The Advanced Particle-astrophysics Telescope (APT): Computation in Space

James H. Buckley
Jeremy Buhler
Roger D. Chamberlain

James H. Buckley, Jeremy Buhler, and Roger D. Chamberlain, "The Advanced Particle-astrophysics Telescope (APT): Computation in Space," in *Proc. of 21st ACM International Conference on Computing Frontiers Workshops and Special Sessions*, May 2024, pp. 122-127.
DOI: 10.1145/3637543.3652980

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

The Advanced Particle-astrophysics Telescope (APT): Computation in Space

Invited Paper

James H. Buckley
Dept. of Physics
Washington University in St. Louis
St. Louis, Missouri, USA

Jeremy Buhler
McKelvey School of Engineering
Washington University in St. Louis
St. Louis, Missouri, USA

Roger D. Chamberlain
McKelvey School of Engineering
Washington University in St. Louis
St. Louis, Missouri, USA

ABSTRACT

The Advanced Particle-astrophysics Telescope (APT), which is being designed for a Sun-Earth Lagrange L_2 orbit, and its balloon-borne Antarctic Demonstrator (ADAPT) represent substantial challenges in computational sensing. Since communications with earth-bound computing resources are intermittent and are constrained to low data rates, the bulk of the computational science must be accomplished aboard the instrument. We describe the computational requirements for APT and our current plan for achieving those requirements through a combination of custom hardware, FPGAs, and embedded processor cores.

CCS CONCEPTS

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

KEYWORDS

Gamma-ray astronomy, multi-messenger astrophysics

ACM Reference Format:
James H. Buckley, Jeremy Buhler, and Roger D. Chamberlain. 2024. The Advanced Particle-astrophysics Telescope (APT): Computation in Space: Invited Paper. In *Proceedings of the 21st ACM International Conference on Computing Frontiers Workshops and Special Sessions (CF '24 Companion), May 7–9, 2024, Ischia, Italy.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3637543.3652980

1 INTRODUCTION

Multi-messenger astrophysics seeks to learn about the universe by combining information derived from the electromagnetic spectrum, gravitational waves, neutrinos, and cosmic rays [6, 24, 26]. For transient events such as neutron-star mergers, prompt localization of the event is crucial, enabling narrow field-of-view (FoV) instruments to point at it effectively. The Advanced Particle-astrophysics Telescope (APT) [10] is a planned instrument that will have close to a full omnidirectional FoV with a goal of transient event detection and localization to less than 1° of uncertainty. Additional science goals include helping determine the nature of dark matter and detecting rare ultra-heavy cosmic-ray nuclei.

Presently, the primary omnidirectional observational modality is the set of gravity-wave detectors, LIGO [1, 2] and Virgo [8]. These instruments' localizations have fairly large uncertainties (greater than 20°), which makes electromagnetic follow-up observation difficult [5]. APT seeks to reduce localization uncertainty sufficiently for follow-up observation to be straightforward.

The APT mission presents substantial challenges for computational sensing. First, it is being designed to be in a Sun-Earth Lagrange L_2 orbit, which implies that the data rate to earth is limited. As such, it is infeasible to send raw observation data to earth for processing, and so much of the computational science must be accomplished on the instrument. Second, because it will be operating in space, there are severe limits on size, weight, and power for the computing system. Third, because APT is charged with detecting transient events, the latency of the computational science pipeline must be low enough so as to effectively inform narrow-FoV instruments for follow-up observations.

The APT Collaboration team is currently designing and building a balloon-borne prototype instrument, the Antarctic Demonstrator for APT (ADAPT), that is scheduled for flight in the 2025-26 season. The team comprises over 45 individuals across 10 organizations, representing 3 countries, 6 states, and the District of Columbia. Sensing and computation include custom ASICs, FPGAs, and embedded processors, which are being designed in Bari, Italy; Honolulu, Hawaii, USA; and St. Louis, Missouri, USA.

The mission's sensing and computation operations include numerous components, as listed below. For each of these, we describe the algorithmic challenges, current planned approach, and future investigations.

- *Raw signal acquisition, multiplexing, amplification, and triggering.* These steps are performed using custom ASICs and discrete electronics.
- *Waveform buffering and analog-to-digital conversion.* These operations are executed on a custom ASIC.
- *Digital signal preprocessing, including pedestal subtraction, integration, zero suppression, island detection, and centroiding.* These operations are executed on FPGAs [35].
- *Event building.* These operations are allocated to both FPGAs and on-board embedded processor cores.
- *Compton reconstruction and localization.* These operations are executed using on-board embedded processor cores on ADAPT [17]; however, parts of them might utilize GPUs or FPGAs on APT [38].

CF '24 Companion, May 7–9, 2024, Ischia, Italy

ACM ISBN 979-8-4007-0492-5/24/05.
https://doi.org/10.1145/3637543.3652980

- *Gamma-ray burst identification and notification.* These operations execute using on-board embedded processor cores.

In addition, the instrument must distinguish between gamma rays that Compton-scatter versus those that generate pair-production events; identify and classify cosmic rays; and reject confounding background noise sources that obscure the signals of interest. All of the above is in addition to the required flight control software, which will be executing in on-board embedded processor cores.

Once data have been returned to Earth, there is still analysis that must be performed. The APT Collaboration has recently engaged with the COSI Collaboration (the Compton Spectrometer and Imager [36]) to investigate acceleration of data analysis techniques for diffuse and weak gamma-ray sources.

2 INSTRUMENT DESCRIPTION

The core sensing components of the instrument are illustrated in Figure 1. Figure 1a shows four layers (for ADAPT) of 3 by 3 tiles of CsI:Na scintillator, each of which has associated wavelength-shifting (WLS) fibers. These components comprise ADAPT's *calorimeter*. Figure 1b shows how the optical light that results from an energy deposition in an individual calorimeter layer propagates and is detected. Square wavelength-shifting fibers oriented in orthogonal directions above and below each CsI:Na layer (225 fibers in each of the x- and y-orientations) collect light that is then detected using silicon photomultipliers (SiPMs). These signals provide positional information for the energy deposition.

On the edge of the crystal opposite the fiber SiPMs are additional *edge-detector* SiPMs that assist in the determination of the energy associated with each deposition. Also incorporated are four more layers of CsI:Na, called *tail counters*, that only include edge detectors and whose purpose is to increase the likelihood that gamma rays are fully absorbed and their energy accurately measured.

Associated with each calorimeter layer are scintillating-fiber tracker hodoscope layers. The tracker layers are each composed of 1.5 mm round scintillating fibers with two interleaved layers for both x- and y-coordinate determination.

The performance both of ADAPT's calorimeter and of the instrument as a whole have been extensively modeled in simulation [3, 12, 14], and substantial portions have also been characterized in the lab [18]. Figure 2 shows the computational pipeline that forms the majority of the signal processing tasks on the telescope, focusing on the calorimeter. The figure shows a single channel; however, the parallel component counts are articulated in the figure caption. Each element of the pipeline will be described in the sections that follow.

3 SIGNAL ACQUISITION AND AMPLIFICATION

SiPMs at the ends of the optical fibers and at the layer edges convert individual photons to electrical current pulses [29]. Discrete components multiplex fiber signals from each of the three tiles in a layer, summing them into one signal path. This reduces the 225 fiber signals in each orientation on each layer into 75 channels/orientation/layer for data acquisition.

Current pulses are converted to voltage signals using a SMART trans-impedance pre-amplifier [4]. The SMART chips each have 16 parallel amplification channels, and they are adjustable in SiPM bias voltage, gain, bandwidth, and pole-zero network filter for tail suppression.

While the responses of these elements are linear, they each contribute electrical noise into the signal pipeline. The SiPMs have dark counts that are a strong function of chip temperature, and (as is typical) the input-referenced electrical noise on the SMART chips is bounded by the gain-bandwidth product of the first amplification stage. We provide quantitative models of these elements in [33].

Triggering to recognize when a photon interacts with the calorimeter is accomplished by forming an analog sum of multiple channels, utilizing the previously developed CT5TEA ASICs [15] which are used in the Cherenkov Telescope Array [39]. For channel sums above a configurable threshold, a trigger signal is generated. While optical signal levels are strong enough for sufficiently energetic gamma-ray interactions to give reliable triggers, lower-energy gamma-ray interactions pose a greater challenge for triggering. This is an area of ongoing investigation.

4 BUFFERING AND A/D CONVERSION

Signals from the calorimeter fibers, edge-detectors, tail counters, and tracker fibers are continuously saved as an analog waveform on a series of switched-capacitor sample-and-hold circuits that operate as a ring buffer. Storing the signal in analog form requires dramatically less power than converting it to a digital signal. Upon receipt of a trigger, the triggered channel(s) undergo analog-to-digital conversion (with 12-bit resolution) and communication to FPGAs for signal analysis. This technique has previously been deployed in terrestrial telescopes using the TARGET ASIC [7] and in particle physics experiments using the DRS4 ASIC chip [30].

For ADAPT, an ALPHA ASIC chip that performs buffering and A/D conversion is currently in testing. The ALPHA supports 16 signal channels per chip. Given that the ALPHA chip is a new design, currently undergoing testing, we need a backup plan in the event that ALPHAs are not available in time for instrument fabrication. Alternative chips that will be considered should the ALPHA not be ready include the above-mentioned TARGET chip and its derivative HDSoC ASIC chip [25].

5 DIGITAL SIGNAL PREPROCESSING

The task of digital signal preprocessing is to take the raw waveforms from A/D conversion and, for each trigger event, discern the location and magnitude of the associated gamma ray's energy depositions. Power efficiency in the computation is crucial, as there is a significant quantity of raw data to be examined. The preprocessing steps have been assigned to three FPGAs per detector layer to minimize power consumption and ensure consistent execution times. Two FPGAs handle the x- and y-dimensions for the layer, respectively, while a third FPGA combines data from the first two and communicates detected events to the CPU via a network link.

The digital signal preprocessing steps are indicated in yellow in Figure 2, including: pedestal subtraction, integration, zero suppression, island detection, centroiding, and event building [33]. Pedestal subtraction is required to normalize the sampled data relative to individual analog sample-and-hold cells in the ALPHA chips. Other

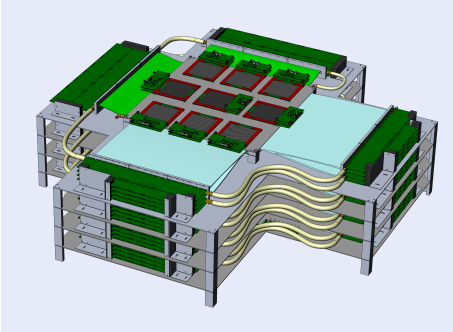

(a) The ADAPT detector stack.

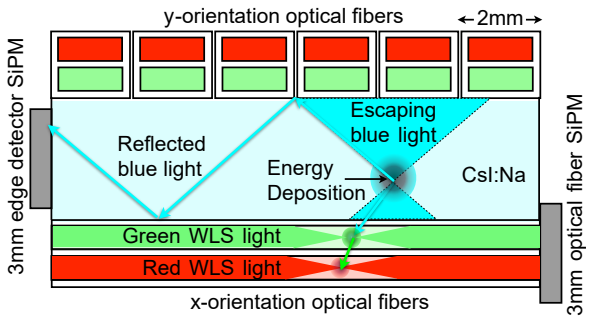

(b) A single calorimeter layer, illustrating optical light propagation and detection [10].

Figure 1: The ADAPT instrument sensor.

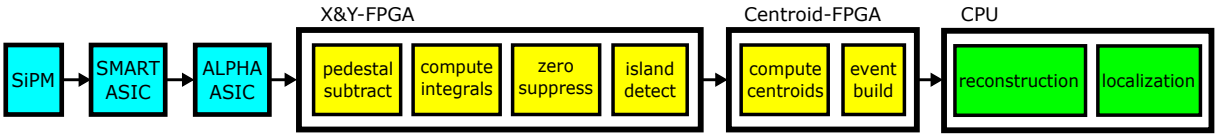

Figure 2: Computational pipeline for APT and ADAPT. Each calorimeter layer has 1 Centroid-FPGA, 2 X&Y-FPGAs (one in X, one in Y), 12 ALPHA chips, and 12 SMART chips. Reconstruction and localization use one multicore CPU for all layers.

than being careful about indexing (so that we are subtracting the correct pedestal value), this is a straightforward operation.

We compute four time integrals on the sampled waveforms (computed as sums, given that the data have already been converted to discrete time bins). This allows us to separately consider the pre-trigger pedestal, the initial rising waveform, the tail, and the complete waveform. While the ultimate use we will make of these separate integrals is still under consideration, the initial rising waveform appears to be sufficient to determine the position of an energy deposition (i.e., its x- or y-coordinate), while the complete waveform is the current best candidate to quantify the energy deposited. Currently, the z-coordinate is simply determined by which layer observes the energy deposition; however, we are exploring a refinement of the position estimate by employing the ratio of light collected above and below the scintillator crystal (in a layer's x and y fibers).

The next three steps have current implementations [35] but are undergoing investigation to see if alternative approaches yield better science. Zero suppression is accomplished by comparing the complete waveform integral to a threshold; only those channels that are above threshold are sent downstream for further processing. Island detection is responsible for identifying contiguous groups of fibers that are likely to have originated from a single energy deposition. It is currently restricted to recognizing a single island in each dimension (either x or y), with expansion to recognizing multiple distinct hits being the subject of ongoing work. Centroiding performs a dot product between the fiber position information and the signal values from an island.

We are investigating a range of algorithms for these three stages. We are considering approaches to zero suppression that require a higher-than-threshold set of samples prior to integration (possibly implemented as a threshold compared to a tight-time integral). Also, we are examining an alternative approach to merge island detection and centroiding using one or more matched filters; this approach is potentially more resilient to dark counts in the photomultipliers and electronic noise in the analog front end.

Independent of the algorithms ultimately chosen for zero suppression, island detection, and centroiding, the event building step is responsible for aggregating all of the data associated with a single gamma-ray event, packaging it together, and sending it to the processor cores for reconstruction and localization. This bookkeeping function is the responsibility of the Centroid-FPGA.

6 COMPTON RECONSTRUCTION AND LOCALIZATION

Given the stream of events (individual gamma-ray photons) emerging from digital signal preprocessing, we seek to identify gamma-ray bursts (GRBs), which manifest as a large number of photons emanating from a common point source in the sky over a short period. When a GRB occurs, we must pinpoint its location in the sky rapidly enough to notify secondary instruments for follow-up observations. ADAPT and APT target bursts of relatively low brightness (< 1 to a few MeV/cm^2 incident light energy over roughly a second), which are likely to be visible to follow-up instruments for only seconds or minutes; hence, minimizing reporting latency is an important design goal. Moreover, the narrow FoV of most follow-up instruments demands highly accurate localization — ideally to within 1° or less — of the burst to avoid wasting time scanning the sky to find it.

We divide the determination of a burst's location into two stages: *reconstruction* of individual gamma-ray photon tracks in the detector, each of which constrains the direction of the source with

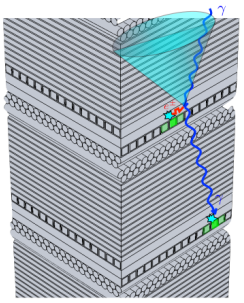

Figure 3: A Compton ring defined by two interactions of a gamma-ray photon with the calorimeter. The vector connecting the interactions determines the ring's orientation, while its angular radius depends on the photon's total energy and the energy deposited by the first scattering [10].

respect to the detector, and *localization* of the burst by combining information from multiple photon tracks.

Reconstruction. A Compton-regime photon entering the detector interacts one or more times in the calorimeter layers, depositing energy and changing direction each time it scatters, until it finally loses all its energy and is absorbed. Centroiding yields a list of these interactions, with deposited energy and spatial coordinates for each, which is *unordered* because interactions happen too fast to resolve temporally. Reconstruction therefore searches through all possible photon trajectories (i.e., interaction orderings) to find one for which the angle of each scattering best agrees with the amount of energy it deposits as predicted by the well-known Compton Law. Given the first and second interactions and the total energy of the incident photon (summed over all interactions), its source in the sky may be restricted to a *Compton ring* of known angular radius centered on a line through the two interactions, as shown in Figure 3.

Our reconstruction algorithm is based on a statistical method [9] that accounts for uncertainties in the measured energy and position of each interaction. We reorganized this method from a simple iteration over permutations into a tree search [32] that eliminates redundant computation and permits rapid pruning of whole groups of ill-fitting trajectories at once. This approach is effective for picking a single most plausible trajectory, and hence a single Compton ring, for photons that scatter at least twice before being absorbed by the detector. For photons that scatter only once before absorption, the Compton Law alone cannot distinguish the two possible orderings, so we report both possibilities to the subsequent localization stage.

Each Compton ring emitted by reconstruction includes an estimate $d\eta$ of uncertainty propagated from the uncertainties in the positions and energies of the scatterings in its photon's trajectory. Due to experimental noise, the true source direction $\mathbf{s}$ of the photon is assumed to lie *near*, but not necessarily on, its ring. For a ring centered on vector $\mathbf{c}$ with angular radius θ, the cosine of the angle between $\mathbf{s}$ and $\mathbf{c}$ is treated as a Gaussian random variable with mean $\eta = \cos\theta$ and standard deviation $d\eta$.

Localization. To localize a GRB, we combine Compton rings from hundreds or thousands of detected photons to find a direction (i.e., the GRB source $\mathbf{s}$) in which they intersect. The problem is similar to that of multilateration in navigation, in which a point is inferred from a large number of noisy range measurements (here, the angular radii of the rings). In our application, instrument noise, incomplete observations of photon trajectories, and the ambiguity of photons with two interactions cause at least half of all inferred rings not to pass near the actual GRB source; hence, localization must be highly resilient to noisy input.

Our localization algorithm [17, 32] first uses a randomized approach to identify an initial candidate source direction that lies near a large fraction of the input rings. We then refine this initial guess using a least-squares approach equivalent to maximizing the joint likelihood of the source direction given the rings (with their associated uncertainties). To improve noise resilience, we exclude from the least-squares calculation all rings that do not pass sufficiently near the candidate source. Because the set of excluded rings changes as the estimated source direction is updated, we iterate between performing least-squares and discarding rings until the source estimate converges.

Performance. Our reconstruction and localization pipeline targets a low-power multicore processor. To approximate the likely computing power available to us on the APT instrument, we use a Raspberry Pi 3B+ with a 4-core ARM A53 processor at 1.4 GHz. For a GRB producing tens of thousands of photons per second, we can compute a source estimate within a few hundred milliseconds of seeing the photons [17]. Depending on future computational needs, we can also accelerate the highly vectorizable localization stage using a low-power GPU [38].

We validate the accuracy of our methods using detailed physical simulations of incident gamma-ray photons and their interactions with the detector, as well as the instrument's front-end electronics. Based on our most complete model of ADAPT to date [16, 17], Table 1 shows the predicted localization error for both 68% and 95% containment (i.e., the actual source is within the error 68% or 95% of the time). For the full APT instrument, with its greater light-gathering ability, we expect to localize GRBs of similar brightness to within one degree or less.

Table 1: ADAPT Localization error (degrees) [17].

Fluence (MeV/cm^2)	68% containment	95% containment
0.5	6.63 ± 0.02	16.33 ± 1.15
1	2.58 ± 0.04	6.74 ± 0.04
2	1.54 ± 0.03	3.15 ± 0.11
3	1.23 ± 0.02	2.33 ± 0.05
4	1.05 ± 0.02	1.98 ± 0.04

Adoption of Machine Learning. As we develop reconstruction and localization, we are identifying opportunities to replace or augment parts of our pipeline with machine learning to improve accuracy. As one example, the error $d\eta$ in angular radius of each ring is not well-predicted by our current error-propagation approach, but a multilayer feed-forward neural network that considers the energies and locations of the interactions that define a Compton ring yields much better predictions. We have observed improvements of up to half a degree in localization accuracy from this approach. As

another example, recall that when a photon scatters only once in the detector before being absorbed, the Compton Law cannot distinguish between two distinct possibilities for the ring containing its source. However, a feed-forward neural network can select the correct ring with 80-90% accuracy.

We expect that machine-learning inference will increasingly be a component of astrophysics analysis pipelines that deploy on high-altitude balloons and spacecraft. Accommodating the computational needs of such inference on low-power, resource-constrained platforms will require optimizing ML models for deployment. Model quantization and exploitation of parallelism in model inference will likely be needed to achieve the throughput and latency required by ADAPT and APT. Optimized models may be deployed on low-power GPUs or FPGAs to meet the resource constraints of our experimental platforms.

Trading Off Accuracy vs. Efficiency. GRB localization entails a tradeoff between accuracy and resource usage. Our analysis pipeline may need to reduce its use of computational resources, either to conserve power or to accommodate other computations, such as flight control, using the same hardware. We can reduce the resource usage by using simpler algorithms or fewer iterations of least-squares, or by using a smaller sample of incident photons. These strategies yield a less accurate location for the GRB source, which in turn reduces the utility of any follow-up observations.

We say that the GRB analysis computation is *elastic*, in that its resource utilization can be compressed at some cost to result quality. Using empirical observations of result quality and computational cost for a variety of pipeline configurations, we have developed techniques [34] to rapidly pick a configuration with the highest localization accuracy that can be achieved for given resource usage. In the deployed ADAPT and APT instruments, configurations for this pipeline, along with other competing computations, can be controlled by an elastic real-time scheduler [28] that decides both which tasks to run and how heavily they must be compressed to meet latency constraints. Incorporating task-specific definitions of utility into the decisions of such a scheduler is an open research question that we plan to explore.

7 GRB IDENTIFICATION AND NOTIFICATION

While the computational pipeline described to this point is primarily tasked with localizing a GRB, we also need to decide whether or not a GRB is currently present in the data. While the initial approach to this problem will compare event rates (i.e., gamma-ray photon arrivals from a common direction) to a threshold, more sophisticated algorithms are under consideration.

Upon determination that a GRB has been detected, we need to notify secondary instruments that a burst has occurred, along with its direction in the sky and its spectral characteristics. One of the objectives of the ADAPT prototype mission is to test its ability to perform fast alerts to ground-based instruments. Historically, fast alerts were measured in hours [11], as the alert was not generated until data were transferred to the ground and analyzed there. Here, we have much lower latency expectations, on the order of seconds. Recall that low latency is important because low-intensity bursts fade within seconds to minutes, and we want to measure as much of the evolution of the burst's light curve as possible.

Upon detection, alerts will be distributed via NASA's General Coordinates Network (GCN)[1], formerly the Gamma-ray Coordinates Network. GCN is designed to share information about transient phenomena to whomever wishes to subscribe. Data distribution is over the public network, via Apache Kafka [21]. GCN has a history going back more than 3 decades, with phone alerts starting in the early 1990s, and has recently been used to distribute alerts for multi-messenger astronomy based on gravitational waves [23], in some cases with latency less than one minute.

For APT, communication with the ground incurs a minimum 5 s speed-of-light delay. Hence, for absolute minimum latency co-observation of a GRB, a collaborating follow-up instrument will need to be either aboard the APT satellite or nearby in space.

8 EARTHBOUND COMPUTATION

Although our efforts on ADAPT and APT have focused on localization of bright transients, these events, while scientifically important, will occupy only a small fraction of the instruments' total observing time. Most of the time, they will observe a much lower flux of gamma rays from diffuse sources, such as the galactic disk, and weak emitters such as the Crab Nebula. By gathering observations over a span of days to months, the instruments will produce a data set that allows detailed imaging of such sources. Imaging of this type has been done with prior gamma-ray observatories such as the Compton Spectrometer and Imager (COSI) [36].

Unlike transient localization, imaging weak sources is not time-critical, so it is feasible to send observations to the ground for offline analysis. The data rates involved are not particularly high and so do not require specialized data-reduction hardware, unlike, e.g., observations by ground-based telescope arrays [37]. However, the necessary imaging computations are both compute-intensive and sensitive to the exact methods and parameter settings used. For this reason, the COSI Consortium is hosting a series of "data challenges" to crowd-source image analysis and parameter tuning for these tasks [19]. To ensure that the data challenges are as accessible as possible, the COSI Consortium has released supporting software and candidate analysis methods to the astrophysics community as portable Python code.

In order to effectively engage the community, software must be both highly accessible (a goal achieved by using Python) *and* highly efficient. Imaging requires iterative methods [20] that we found to take hours to run even on relatively small Compton imaging data sets in a pure-Python implementation. Fortunately, JIT-based parallelization strategies for Python [22, 27] exist to recover performance without compromising portability. We recently demonstrated acceleration of 50-60× for imaging tasks in the first COSI Data Challenge on multicore CPUs [31]. We plan to work with the COSI team in the future to jointly accelerate future imaging computations on both multicores and accelerators such as GPUs.

9 CONCLUSIONS

We have described the computational requirements, algorithmic approaches, execution hardware, and future needs for the Advanced Particle-astrophysics Telescope (APT) and its Antarctic Demonstrator (ADAPT). ADAPT is scheduled for a balloon flight in the 2025-26

[1] https://gcn.nasa.gov/

season. Its computational pipeline is functional today, although we continue to pursue improvements.

There are a host of questions still to be addressed for APT, however. We have initial models of the anticipated background effects at the Sun-Earth Lagrange L_2 orbit [13] that warrant further investigation. We will investigate how the instrument scales from 4 calorimeter layers to 20. And finally, lessons learned from the balloon flight of ADAPT will be incorporated into the full instrument.

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, and the Peggy and Steve Fossett Foundation.

REFERENCES

[1] J Aasi et al. 2013. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics* 7, 8 (July 2013), 613–619. https://doi.org/10.1038/nphoton.2013.177

[2] BP Abbott et al. 2009. LIGO: The laser interferometer gravitational-wave observatory. *Reports on Progress in Physics* 72, 7 (2009), 076901. https://doi.org/10.1088/0034-4885/72/7/076901

[3] C Altomare, JH Buckley, et al. 2022. Simulation of a Compton-pair imaging calorimeter and tracking system for the next generation of MeV gamma-ray telescopes. In *Journal of Physics: Conference Series*, Vol. 2374. IOP Publishing, 012116. https://doi.org/10.1088/1742-6596/2374/1/012116

[4] C Aramo, E Bissaldi, M Bitossi, et al. 2023. A SiPM multichannel ASIC for high Resolution Cherenkov Telescopes (SMART) developed for the pSCT camera telescope. *Nucl. Instrum. Methods Phys. Res. A* 1047 (2023), 167839. https://doi.org/10.1016/j.nima.2022.167839

[5] 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. https://doi.org/10.1093/mnras/stz3634

[6] Imre Bartos and Marek Kowalski. 2017. *Multimessenger Astronomy*. IOP Publishing. https://doi.org/10.1088/978-0-7503-1369-8

[7] K Bechtol et al. 2012. TARGET: A multi-channel digitizer chip for very-high-energy gamma-ray telescopes. *Astroparticle Physics* 36, 1 (2012), 156–165. https://doi.org/10.1016/j.astropartphys.2012.05.016

[8] Diego Bersanetti et al. 2021. Advanced Virgo: Status of the detector, latest results and future prospects. *Universe* 7, 9 (2021), 322. https://doi.org/10.3390/universe7090322

[9] SE Boggs and P Jean. 2000. Event reconstruction in high resolution Compton telescopes. *Astronomy and Astrophysics Supplement Series* 145, 2 (2000), 311–321. https://doi.org/10.1051/aas:2000107

[10] James Buckley et al. 2021. The Advanced Particle-astrophysics Telescope (APT) Project Status. In *Proc. of 37th Int'l Cosmic Ray Conference*, Vol. 395. Sissa Medialab, 655:1–655:9. https://doi.org/10.22323/1.395.0655

[11] A Bulgarelli et al. 2014. The Agile alert system for gamma-ray transients. *The Astrophysical Journal* 781, 1 (Jan. 2014), 13 pages. https://doi.org/10.1088/0004-637X/781/1/19

[12] Wenlei Chen et al. 2021. 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. Sissa Medialab, 590:1–590:9. https://doi.org/10.22323/1.395.0590

[13] Wenlei Chen et al. 2023. The Advanced Particle-astrophysics Telescope: Reconstruction of the MeV gamma-ray sky and estimation of point-source sensitivity in the presence of the background. In *20th Divisional Meeting of the High Energy Astrophysics Division*. American Astronomical Society.

[14] Wenlei Chen, James Buckley, 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. of 38th Int'l Cosmic Ray Conference*, Vol. 444. Sissa Medialab, 841:1–841:9. https://doi.org/10.22323/1.444.0841

[15] S Funk et al. 2017. TARGET: A digitizing and trigger ASIC for the Cherenkov telescope array. In *AIP Conference Proceedings*, Vol. 1792. American Institute of Physics, 080012. https://doi.org/10.1063/1.4969033

[16] Ye Htet et al. 2023. Localization of Gamma-ray Bursts in a Balloon-Borne Telescope. In *Proc. of Workshops of Int'l Conf. on High Performance Computing, Network, Storage, and Analysis*. ACM, 395–398. https://doi.org/10.1145/3624062.3624107

[17] 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 Conference*, Vol. 444. Sissa Medialab, 956:1–956:9. https://doi.org/10.22323/1.444.0956

[18] Zachary Hughes et al. 2021. Characterization of a prototype imaging calorimeter for the Advanced Particle-astrophysics Telescope from an Antarctic balloon flight and CERN beam test. In *Proc. of 37th Int'l Cosmic Ray Conference*, Vol. 395. Sissa Medialab, 137:1–137:9. https://doi.org/10.22323/1.395.0137

[19] C Karwin, S Boggs, J Tomsick, A Zoglauer, et al. 2022. The COSI Data Challenges and Simulations. *Bulletin of the AAS* 54, 3 (April 2022). https://baas.aas.org/pub/2022n3i108p30

[20] J Knödlseder et al. 1999. Image reconstruction of COMPTEL 1.8 MeV ^{26}AL line data. *Astronomy and Astrophysics* 345 (May 1999), 813–825.

[21] Jay Kreps, Neha Narkhede, and Jun Rao. 2011. Kafka: A distributed messaging system for log processing. In *Proc. of 6th Workshop on Networking Meets Databases (NetDB)*. ACM, 7 pages.

[22] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. 2015. Numba: a LLVM-based Python JIT compiler. In *Proc. of 2nd Workshop on the LLVM Compiler Infrastructure in HPC*. ACM, Article 7, 6 pages. https://doi.org/10.1145/2833157.2833162

[23] LIGO Scientific Collaboration, Virgo Collaboration, et al. 2019. Low-latency Gravitational-wave Alerts for Multimessenger Astronomy during the Second Advanced LIGO and Virgo Observing Run. *The Astrophysical Journal* 875, 2 (2019), 20 pages. https://doi.org/10.3847/1538-4357/ab0e8f

[24] Péter Mészáros, Derek B Fox, Chad Hanna, and Kohta Murase. 2019. Multi-messenger astrophysics. *Nature Reviews Physics* 1, 10 (2019), 585–599. https://doi.org/10.1038/s42254-019-0101-z

[25] M Mishra et al. 2022. Application of High Density Digitizer System-on-Chip (HDSoC) prototype for acquiring fast silicon photomultiplier signals. In *Proc. of Nuclear Science Symposium and Medical Imaging Conference*. IEEE, 3 pages. https://doi.org/10.1109/NSS/MIC44845.2022.10399253

[26] Andrii Neronov. 2019. Introduction to multi-messenger astronomy. In *Journal of Physics: Conference Series*, Vol. 1263. IOP Publishing, 012001. https://doi.org/10.1088/1742-6596/1263/1/012001

[27] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis. 2017. CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations. In *Proc. of Workshop on Machine Learning Systems*. 7 pages. http://learningsys.org/nips17/assets/papers/paper_16.pdf

[28] James Orr, Chris Gill, Kunal Agrawal, Jing Li, and Sanjoy Baruah. 2019. Elastic Scheduling for Parallel Real-Time Systems. *Leibniz Trans. on Embedded Systems* 6, 1 (May 2019), 05:1–05:14. https://doi.org/10.4230/LITES-v006-i001-a005

[29] AN Otte et al. 2017. Characterization of three high efficiency and blue sensitive silicon photomultipliers. *Nucl. Instrum. Methods Phys. Res. A* 846 (2017), 106–125. https://doi.org/10.1016/j.nima.2016.09.053

[30] Stefan Ritt. 2008. Design and performance of the 6 GHz waveform digitizing chip DRS4. In *IEEE Nuclear Science Symposium Conference Record*. IEEE, 1512–1515. https://doi.org/10.1109/NSSMIC.2008.4774700

[31] Jamie Shin et al. 2024. Accelerating Compton Imaging of Astrophysical Sources in Python. In *21th Divisional Meeting of the High Energy Astrophysics Division*. American Astronomical Society.

[32] 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 Conference*, Vol. 395. Sissa Medialab, 588:1–588:9. https://doi.org/10.22323/1.395.0588

[33] 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 Conference*, Vol. 444. Sissa Medialab, 764:1–764:9. https://doi.org/10.22323/1.444.0764

[34] Marion Sudvarg, Jeremy Buhler, Roger D. Chamberlain, Chris Gill, James Buckley, and Wenlei Chen. 2023. Parameterized Workload Adaptation for Fork-Join Tasks with Dynamic Workloads and Deadlines. In *Proc. of IEEE 29th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 232–242. https://doi.org/10.1109/RTCSA58653.2023.00035

[35] Marion Sudvarg, Chenfeng Zhao, Ye Htet, Meagan Konst, Thomas Lang, Nick Song, Roger D. Chamberlain, Jeremy Buhler, and James H. Buckley. 2024. HLS Taking Flight: Toward Using High-Level Synthesis Techniques in a Space-Borne Instrument. In *Proc. of 21st International Conference on Computing Frontiers*. ACM, 12 pages. https://doi.org/10.1145/3649153.3649209

[36] John Tomsick et al. 2021. The Compton Spectrometer and Imager Project for MeV Astronomy. In *Proc. of 37th Int'l Cosmic Ray Conference*, Vol. 395. Sissa Medialab, 652:1–652:9. https://doi.org/10.22323/1.395.0652

[37] TC Weekes et al. 2002. VERITAS: the very energetic radiation imaging telescope array system. *Astroparticle Physics* 17, 2 (2002), 221–243. https://doi.org/10.1016/S0927-6505(01)00152-9

[38] Jacob Wheelock, William Kanu, Marion Sudvarg, Zhili Xiao, Jeremy D. Buhler, Roger D. Chamberlain, and James H. Buckley. 2021. Supporting Multi-messenger Astrophysics with Fast Gamma-ray Burst Localization. In *Proc. of HPC for Urgent Decision Making Workshop*. IEEE, 8 pages. https://doi.org/10.1109/UrgentHPC54802.2021.00008

[39] J Zorn et al. 2018. Characterisation and testing of CHEC-M—A camera prototype for the small-sized telescopes of the Cherenkov telescope array. *Nucl. Instrum. Methods Phys. Res. A* 904 (2018), 44–63. https://doi.org/10.1016/j.nima.2018.06.078