Parallel GRB Source Localization Pipelines for the Advanced Particle-astrophysics Telescope

MARION SUDVARG, JACOB WHEELOCK, JEREMY BUHLER, and JAMES BUCKLEY, Washington University in St. Louis, USA
WENLEI CHEN, University of Minnesota, USA

CCS Concepts: • **Computing methodologies → Parallel algorithms**; • **Applied computing → Astronomy**.

Additional Key Words and Phrases: multi-messenger astronomy, gamma-ray bursts, parallel algorithms, GPU acceleration

ACM Reference Format:
Marion Sudvarg, Jacob Wheelock, Jeremy Buhler, James Buckley, and Wenlei Chen. 2021. Parallel GRB Source Localization Pipelines for the Advanced Particle-astrophysics Telescope. 1, 1 (October 2021), 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

The Advanced Particle-astrophysics Telescope (APT) [2] is a planned space-based observatory to survey the entire sky for gamma-ray bursts (GRBs). It seeks to promptly detect these transient events, then communicate with narrow-band instruments for follow-up observations. To this end, we are developing analytical methods for real-time detection and localization of GRBs, then parallelizing and accelerating the software pipeline to maintain sufficient throughput for computing hardware that might fly onboard the orbiting platform.

As described in [6], we focus on detecting events for which a GRB's photons Compton-scatter one or more times within the instrument until they are eventually photoabsorbed. All scatterings for one photon appear simultaneous at the time resolution of the detector. Our localization pipeline, then, has three primary stages: (1) reconstructing each photon's trajectory in the instrument to estimate an annulus containing the photon's source direction; then combining the annuli from all detected photons to estimate the most likely direction by (2) finding a rough approximation of the direction from a set of initial candidates according to a maximum-likelihood approach; then (3) performing iterative least-squares refinement to produce a final estimate of source direction. Such analysis must be simultaneously accurate and fast, even on a low-power, embedded computational platform. In [6], we parallelize the pipeline to target an ARM Cortex-A53 processor, and measure execution times for each stage over a range of input data.

In this work, we present a GPU-accelerated version of the localization stages of our pipeline, i.e. (2) and (3). The initial reconstruction stage, i.e. (1), generates N annuli $(\mathbf{c}, \phi, \sigma)$, each constraining the estimated source direction to a circle with a center described by the vector $\mathbf{c}$ and opening angle ϕ, as well as an error term σ that "smears" the estimate.

For the initial approximation of source direction, we select 20 annuli at random, then run a CUDA kernel in which each one is assigned to a distinct block. For each annulus i, we test a set of 720 candidate source directions $\mathbf{s}_i$, evenly

---

Authors' addresses: Marion Sudvarg, msudvarg@wustl.edu; Jacob Wheelock, jacobwheelock@wustl.edu; Jeremy Buhler, jbuhler@wustl.edu; James Buckley, buckley@physics.wustl.edu, Department of Computer Science & Engineering, Washington University in St. Louis, 1 Brookings Drive, St. Louis, Missouri, USA, 63130; Wenlei Chen, chen6339@umn.edu, University of Minnesota, Minneapolis, Minnesota, USA, 55455.

Manuscript submitted to ACM

spaced on the circle $(\mathbf{c}_i, \phi_i)$. The kernel calculates the joint log-likelihood for each one with respect to all input annuli, distributing work across GPU threads in the block. The candidates with the highest likelihood for each annulus are averaged, weighted by likelihood, to produce an initial approximation of source direction, $\mathbf{s}_0$, which is passed to the refinement stage.

The iterative least-squares refinement stage uses a CUDA kernel to test if each annulus i lies within $3\sigma_i$ of an estimated source direction $\mathbf{s}$, initialized to the approximation $\mathbf{s}_0$. From those annuli that do, we generate a linear least-squares problem with a unit-norm constraint. As described in [4], this can be reduced to a quadratic eigenvalue problem in $O(N^2)$ time; this is accelerated in another CUDA kernel. The eigenvalue problem is then solved on the CPU using the Eigen library [5] to produce a refined estimate for $\mathbf{s}$. This process is repeated 20 times, with each iteration using the result of the previous as the estimated source direction. The final solution $\mathbf{s}$ is taken as the source direction of the GRB. More information on implementation and reproducibility can be found in [7].

To test our methods, we simulated a long GRB with a typical spectrum, using Geant4 [1] to generate 10^6 uniformly distributed gamma-ray photons from a normally incident, collimated beam with a cross-section of 18 m^2 to fully cover the APT detector. The APT instrumentation and analog-to-digital conversion were simulated using APTSoft (described in [3]). We measured the performance of both localization stages on an NVIDIA GeForce RTX 2080 GPU, then used this to estimate the performance on an NVIDIA Jetson NX Xavier,[1] for which the 10-watt power requirement makes it comparable to what might fly onboard the APT platform. The GPU-accelerated initial source approximation demonstrated an estimated speedup of 3.5 for a large input dataset. For the final stage, execution switches between the CPU and GPU at each refinement iteration, incurring significant execution time overhead. However, GPU acceleration slows the increase of execution time with input size when compared to the CPU times. This initial work suggests that GPU acceleration for source localization can improve execution time performance, and we continue to develop our kernels to reduce overhead and optimize further.

ACKNOWLEDGMENTS

This work was supported by NASA awards 80NSSC19K0625 and 20-APRA20-0148, NSF award CNS-1763503, Washington U.'s McDonnell Center for the Space Sciences, and the Peggy and Steve Fossett Foundation. We are grateful to Washington U. technical staff including Richard Bose, Dana Braun, and Garry Simburger for contributions to the prototype detectors and the Antarctic flight of the APTlite instrument.

REFERENCES

[1] S. Agostinelli, J. Allison, K. Amako, and others. 2003. Geant4 — a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506, 3 (2003), 250–303. https://doi.org/10.1016/S0168-9002(03)01368-8

[2] James Buckley. 2021. The Advanced Particle-astrophysics Telescope (APT) Project Status. In *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, Vol. 395. 655. https://doi.org/10.22323/1.395.0655

[3] Wenlei Chen, James Buckley, S. Alnussirat, and others. 2021. The Advanced Particle-astrophysics Telescope: Simulation of the Instrument Performance for Gamma-Ray Detection. In *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, Vol. 395. 590. https://doi.org/10.22323/1.395.0590

[4] W. Gander, G. H. Golub, and U. Von Matt. 1989. A constrained eigenvalue problem. *Linear Algebra Appl.* 114-115 (1989), 815–839. https://doi.org/10.1016/0024-3795(89)90494-1

[5] Gaël Guennebaud, Benoît Jacob, and others. 2010. Eigen v3. http://eigen.tuxfamily.org.

[6] Marion Sudvarg, Jeremy Buhler, James H. Buckley, and others. 2021. A Fast GRB Source Localization Pipeline for the Advanced Particle-astrophysics Telescope. In *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, Vol. 395. 588. https://doi.org/10.22323/1.395.0588

[7] Jacob Wheelock, William Kanu, Marion Sudvarg, and others. 2021. Supporting Multi-messenger Astrophysics with Fast Gamma-ray Burst Localization. UrgentHPC 2021, to appear.

---

[1] The RTX 2080 has 46 SMs, while the Jetson Xavier has only 6. We conservatively estimate each SM of the RTX 2080 to perform twice as quickly as an SM of the Xavier NX, giving us a scaling factor of about 15 for the execution time estimates.