Call for Collaboration: Contributing to Multi-Messenger Astrophysics

Marion Sudvarg*†, Ye Htet*, Roger D. Chamberlain*, Jeremy D. Buhler*, James H. Buckley†
*Department of Computer Science and Engineering, †Department of Physics
Washington University in St. Louis
{msudvarg, htet.ye, roger, jbuhler, 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-astrophyics 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. In this presentation, we describe the current status of on-instrument computations for the Antarctic Demonstrator for APT (ADAPT) pursuant to guiding prompt follow-up observations of transient events. We also describe open problems and the challenges of extending ADAPT's computation to the future APT instrument. We encourage contributions and collaboration from members of the real-time systems community.**

I. Introduction

Background and Motivation. The astrophysics community has a strong interest in observing transient astrophysical phenomena using multiple modalities. This *multi-messenger* approach may include, e.g., gravitational waves, electromagnetic waves, neutrinos, and cosmic rays [1], [2]. Because these transients can be short-lived [3], fast detection and localization is key to supporting cooperative multi-modal observation.

The Advanced Particle-astrophysics Telescope (APT) [4] 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 second [5].

Nonetheless, APT may supplement point localization with more detailed *likelihood maps* that provide a spatial probability distribution over possible locations. This is especially important for its Antarctic Demonstrator (ADAPT), which has greater localization uncertainty. The most likely regions of the map can then be searched by fast-slewing optical telescopes to localize a source for subsequent observations.

Factors contributing to delays between detection and secondary observations include time to localize/map the transient, communication latency to follow-up telescopes (which is particularly challenging for space-based instruments like APT), and the time for these telescopes to physically search the sky as directed by a likelihood map. The computational pipeline that transforms raw sensor data into localizations (see Fig. 1) must therefore execute on the instrument (in space) [4] and meet significant size, weight, and power (SWaP) constraints.

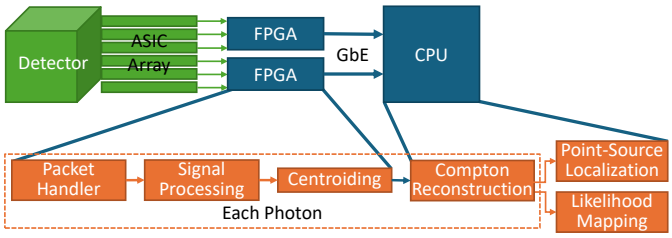

Fig. 1. ADAPT's computational pipeline for GRB localization.

Current Status. This presentation describes several important computational elements associated with APT and their ongoing development for ADAPT. §II outlines our progress on improved trajectory reconstruction for individual gamma rays, a process which uses FPGAs to read out and process signals from digitizer ASICs. From resulting estimates of the positions and energies of interactions in the detector, each gamma-ray photon's initial trajectory is constrained to a ring in the sky. §III details our method for point-source localization using the resulting collection of rings. While our original method based on least-squares refinement is effective in simulation for APT, ADAPT's smaller size and exposure to atmospheric background radiation give rise to greater uncertainty. To address this, we augment our iterative approach with machine learning. Furthermore, in §IV, we also plan to generate likelihood maps of the GRB's location; these will then be transmitted to optical telescopes for subsequent physical search.

Open Problems. Challenges persist in developing our computational pipeline for ADAPT, and there remain open problems related to deployment on the future APT instrument, which is larger, has more sensors, and demands higher readout rates. In §II, we motivate the need for efficient, FPGA-based noise-suppression and photon-counting algorithms. The future APT mission will need improved capabilities to transmit result data from several dozen FPGAs for aggregation on a CPU. In §III, we describe ADAPT's iterative ML-based reconstruction and localization loop, and connect it to existing work on concurrent, real-time execution of neural network models and recent work on IDK cascades [6]. In §IV, we discuss opportunities for

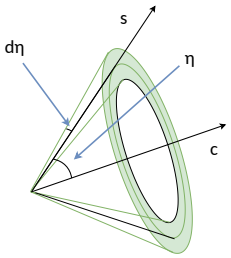

Fig. 2. A reconstructed Compton ring (from [11]).

tradeoffs between computational workload and precision of the probability maps that ADAPT will generate and consider the implications when portions of the spectral fitting and mapping algorithms are offloaded to GPU or FPGA accelerators, which may need to be shared with mission- or safety-critical recurrent tasks. For all of these problems, we welcome contributions and collaboration with other research groups in the real-time systems and operating systems communities.

Attribution. Many of the details of ADAPT's current status closely follow an invited paper at this year's CompSpace special session at the Computing Frontiers conference [7].

II. Trajectory Reconstruction

The upcoming ADAPT and proposed APT telescopes infer a GRB's direction by combining the trajectories of individual gamma-ray photons that interact with them. Here we give an overview of the instrument designs and methods to reconstruct gamma-ray trajectories from their raw sensor data.

The APT and ADAPT Detectors are constructed with layers of scintillating tiles that emit visible light when incoming gamma-ray photons scatter within them. This light is first 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 their ends [8], [9]. This overlay of 1-dimensional fiber arrays into a 2-dimensional mesh, with the relative position of the tile, allows us to resolve the 3-dimensional position $\mathbf{r} = (x, y, z)$ of each 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 will fly with onboard computational hardware, including waveform digitizer ASICs to sample and digitize analog signals from the SiPMs [10]. FPGAs process the ASIC data, reducing it to spatial coordinates and energy measurements. A processor builds the final set of interactions $(\mathbf{r}_i, E_i)$ 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 Fig. 2.

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 $\geq$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.

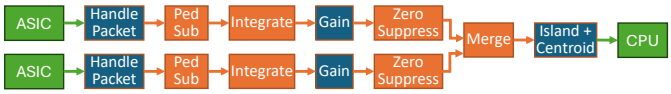

Fig. 3. ADAPT's FPGA pipeline.

We refer to the collection of data from a single gamma ray's interactions as an "event." To handle streaming event arrivals, a finite state machine (FSM) **packet handler** reads out the serial interface from the digitizer ASICs, providing flexibility to handle multiple types of waveform digitizers. ADAPT is expected to demonstrate the capabilities of two ASICs: the ALPHA [12], developed by collaborators at the University of Hawai'i at Mānoa, and the HDSoC from Nalu Scientific, both based on the TARGET ASIC [13].

Our pipeline first performs **pedestal subtraction**, removing the unique capacitive charge pedestal inherent to each of the ASIC's analog memory cells from the digitized readouts to yield the true sampled signal values. To infer the number of photons captured by each fiber or edge detector's SiPMs, a **signal integration** stage sums over digitized output values. To estimate the number of photons captured, a **gain correction** stage multiplies the integrated value by a per-channel fractional gain, then subtracts the expected *dark count* (spontaneous SiPM impulses due to thermally-generated electrons) for the duration of the integration window. **Zero-suppression** then sets sufficiently low photon counts to zero, under the assumption that these values may be caused by noise effects in the circuit or variation in dark counts.

Zero-suppressed photon counts from multiple ASICs are merged into a single array for each WLS fiber plane. To then derive interaction coordinates from each array, **island detection and centroiding** take the mean of WLS fiber positions, weighted by photon intensity, over islands of adjacent non-zero channels. The complete pipeline is illustrated in Fig. 3.

In [14], we described several HLS-based optimization techniques for an earlier version of this pipeline, achieving a throughput of $>2 \times 10^5$ events per second in simulation, even with a conservative 100 MHz system clock. New (or modified, in the case of island detection and centroiding) components since that work are marked in dark blue in Fig. 3.

Compton Reconstruction. Using a CPU, we next build a set of interactions, or *hits*, $(\mathbf{r_i}, E_i)$ 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* [15]. For a gamma-ray photon that scatters following an interaction with an electron, the Compton law gives the relationship between the cosine η of its scattering angle and its energy before and after the interaction. Given the vector $\mathbf{c}$ between its two interactions and this η value, we can constrain the gamma ray's source direction $\mathbf{s}$ to a circle projected on the unit sphere, as illustrated in Fig. 2. 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 [16], [17].

Reconstruction is challenging because the set of hits is temporally unordered due to the gamma ray's speed-of-light travel within the instrument. We therefore use the methods described

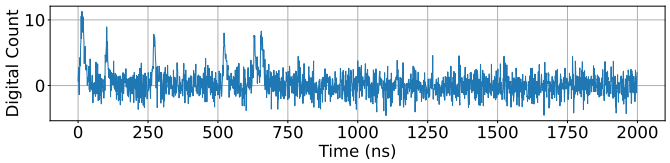

Fig. 4. 2 μs digitized waveform from an optical fiber that captured 5 photons.

in [15] to infer the most likely ordering of hits. We developed an accelerated branch-and-bound tree search algorithm [10], [16] that achieves a throughput of around 3×10^5 events per second when utilizing all 4 cores of ADAPT's 1.92 GHz Intel Atom 3845 CPU-based flight computer [18].

Open Challenge: Photon Counting. The energy deposited by a scattering gamma ray is low enough that only a handful of optical photons (typically $<$20) are transmitted by each optical fiber. For example, Fig. 4 shows a simulated waveform read out from a fiber; of the six peaks, five correspond to captured photons, while one is a dark count. Photon counts inferred from signal integration are thus inaccurate given the low SNR. Alternative FPGA-based techniques for noise suppression and photon counting — e.g., combinations of threshold-based methods [19], filters, and deep-learning based approaches — remain an open challenge. In particular, we would like to explore combinations of such methods, including synthesis of multiple kernels with different combinations of filters to be dynamically swapped out depending on resource availability and latency/throughput requirements.

Open Challenge: Data Rates. ADAPT will produce around 200 KB–1 MB of raw data for a single event; during a burst, several thousand events trigger every second. To sustain these high data rates, a hierarchy of FPGAs implements the preprocessing pipeline to handle thousands of readout channels. Ultimately, 13 FPGAs transmit reduced data to a CPU via Gigabit Ethernet. The larger APT instrument will produce $\sim$100$\times$ more data per event while triggering $\sim$10–100$\times$ faster. We anticipate that at least 60 FPGAs will send data, though perhaps not all of them for every event. Addressing this challenge within the constraints imposed by the space-based computational environment may require Time-Sensitive Networking (TSN). We will alternatively consider a dedicated FPGA with access to unified CPU memory or cache fabric that performs event building and dispatches Compton reconstruction workloads to CPU cores directly, building upon the principles of CAESAR [20].

III. Point-Source Localization

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

Approach. Localization fixes a GRB's source direction $\mathbf{s}$ by "intersecting" multiple Compton rings. In principle, three rings suffice to fix $\mathbf{s}$; 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 [5], ADAPT's localization operates in two stages. The first stage, *approximation*, selects the most likely direction $\mathbf{s_0}$ from a set of candidates. The second stage, *refinement*, uses an iterative least-squares approach to adjust $\mathbf{s_0}$ until it converges to a maximum-likelihood estimate $\mathbf{s}$.

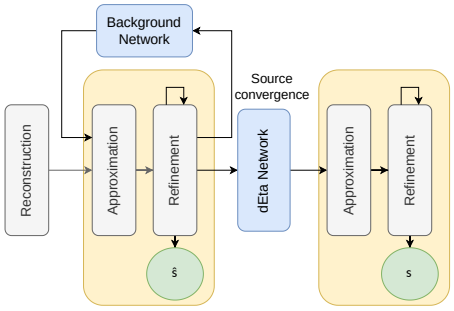

Fig. 5. ADAPT GRB localization pipeline (from [11]).

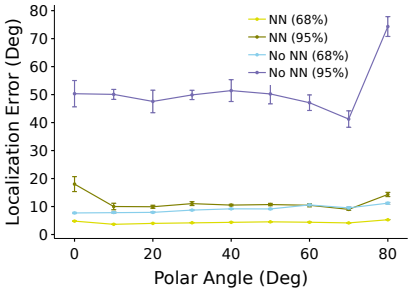

Fig. 6. Localization accuracy vs. polar angle with/without ML (from [11]).

Use of Machine Learning. Our recent work [11] 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 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.

Using the polar angle as an input has proved essential to model accuracy, but it is not known a priori. We therefore iterate between the basic localization computation — which produces an estimated source direction $\hat{\mathbf{s}}$ — and then discarding any input rings classified as background given $\hat{\mathbf{s}}$. Once the estimate $\hat{\mathbf{s}}$ converges, we re-estimate $d\eta$ for all surviving rings and obtain $\mathbf{s}$ with a final run of the core algorithm.

Validation. Fig. 6 summarizes experiments from [11] that measure the impact of machine learning. These results measure the accuracy with which ADAPT can localize simulated GRBs. Using ML consistently improved localization accuracy, both in the common case (68%ile error) and especially for outliers (95%ile error). We also measured the computational cost on ADAPT's flight computer. Reconstruction and localization for a representative bright, short burst required $\sim$220 ms.

Open Problem Area: ML in RT As evidenced by the recent ML-RT Agenda (ECRTS'24 and '25) and WMC (RTSS '24) workshops, there is growing interest in using machine learning safely and predictably in real-time systems. Our iterative approach to ML-based GRB localization allows us to trade off between accuracy and efficiency, but exploring this

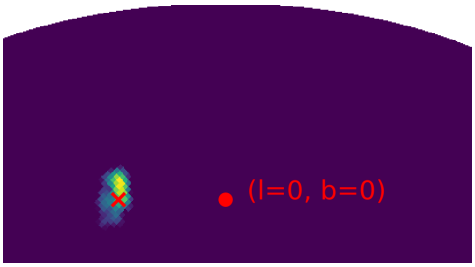

Fig. 7. Partial likelihood map for GRB 140329295 from Fermi GBM catalog, showing 99% containment region (from [7]). Lighter-colored pixels are more likely to contain the GRB source. The red cross denotes the actual source.

tradeoff space is an open problem, especially when the neural network models (which can exploit intra-task parallelism and vectorization) must run concurrently with each other and with safety-critical instrument- or satellite-control tasks. Several of these concerns have been outlined in a recent paper by Buttazzo [21]; we encourage members of the community with expertise in these areas to collaborate. Moreover, there may be opportunity to use IDK classifiers [6] for background rejection. Effective scheduling for sequences of multiple IDK classifiers (cascades) has garnered recent attention [22]–[25].

IV. Likelihood Mapping

The techniques of §III identify one likely direction for a GRB. To allow telescopes to search for an optical counterpart, we will also communicate a *likelihood map* over its possible location in the sky, as illustrated in Fig. 7.

Approach. Our mapping computation follows that of the `cosipy` library [26] released for the planned Compton Spectrometer and Imager (COSI) mission. Mapping, like point-source localization, begins with a set D of Compton rings, each with a center vector, radius, and measured energy E_m. These parameters define the *Compton data space* (CDS) of possible rings, which 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.

To generate a likelihood map, we compare for each source direction $\mathbf{s}$ the hypothesis H_1^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^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^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^s versus H_0 given D.

Computational Cost. A key question is whether likelihood map generation 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 events, and the response R and background B are kept as arrays in DRAM for speed of access. We also implemented multiresolution mapping [27], [28], 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. Selecting an appropriate granularity of subdivision to provide sufficient map precision to optical telescopes, while not introducing undue delays in the transmission of the map, remains an open problem.

We tested our implementation on 17 simulated GRBs from the 3rd COSI Data Challenge [29], 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 an 8-core Arm Cortex-A78AE (Nvidia Jetson Orin NX), map generation consistently completed in under 200 ms. On 4 performance cores on ADAPT's Intel Core i7-13700TE CPU, it completes in under 100 ms. Further improvements may arise from porting our implementation to C++ and exploiting GPU or FPGA acceleration.

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. `Cosipy`'s spectral estimation fits a model to maximize the likelihood of the observed Compton rings, which requires nonlinear optimization. It 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. 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.

V. Open Problems and Call for Collaboration

As we continue to develop the upcoming ADAPT instrument, work toward the proposed APT mission, and coordinate with ground-based optical telescopes for real-time follow-up observations of astrophysical transients, several open problems remain. These include (but are not limited to) development of efficient FPGA-based signal processing algorithms; real-time transmission of data from dozens of FPGAs to a CPU; computational offloading of ML models and likelihood mapping to GPU or FPGA accelerators; efficient representation and inference of the instrument response R; and runtime platforms to coordinate, schedule, and execute these with timing guarantees. The real-time systems community, with expertise in running latency-constrained applications atop SWaP-constrained hardware, is particularly suited to addressing these problems. We invite interested researchers to collaborate!

ACKNOWLEDGMENT

The authors would like to acknowledge the entire APT Collaboration (see https://adapt.physics.wustl.edu/). Support 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] P. Mészáros, D. B. Fox, C. Hanna, and K. Murase, "Multi-messenger astrophysics," *Nature Reviews Physics*, vol. 1, no. 10, pp. 585–599, 2019.

[2] K. Murase and I. Bartos, "High-energy multimessenger transient astrophysics," *Annu. Rev. Nucl. Sci.*, vol. 69, no. 1, pp. 477–506, 2019.

[3] R. Abbasi *et al.*, "All-sky search for transient astrophysical neutrino emission with 10 years of IceCube cascade events," *ApJ*, vol. 967, no. 1, p. 48, 2024.

[4] J. H. Buckley, J. Buhler, and R. D. Chamberlain, "The Advanced Particle-astrophysics Telescope (APT): Computation in space," in *Proc. of 21st Int'l Conference on Computing Frontiers Workshops and Special Sessions*, May 2024, pp. 122–127.

[5] Y. Htet *et al.*, "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, Jul. 2023, pp. 956:1–956:9.

[6] T. Trappenberg and A. Back, "A classification scheme for applications with ambiguous data," in *Proc. of IEEE-INNS-ENNS Int'l Joint Conference on Neural Networks*, vol. 6, Jul. 2000, pp. 296–301.

[7] D. Wang, Y. Htet, M. Sudvarg, R. Chamberlain, J. Buhler, and J. Buckley, "Coordinating instruments for multi-messenger astrophysics," in *Proc. of 22nd ACM Int'l Conference on Computing Frontiers Workshops and Special Sessions*, May 2025, pp. 213–218.

[8] W. Chen *et al.*, "The Advanced Particle-astrophysics Telescope: Simulation of the Instrument Performance for Gamma-Ray Detection," in *Proc. of 37th Int'l Cosmic Ray Conference*, vol. 395, 2021, pp. 590:1–590:9.

[9] ——, "Simulation of the Instrument Performance of the Antarctic Demonstrator for the Advanced Particle-astrophysics Telescope in the Presence of the MeV Background," in *Proc. of 38th Int'l Cosmic Ray Conf.*, vol. 444, Jul. 2023, pp. 841:1–841:9.

[10] M. Sudvarg *et al.*, "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, Jul. 2023, pp. 764:1–764:9.

[11] Y. Htet *et al.*, "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*, Nov. 2024, pp. 4–10.

[12] M. Kuwahara and G. S. Varner, "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*, 2023.

[13] K. Bechtol *et al.*, "TARGET: A multi-channel digitizer chip for very-high-energy gamma-ray telescopes," *Astropart. Phys.*, vol. 36, no. 1, pp. 156–165, 2012.

[14] M. Sudvarg *et al.*, "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*, 2024, pp. 115–125.

[15] S. Boggs and P. Jean, "Event reconstruction in high resolution Compton telescopes," *A&AS*, vol. 145, no. 2, pp. 311–321, 2000.

[16] M. Sudvarg *et al.*, "A Fast GRB Source Localization Pipeline for the Advanced Particle-astrophysics Telescope," in *Proc. of 37th Int'l Cosmic Ray Conf.*, vol. 395, Jul. 2021, pp. 588:1–588:9.

[17] Y. Htet *et al.*, "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*, Nov. 2023, pp. 395–398.

[18] M. Sudvarg *et al.*, "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*, 2023, pp. 232–242.

[19] L. Huang, "New Signal Identification Algorithms for Enhanced Gamma-Ray Burst Detection in the Advanced Particle-Astrophysics Telescope," Master's thesis, Dept. of Electrical and Systems Engineering, Washington University in St. Louis, May 2025.

[20] S. Roozkhosh, D. Hoornaert, and R. Mancuso, "CAESAR: Coherence-aided elective and seamless alternative routing via on-chip FPGA," in *Proc. of Real-Time Systems Symposium (RTSS)*. IEEE, 2022, pp. 356–369.

[21] G. Buttazzo, "Toward predictable AI-based real-time systems," *Real-Time Systems*, pp. 1–16, 2025.

[22] T. Abdelzaher, K. Agrawal, S. Baruah, A. Burns, R. I. Davis, Z. Guo, and Y. Hu, "Scheduling IDK classifiers with arbitrary dependences to minimize the expected time to successful classification," *Real-Time Systems*, vol. 59, no. 3, pp. 348–407, 2023.

[23] S. Baruah, A. Burns, R. I. Davis, and Y. Wu, "Optimally ordering IDK classifiers subject to deadlines," *Real-Time Systems*, vol. 59, no. 1, pp. 1–34, 2023.

[24] S. Baruah, A. Burns, and R. I. Davis, "Optimal synthesis of robust IDK classifier cascades," *ACM Transactions on Embedded Computing Systems*, vol. 22, no. 5s, pp. 150:1–150:26, 2023.

[25] S. Baruah, I. Bate, A. Burns, and R. I. Davis, "Optimal synthesis of fault-tolerant IDK cascades for real-time classification," in *Proc. of 30th Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2024, pp. 29–41.

[26] I. Martinez *et al.*, "The cosipy library: COSI's high-level analysis software," in *Proc. of 38th Int'l Cosmic Ray Conf.*, vol. 444, 2023, pp. 858:1–858:8.

[27] P. Fernique, M. Allen, T. Boch, A. Oberto, F. Pineau, D. Durand, C. Bot, L. Cambrésy, S. Derriere, F. Genova *et al.*, "Hierarchical progressive surveys: Multi-resolution HEALPix data structures for astronomical images, catalogues, and 3-dimensional data cubes," *Astronomy & Astrophysics*, vol. 578, p. A114, 2015.

[28] L. P. Singer and L. R. Price, "Rapid Bayesian position reconstruction for gravitational-wave transients," *Phys. Rev. D*, vol. 93, p. 024013, Jan 2016.

[29] The Compton Spectrometer and Imager (COSI) Collaboration, "COSI Data Challenges," Apr. 2025. [Online]. Available: https://doi.org/10.5281/zenodo.15126188