telescope should direct a partner instrument toward the GRB's location as quickly as possible.

ADAPT can communicate an estimated source location to partner instruments as an *alert* message with negligible (10s of ms) light-speed delay. While ADAPT's detector can survey the entire visible sky at once, an optical telescope typically has a narrow field of view – at most a few degrees, and often $<1^\circ$ – and so must be given a precise source direction to observe. When such a telescope receives a location alert, it immediately slews its detector toward the source. Modern telescope mounts [7] support slewing a telescope in an arbitrary direction at an angular velocity up to tens of degrees per second.

ADAPT and APT differ from prior balloon-borne and satellite gamma-ray instruments in that they are designed to perform localization entirely with on-board computation. On-board analysis eliminates communication bottlenecks associated with transferring raw observations to ground-based computation, enabling localization and alert generation that utilizes all available data on sub-second time scales. This capability makes feasible the progressive localization strategies that we explore in the next section.

3 MODELING AND OPTIMIZATION APPROACH

In this section, we first frame the progressive GRB localization problem as an optimization task, then develop supporting models that characterize the computational cost and accuracy of localization as

a function of the number of input Compton rings. While a simple performance model captures the computational cost of localization, more detailed fitting to empirical observations is required to characterize its accuracy.

3.1 The Progressive Localization Problem

Assume that a GRB occurs over a period of T seconds. ADAPT detects the burst at time 0, when it first becomes bright enough in gamma rays to trigger the detector, and continues to observe it until time T , when the flux of incident gamma rays returns to the baseline level associated with atmospheric background radiation. In the time range $[0, T]$, GRB photons enter the instrument at a rate determined by the (time-dependent) gamma-ray flux. A small fraction of these photons (those in which the full photon energy is absorbed by the detector via interactions at two or more separately-resolved locations) result in a Compton ring. During this period, background particles may also appear and generate spurious Compton rings. At time $t \leq T$, the instrument has seen $n_{\text{src}}(t)$ rings from the GRB source and $n_{\text{bkg}}(t)$ rings from the background.

Over the course of the GRB, ADAPT performs $m + 1$ successive localization computations at times $t_0 \dots t_m$. The last localization is always launched at time $t_m = T$, when the burst ends, and so uses all available data from it. Earlier localizations may be launched at any time $< T$ and use only the set of Compton rings available at the time of launch. A localization launched at time t_i runs for a length of time $t_{\text{loc}}(n_{\text{src}}(t_i), n_{\text{bkg}}(t_i))$, after which it emits *location alert* a_i . Location alerts are direction vectors pointing from ADAPT to the GRB point-source, expressed in polar coordinates (θ, ϕ) denoting co-latitude and longitude, respectively.

Cooperative Observation. ADAPT interacts with an optical telescope, the *partner*, that seeks to observe the GRB's location as soon as possible after it is detected. The partner is initially pointing in direction s_0 at time 0 and must eventually slew to point in direction a_m . We assume that the partner (on the ground) receives location alerts from ADAPT (in the upper atmosphere) with negligible communication delay. Even so, the partner will not receive the final direction a_m until after time T , when ADAPT has processed all available gamma-ray data from the burst.

To exploit ADAPT's progressive alerts, the partner operates as follows. When alert a_0 from the first localization is received, the partner starts to slew toward direction a_0 at a constant angular velocity (the "slew rate"). When the next alert a_1 is received, the partner is pointing in some direction s_1 between s_0 and a_0 ($s_1 = a_0$ if it arrives early); it immediately begins to slew from s_1 toward a_1 . This process repeats for all subsequent alerts; in particular, when the final alert a_m is received, the partner is pointing in direction s_m and must slew to a_m .

Objective and Constraints. The objective of progressive localization is to minimize the time until the partner points to the final computed source direction a_m . If we neglect the possibility (rare in practice) that the partner happens to reach a_m before ADAPT's last, most accurate alert is sent, then the time to reach a_m is determined by the time to move from s_m to a_m . Hence, we seek to minimize the angular distance $\arccos(s_m \cdot a_m)$ (hereafter referred to as the *loss*). The position s_m is itself determined by the slew rate of the partner

(currently fixed at 50° per second for our experiments), and the timing and accuracy of prior alerts from localization runs launched at times $t_0 \dots t_{m-1}$.

While we can control the times at which localization runs, the accuracies of the resulting alerts are not deterministic. Even given a fixed *number* of Compton rings at time t , the resulting location estimate may have greater or lesser error depending on the characteristics of the photons that generated those rings. Hence, the alert locations $a_0 \dots a_m$ are *random variables*, which cause the partner positions $s_1 \dots s_m$, and hence the loss, to be stochastic as well. We estimate the loss associated with a set of localization run times $t_0 \dots t_m$ by performing multiple trials with different simulated photons. Each trial produces a set of alerts $a_0 \dots a_m$ and computes the partner's resulting trajectory and its loss. We seek to minimize the mean loss across all trials.

The localization run times $t_0 \dots t_{m-1}$ are constrained to lie between 0 and T . Moreover, ADAPT's limited computational resources permit only one instance of localization to run at a time; hence, we require that $t_{i+1} \geq t_i + t_{\text{loc}}(n_{\text{src}}(t_i), n_{\text{bkg}}(t_i)) + c$ for $0 \leq i < m$. Here, c is a "slack" value (100 ms in this work) that allows for inaccuracy in estimating t_{loc} .

Physical Inputs. We solve the progressive localization problem for a short-duration GRB that is typical of the bursts for which ADAPT's detector pipeline is optimized. Our test burst lasts $T = 1$ second and exhibits a Band energy spectrum [1] with parameters $\alpha = -0.5$, $\beta = -2.35$, and $E_{\text{peak}} = 490$ keV. The burst generates 44,600 incident gamma-ray photons, corresponding to an overall fluence of 1 MeV/cm^2 for the given spectrum. The incident photon flux varies with time according to a Gaussian curve with $\mu = 508.12$ and $\sigma = 233.12$. In contrast, background particles arrive at a constant rate of 4.48×10^5 per second throughout the burst, with particle types and energies following the background model of [3]. The total incident flux is the sum of the GRB and background flux.

The characteristics of the Compton rings obtained from incident gamma rays varies with the actual direction of the burst, due to the detector's geometry and algorithmic biases in event processing and reconstruction. This variation impacts the timing and accuracy of localization. We therefore tested our approach for three representative source directions: normal to the detector ($\theta = 0^\circ$), as well as $(\theta = 30^\circ, \phi = 45^\circ)$ and $(\theta = 60^\circ, \phi = 45^\circ)$. In what follows, we refer to these three cases by their respective θ angles. To ensure that the partner instrument is required to slew in all cases, we set s_0 equal to $(60^\circ, 45^\circ)$ for $\theta = 0^\circ$ and to $(0^\circ, 0^\circ)$ for $\theta \in \{30^\circ, 60^\circ\}$ as shown in Figure 2.

To generate incident particles and simulate their interactions with the detector, we used the GEANT4 particle physics simulator together with a digital twin of the detector hardware [11]. Simulated interactions were fed to ADAPT's computational pipeline to produce reconstructed Compton rings, which were then used for localization as described in [5].

3.2 Modeling the Computational Cost of Reconstruction and Localization

We modeled the cost of running ADAPT's analysis pipeline, including both reconstruction and localization, on ADAPT's flight instrument computer, which includes a quad-core Intel Atom E3845




Figure 2: Three representative scenarios.





Figure 3: Accumulation of Compton rings vs. time.

CPU running at 1.92 GHz. Reconstruction and localization can be parallelized across multiple CPU cores.

We measured the cost t_{recon} of reconstruction as a function of the number of incoming events. Reconstruction received, over 300 trials each simulating a one-second burst, an average of $N = 31,746$ events caused by gamma-ray photons and background particles that produced at least one detectable interaction in the detector; the number of events exhibited minimal variation (1 s.d. = 154). With one core, the average time $t_{\text{recon}}(N)$ to reconstruct all these events was 134 ms, with a maximum observed time of 140 ms. Because one core proved more than sufficient to reconstruct incident photons in real time at the flux dictated by our test burst, we chose a simplified architecture in which one core was dedicated to reconstruction, while the other three were used for localization. Future work will consider alternative resource partitions, particularly for the case of more intense GRBs (up to 10-fold brighter than our test burst) where reconstruction of incident photons could become a bottleneck.

Only a small fraction of incident photons yield Compton rings, due to computational filters that discard low-quality or unreconstructable signals. Figure 3 illustrates the number of Compton rings accumulated from our test burst and from different types of background particles over 300 random trials. The average number of rings accumulated in one second was 581 with limited variance (s.d. = 24).

We measured the cost $t_{\text{loc}}(n_{\text{src}}, n_{\text{bkg}})$ of localization for values of n_{src} and n_{bkg} in the range $[10..500]$, which covers the range of values seen in our simulations. We used a two-dimensional input,

Figure 4: Localization time t_{loc} vs. number of source and background Compton rings.

rather than just the total number of rings $n_{\text{src}} + n_{\text{bkg}}$, because localization converges faster when given a higher fraction of rings from the GRB source vs. the background. For each set of parameters, we timed 100 localization trials with randomly generated photons. We then fit the median of these running times for each set of tested parameters using a radial basis function (RBF) interpolator with a linear kernel, so as to allow estimation of times for parameters that were not tested. Figure 4 illustrates these fits assuming different numbers of cores are dedicated to localization; using three cores exhibits substantial speedup over using just one core. We used the three-core configuration in all subsequent modeling and optimization.

To determine computational cost as a function of the time t at which a localization run starts, we estimated the number of Compton rings available at time t from the curve fits in Figure 3, then applied the RBF model of Figure 4 to estimate t_{loc} . Although the number of rings available at a given time t is actually stochastic rather than fixed, its variance is small enough that the value inferred from the RBF fit is a useful estimate in practice.

3.3 Modeling Localization Alert Accuracy

To estimate the accuracy of a localization run starting at time t , we must (1) extract the numbers of GRB and background Compton rings present at time t from the arrival curves of Figure 3, and (2) repeatedly simulate incident photons to generate this many rings, process them through our pipeline, and use the results to estimate the accuracy of localization. However, the cost of simulation is high and would be incurred anew for each value of t tested. To reduce the cost of estimating accuracy, we developed empirically fitted accuracy distributions as a proxy for simulation.

For each of the $n_{\text{src}}, n_{\text{bkg}}$ combinations tested in Figure 4, we performed 500 trials of localization with sets of Compton rings derived from simulated photons, recorded the resulting alert directions (θ, ϕ) for each trial, and then fit parametric distributions to each of θ and ϕ (assuming these values were uncorrelated). To estimate these distributions for input ring counts that were not explicitly simulated, we linearly interpolated the parameters of the distributions for the nearest simulated counts.

Before fitting, θ is normalized to between 0 and 1, while ϕ (in radians) is shifted by adding π , then normalized to between 0 and 1. Empirically, we found that using a beta distribution (α, β with fixed $\text{loc} = 0$, $\text{scale} = 1$) for θ and a shifted generalized normal distribution ($\beta, \text{loc}, \text{scale}$) for ϕ worked best for bursts at 30° and 60° . However, for the burst at 0° , the empirical values of θ were multimodal, with a peak close to the expected $\theta = 0^\circ$ and a second, much more dispersed, peak at larger values of θ . To improve fitting, we divided the observed density of θ into two regions using a two-component Gaussian mixture model, then separately fit each peak using a beta distribution limited to its region of support.

Examples of these fits – for proportions of GRB/background rings similar to those produced by our test burst with true location $\theta = 0^\circ$ – are shown in Figure 5. The estimated value of θ becomes more sharply concentrated near its true value as the number of GRB-derived Compton rings available to localization increases. The estimated value of ϕ , whose true value is ill-defined when $\theta = 0^\circ$, remains near 45° .

4 EXPERIMENTAL RESULTS

To assess the utility of progressive localization and of our modeling and optimization approach, we used our methods to predict the losses associated with different progressive strategies for the test burst and three source directions described in Section 3.1.

We explored the space of progressive strategies in two ways. First, we varied the number of intermediate localization runs prior to the final run (at $T=1$ sec) between 1 and 4, fixing the times for these runs at evenly spaced times in $[0, T]$ as shown in Table 1. These strategies did not attempt to perform optimization of the times. Second, we implemented the constrained optimization of Section 3.1 using a nonlinear optimizer – SciPy’s implementation of dual annealing [16]. Constraints were as described above; we further constrained all localizations to start at $t \geq 150$ ms, since earlier times had limited numbers of rings and hence very low accuracy. For each candidate strategy (i.e., set of time points) considered by the optimizer, we used our computational cost model to infer t_{loc} for each time point to verify that two localizations do not overlap in time, and computed the strategy’s average loss over 500 trials with

source directions sampled from our accuracy model for each time point. The times for localization chosen by the optimizer for each true source direction and each number of intermediate localizations are shown in Table 2.

m	1	2	3	4
	500	333, 666	250, 500, 750	200, 400, 600, 800

Table 1: Baseline time points (ms).

m	True Source Direction		
	$\theta = 0^\circ$	$\theta = 30^\circ$	$\theta = 60^\circ$
1	385	465	198
2	166, 573	418, 735	150, 412
3	150, 395, 610	400, 598, 762	150, 291, 514
4	151, 320, 457, 619	231, 442, 598, 805	150, 282, 445, 786




Table 2: Optimized time points (ms) for each source direction.

To assess the validity of our models’ predictions, we compared the predicted losses computed for each strategy to “ground truth” losses computed by (much more expensive) direct simulation. For each strategy, the ground truth loss was the average computed over 100 trials, each of which simulated a set of input gamma rays, ran the analysis pipeline to compute the set of Compton rings available at each time point of interest, and finally computed the loss given this set of rings. We estimated the uncertainty in the modeled and simulated averages by computing them multiple times (30 times for modeled loss, 5 times for the more expensive ground-truth loss) with different random seeds and computing the 95% confidence t -test interval on the average of these meta-trials.

Figures 6, 7, and 8 compare the ground truth (“GT”) and modeled losses for unoptimized (“evenly split”) and optimized strategies for the test burst with each of the three source directions. In general, modeled losses did not perfectly reproduce the ground-truth losses, which is not surprising given the approximations made by our accuracy model. A more practical question is whether the models’ guidance about which progressive strategy to use matches that obtained from the more expensive ground-truth estimates.

One form of guidance is the best value for m , the number of intermediate alerts to generate. According to the ground truth estimates for the unoptimized strategy, $m = 3$ suffices for $\theta = 0^\circ$, $m = 2$ for $\theta = 30^\circ$, and $m = 4$ for $\theta = 60^\circ$; these are the largest m for which the error bars for m and $m - 1$ do not overlap. The modeled losses provide similar guidance overall, except for favoring $m = 4$ for $\theta = 0^\circ$.


Another form of guidance is whether the optimized strategies are expected to outperform the baseline. The ground-truth estimates indicate that the optimized strategy does *not* meaningfully outperform the unoptimized one for $\theta \in \{0^\circ, 30^\circ\}$ but does outperform it for $\theta = 60^\circ$. In contrast, the modeled losses predict that the optimized strategy is better in all cases, but the predicted differences in loss are much smaller for $\theta \in \{0^\circ, 30^\circ\}$ than for $\theta = 60^\circ$. This suggests that the modeled losses provide useful qualitative guidance for strategy design after accounting for some distortion introduced by modeling. Moreover, optimization does substantially reduce the loss in at least one scenario ($\theta = 60^\circ$). Future work will identify the range of source directions for which optimization is beneficial.

Figure 5: Examples of accuracy models inferred for test burst with source location $\theta = 0^\circ$.Figure 6: Losses for source at $\theta = 0^\circ$.Figure 7: Losses for source at $\theta = 30^\circ$.

In all scenarios, progressive localization was able to substantially reduce the loss compared to the value obtained by a non-progressive strategy that waits for the final alert a_m . Figure 9 shows the distributions of losses for optimized strategies with $m = 4$ for each true source direction tested. The mean improvements relative to the loss of a non-progressive strategy (60° for $\theta \in \{0^\circ, 60^\circ\}$, or 30° for $\theta = 30^\circ$) were between 2- and 5-fold. Assuming that the partner instrument can slew at $50^\circ/\text{sec}$ (the upper limit of modern telescope mounts), the implied reduction in time to reach the estimated GRB location is 0.75-1 seconds, as shown in Figure 10. In particular, for $\theta = 60^\circ$, the implied time reduction is 1s out of 2.4s, or 42%.

5 CONCLUSIONS AND FUTURE WORK

The ADAPT and APT telescopes will perform source localization and alert generation in real time via on-board computations. The low latency of these computations opens up new possibilities for

Figure 8: Losses for source at $\theta = 60^\circ$.




Figure 9: Distribution of ground-truth losses for optimized strategies with $m = 4$.




Figure 10: Latency improvement (ground-truth values) of optimized progressive strategies with $m = 4$ vs. non-progressive strategy. We assume the non-progressive strategy receives an alert at $t = 1.2s$ and travels to the true target (within 5° of the final alert).

how to use them in cooperative observation of high-energy transient phenomena at multiple wavelengths. We have studied progressive localization, in which the gamma-ray telescope interacts repeatedly with a fast-slewing partner instrument in real time during a GRB to help the latter turn more swiftly toward the GRB's estimated location. We have modeled the computational cost and accuracy of progressive alerts using ADAPT's on-board computing hardware and developed strategies using these models that optimize the times at which localizations are computed. The models circumvent the need for expensive simulations to assess different progressive strategies while providing useful guidance for strategy design. Finally, we have quantified the benefit of progressive localization in several scenarios.

Future work will further develop and assess the utility of progressive localization. We will improve our modeling approach, in particular the accuracy models, to more closely match simulation. Replacing parametric models with a more general fitting approach such as kernel density estimation should more robustly reproduce

the observed features of the accuracy distribution. We will also extend our validation to consider a wider range of source directions and a broader set of GRB parameters. In particular, we will consider brighter GRBs for which one CPU core may not be sufficient to handle reconstruction of all incoming photons; in such a scenario, we may need to consider different strategies for partitioning the available hardware between reconstruction and localization. We anticipate that ADAPT may need to choose dynamically among a library of precomputed strategies based on the observed gamma-ray flux and energy spectrum. Finally, we will improve our model of interaction between ADAPT and its partner to consider longer delays and more realistic models of telescope movement involving acceleration and settling time.

ACKNOWLEDGMENTS

This work was supported by NASA award 80NSSC21K1741 and by a Washington University seed grant.

REFERENCES

- [1] D. Band, J. Matteson, L. Ford, B. Schaefer, D. Palmer, B. Teegarden, T. Cline, M. Briggs, W. Paciesas, G. Pendleton, et al. 1993. BATSE observations of gamma-ray burst spectra. I-Spectral diversity. *Astrophysical J., Part 1, vol. 413, no. 1*, p. 281-292. 413 (1993), 281–292.
- [2] S. E. Boggs and P. Jean. 2000. Event reconstruction in high resolution Compton telescopes. *Astronomy and Astrophysics Suppl. Series* 145, 2 (2000), 311–321. <https://doi.org/10.1051/aas:2000107>
- [3] W. Chen et al. 2023. Simulation of the instrument performance of the Antarctic Demonstrator for the Advanced Particle-Astrophysics Telescope in the presence of the MeV background. In *Proc. 38th Int'l Cosmic Ray Conf.*, Vol. 444. 841:1–841:9. <https://doi.org/10.22323/1.444.0841>
- [4] A. H. Compton. 1923. A quantum theory of the scattering of X-rays by light elements. *Phys. Rev.* 21 (May 1923), 483–502. Issue 5. <https://doi.org/10.1103/PhysRev.21.483>
- [5] Y. Htet et al. 2023. Prompt and accurate GRB source localization aboard the Advanced Particle Astrophysics Telescope (APT) and its Antarctic Demonstrator (ADAPT). In *Proc. 38th Int'l Cosmic Ray Conf.*, Vol. 444. 956:1–956:9. <https://doi.org/10.22323/1.444.0956>
- [6] Y. Htet, M. Sudvarg, D. Butzel, J. D. Buhler, R. D. Chamberlain, and J. H. Buckley. 2024. Machine Learning Aboard the ADAPT Gamma-Ray Telescope. In *Proc. Wkshps. Int'l Conf. High Performance Computing, Network, Storage, and Analysis (SC-W)*. 4–10. <https://doi.org/10.1109/SCW63240.2024.00008> Presented at 5th Workshop on Artificial Intelligence and Machine Learning for Scientific Applications (AI4S), Atlanta, GA, USA.
- [7] J. H. Kim, M. Im, H. Lee, S.-W. Chang, H. Choi, and G. S. H. Paek. 2024. Introduction to the 7-Dimensional Telescope: commissioning procedures and data characteristics. In *Ground-based and Airborne Telescopes X*, Vol. 13094. International Society for Optics and Photonics, SPIE, 130940X. <https://doi.org/10.1117/12.3019546>
- [8] C. Meegan, G. Lichti, P. N. Bhat, et al. 2009. The Fermi Gamma-Ray Burst Monitor. *Astrophysical J.* 702, 1 (Aug. 2009), 791–804. <https://doi.org/10.1088/0004-637x/702/1/791>
- [9] National Academies of Sciences Engineering and Medicine. 2023. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. The National Academies Press, Washington, DC, USA. <https://doi.org/10.17226/26141>
- [10] M. Sudvarg et al. 2021. A fast GRB source localization pipeline for the Advanced Particle-Astrophysics Telescope. In *Proc. 37th Int'l Cosmic Ray Conf.*, Vol. 395. 588:1–588:9. <https://doi.org/10.22323/1.395.0588>
- [11] M. Sudvarg et al. 2023. Front-end computational modeling and design for the Antarctic Demonstrator for the Advanced Particle-Astrophysics Telescope. In *Proc. 38th Int'l Cosmic Ray Conf.*, Vol. 444. 764:1–764:9. <https://doi.org/10.22323/1.444.0764>
- [12] M. Sudvarg et al. 2023. Parameterized Workload Adaptation for Fork-Join Tasks with Dynamic Workloads and Deadlines. In *Proc. 29th Int'l Conf. Embedded and Real-Time Computing Systems and Applications (RTCSA)*. 1–10. <https://doi.org/10.1109/RTCSA55878.2022.00007>
- [13] M. Sudvarg et al. 2025. FPGA-Based Data Processing using High-Level Synthesis on the Antarctic Demonstrator for the Advanced Particle-Astrophysics Telescope (ADAPT). In *Proc. 39th Int'l Cosmic Ray Conf.*, Vol. 501. 852:1–852:9.
- [14] M. Sudvarg, C. Zhao, Y. Htet, M. Konst, et al. 2024. HLS taking flight: Toward using high-level synthesis techniques in a space-borne instrument. In *Proc. 21st ACM Int'l Conf. Computing Frontiers*. 115–125.

- [15] J. Wheelock, W. Kanu, M. Sudvarg, Z. Xiao, J. D. Buhler, R. D. Chamberlain, and J. H. Buckley. 2021. Supporting Multi-messenger Astrophysics with Fast Gamma-ray Burst Localization. In *Proc. IEEE/ACM HPC for Urgent Decision Making Workshop (UrgentHPC)*. 21–28. <https://doi.org/10.1109/UrgentHPC54802.2021>.
- 00008
- [16] Y. Xiang, D. Y. Sun, W. Fan, and X. G. Gong. 1997. Generalized simulated annealing algorithm and its application to the Thomson model. *Physics Letters A* 233, 3 (1997), 216–220. [https://doi.org/10.1016/S0375-9601\(97\)00474-X](https://doi.org/10.1016/S0375-9601(97)00474-X)