PROCEEDINGS OF SCIENCE

Interferometric Neutrino Event Reconstruction in Inhomogeneous Media with the Askaryan Radio Array

M. Beheler-Amass, M. Beydler, A. Karle, J. L. Kelley*, and M-Y. Lu

Wisconsin IceCube Particle Astrophysics Center (WIPAC)
University of Wisconsin–Madison, Wisconsin, U.S.A.
E-mail: jkelley@icecube.wisc.edu

for the ARA Collaboration

Full author list: http://ara.wipac.wisc.edu/collaboration/authors/ara_icrc17

The Askaryan Radio Array (ARA) is an ultra-high-energy (UHE) neutrino detector under construction at the South Pole. Using antennas deployed up to 200m deep in the ice, ARA searches for the impulsive Askaryan radio signals from UHE neutrino interactions. Directional reconstruction of these events is complicated by the inhomogeneity of the polar ice cap, causing the radio waves to bend as they travel through less dense regions close to the surface. To deal with the computational challenges involved, we have developed a spline-based scheme for fast lookup of pre-calculated raytracing results in the ice and air, combined with a parallelized interferometric directional reconstruction. Using these techniques, we demonstrate an angular resolution of $O(1^\circ)$ for neutrino event reconstruction as well as a fast online event filter based on the cross-correlation of antenna signals.

35th International Cosmic Ray Conference (ICRC2017)
10–20 July 2017
Bexco, Busan, Korea

*Speaker.

PoS(ICRC2017)1054

1. Introduction

The Askaryan Radio Array (ARA) is a large-scale neutrino detector under construction near the geographic South Pole in Antarctica [3]. Its primary scientific goal is to discover the predicted flux of cosmogenic neutrinos [1, 2], by deploying clusters of radio antennas ("stations") in the ice and detecting the broadband Askaryan radio emission of UHE neutrino interactions ($E_\nu > 10^{16}$ eV) above a background of thermal noise. An ARA station consists of a cluster of 16 antennas of vertical and horizontal polarization that are deployed up to 200 m below the surface of the snow at the South Pole (Fig. 1). The station's central DAQ electronics on the surface trigger when three of eight antennas of either polarization exceed a power threshold; the antenna signals are then digitized and transmitted to a central server. Three ARA stations are currently deployed, with an inter-station spacing of 2 km, and three more stations will be deployed during the 2017–18 austral summer season.

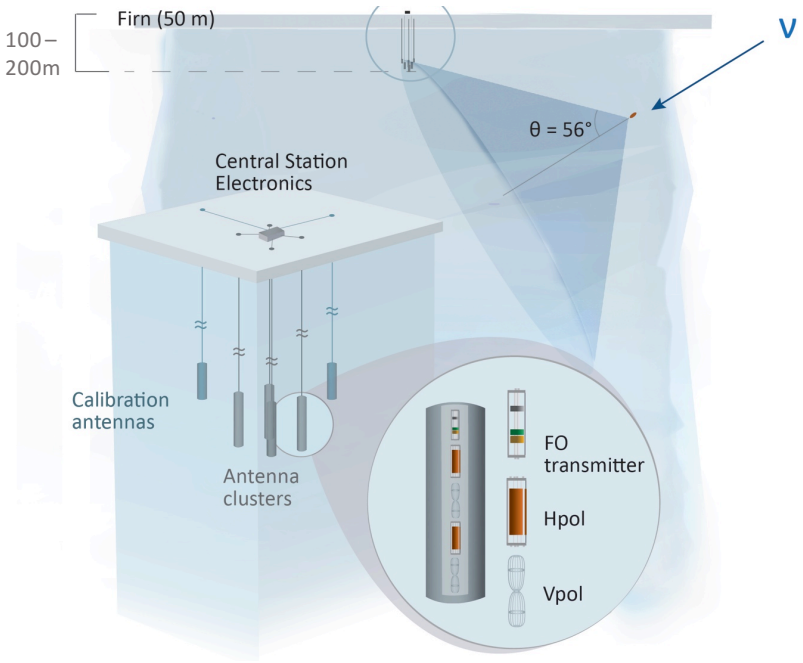

Figure 1: Illustration of an ARA detector station. An incoming neutrino that interacts with the ice will generate a broadband Askaryan radio pulse. The station consists of 16 detector antennas and calibration pulsers deployed below the snow surface and a central station DAQ that triggers upon and digitizes the impulsive signals.

Using the relative timing information of the Askaryan signal as recorded in spatially-separated antennas in the ice, one can reconstruct the direction of the incoming radio signal and, in some cases, the distance to the interaction vertex. This reconstruction is an important initial step in reconstructing key properties of a neutrino such as source direction and energy, but the procedure is complicated by a gradient in the index of refraction of the polar ice cap. The radio signals thus travel in curved paths from the source (neutrino interaction vertex) to the target (the receiving antenna), requiring a complex raytracing calculation to determine the exact trajectory. Such calculations are

feasible if the interaction vertex is known, for example, in neutrino simulations, but quickly become intractable for event reconstruction, where many vertex hypotheses must be tested.

We present here details of the radio propagation in the ice and air, along with a software tool for fast lookup of interpolated, pre-calculated tables that greatly speeds up the calculation of raytrace paths in the ice and air. This fast lookup of raytrace time delays enables an interferometric cross-correlation directional reconstruction, which in turn has also been accelerated using parallel computation techniques. The signal cross-correlation can also be used as an event quality parameter for online noise rejection.

2. Radio Propagation in Ice and Air

The polar ice cap at the South Pole is approximately 2850 m deep and, as one approaches the surface, transitions from ice (index of refraction n_{deep}) to a compacted snow layer known as the firn (n_{surface}). Specifically, we model the index of refraction at 300 MHz as a continuous function of the form $n(z) = n_{\text{deep}} + (n_{\text{surface}} - n_{\text{deep}})\, e^{Cz}$ for $z < 0$, where $n_{\text{deep}} = 1.78$, $n_{\text{surface}} = 1.35$, $C = 0.0132$ m^{-1}, and $z = 0$ at the firn-air boundary (based on measurements in Ref. [4]). As a ray propagates in the ice, it will bend away from the surface (Fig. 2), creating a "firn shadow" within which no solution exists from source to target. The fractional volume of the firn shadow decreases for deeper antennas.

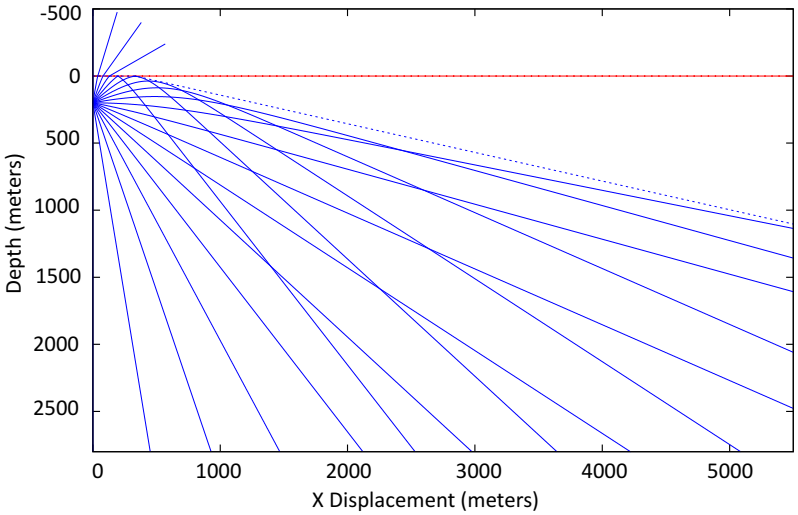

Figure 2: Sample collection of allowed radio paths through the South Pole ice and air, to a target (receiving antenna) at a depth of 200 m. The wedge-shaped forbidden region in the shallow ice is the firn shadow.

The full raytrace calculation involves first numerically solving an equation for the initial "launch" angle, and then using this initial condition to solve the set of propagation differential equations with a Runge-Kutta method [5]. This calculation is relatively time-consuming and becomes intractable when the number of paths is very large, as in event reconstructions that test many directions as hypotheses. To speed the calculation, we can precompute raytrace solutions on a grid, and then interpolate on this grid to allow arbitrary source and target locations.

Because of approximate azimuthal symmetry of the ice, we can reduce the dimensionality of the precomputed result table to three variables: source depth z_s, target depth z_t, and horizontal separation distance r. For event reconstruction, the most immediately useful raytrace result is the direct propagation delay time t without reflections. To smoothly interpolate the surface $t = f(z_s, z_t, r)$, we fit the grid with multidimensional B-splines [6]. However, the existence of the firn shadow and the firn-air boundary introduces several complications, as attempting the spline fit over discontinuities results in ringing that reduces the accuracy of the fit. To work around these issues, we split the solution into three tables: an initial table describing the firn shadow boundary, and two result tables for the raytrace results separated for sources in air and sources in ice (Fig. 3).

First, a spline table describing the boundary of the firn shadow is consulted to determine if any physical solution exists. If so, a table for in-ice or in-air sources is evaluated to determine the propagation delay. The source points for the spline table fits have been extrapolated into the firn shadow in order to avoid artifacts from the discontinuity. The accuracy of the spline approximation at randomly selected points in the parameter space is 0.1 ns in air and 0.2 ns in ice, with a few outliers up to 2–3 ns (Fig. 4). On a typical CPU (2.3 GHz Intel Core i7), the time to compute a single ray improves from 0.21 ms to 0.4 μs. The spline tables and supporting software have been packaged as a C++ project, `radiospline`[1].

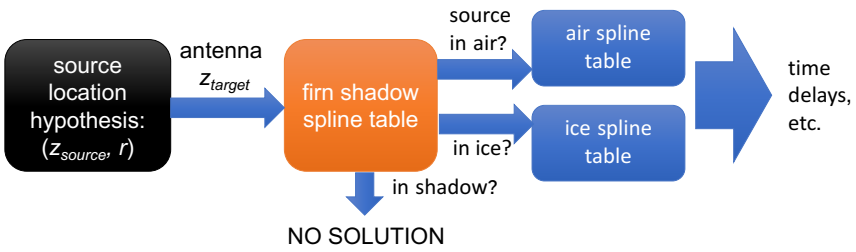

Figure 3: Multi-step lookup of spline-interpolated raytrace solutions.

3. Interferometric Cross-correlation

Since the radio signals recorded in station antennas of matching polarization are similar except for their time offset, interferometric analysis is a powerful technique to determine event quality and to reconstruct the interaction vertex / direction [7]. Several techniques are possible; we focus here on the cross-correlation of pairs of antenna signals.

First, a power-normalized cross-correlation C_{ij} of each antenna pair (i, j) of a given polarization is calculated. This convolution is efficiently done in the frequency domain, with

$$C_{ij}(t) = \frac{v_i(t) * v_j(t)}{\sqrt{P_i P_j}} = \frac{\mathscr{F}^{-1}(\mathscr{F}(v_i) \cdot \mathscr{F}(v_j))}{\sqrt{P_i P_j}} , \tag{3.1}$$

where P_i is the summed power in waveform v_i, and $\mathscr{F}$ is the FFT operation, including a modified Hann window.

[1] http://github.com/WIPACrepo/radiospline

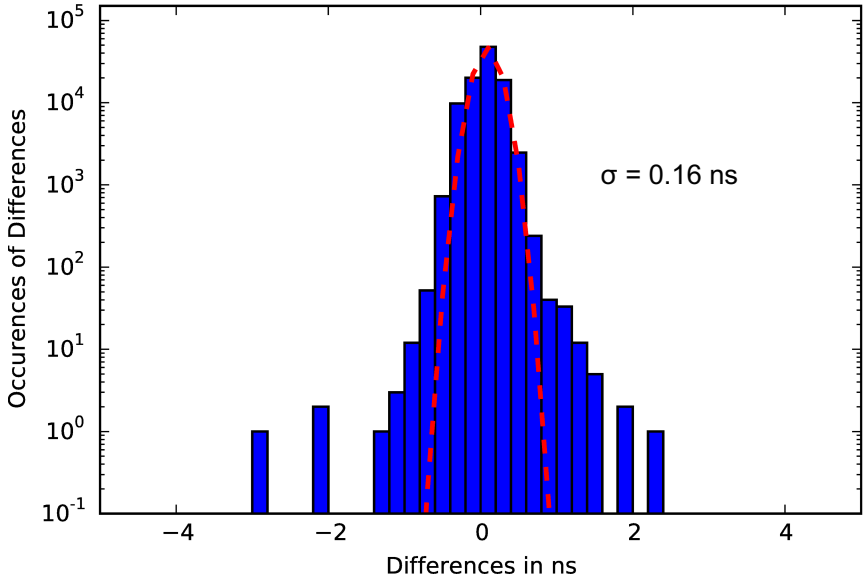

Figure 4: Difference between the full raytrace calculation and spline approximation of radio propagation delay, for sources and targets at random locations in the ice.

Next, given a hypothesized vertex location $\boldsymbol{x}$, we calculate the propagation delay T_i to the ith antenna using `radiospline`. We then calculate the total cross-correlation $C(\boldsymbol{x})$ for that location by summing the cross-correlation at the time differences $T_i - T_j$ for each antenna:

$$C(\boldsymbol{x}) = \frac{1}{N_{\text{pairs}}} \sum_{i<j} C_{ij}(T_i - T_j) \; . \tag{3.2}$$

In an ideal case, C is maximal when the vertex matches the actual neutrino interaction point. In a random direction, the geometric time differences to different antennas will not match the time differences observed in the waveforms, resulting in a low cross-correlation.

The calculation of the cross-correlation is highly parallel as the computation of C for each possible direction / vertex location does not depend on any other, and the computation of any C_{ij} for an antenna pair doesn't depend on any other pair. We have exploited this parallelism to accelerate the computation using the OpenCL[2] heterogeneous computing platform. In particular, this allow us to use the massively parallel floating-point pipelines in modern Graphics Processing Units (GPUs). We will discuss benchmarks of the accelerated algorithm in the following section.

3.1 Neutrino Event Directional Reconstruction

The cross-correlation C shown in Eq. 3.2 can be used for event directional reconstruction, assuming a fixed horizontal distance to the vertex. $C(\theta, \phi)$ is computed for a Healpix[3] grid on the sky, and then the direction with the largest value represents the reconstructed direction.

[2] https://www.khronos.org/opencl/
[3] http://healpix.jpl.nasa.gov

In practice, a combination of noise, antenna response, and limited time resolution can fracture the true maximum into a large number of smaller local maxima nearby. This effect can be smoothed out by replacing the cross-correlation pairs C_{ij} in Eq. 3.1 with its Hilbert envelope (see e.g. [9] for a general description of the envelope technique). An example of a reconstructed simulated neutrino event is shown in Fig. 5. The accuracy of the directional reconstruction for simulated events is $\sim 0.3^\circ$ in azimuth and $\sim 1.3^\circ$ in zenith [8].

Using the OpenCL-accelerated reconstruction, calculating the skymap for a grid with $\sim 2 \times 10^5$ directions requires 130 ms/event on a CPU (Intel Xeon E5-2680 2.70GHz) and 50 ms/event on a single GPU (NVIDIA GeForce GTX 690). The computation can be further accelerated by adding more GPUs, for example, with each covering a fraction of the sky.

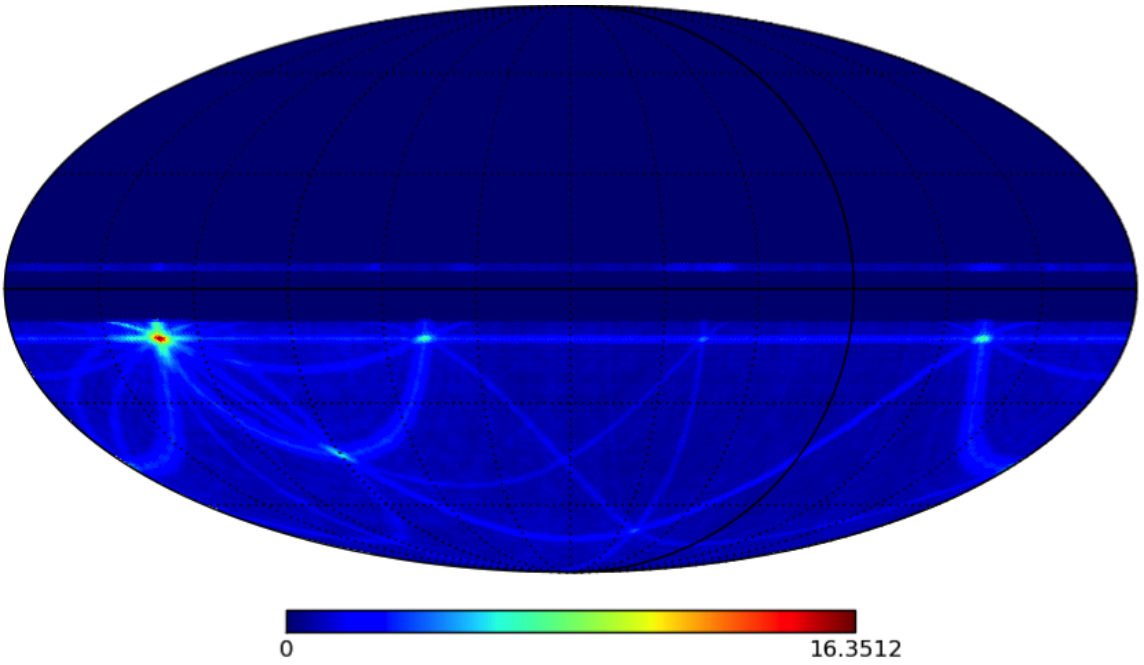

Figure 5: Skymap showing directional reconstruction of a simulated neutrino event of energy 10^{18} eV with one ARA station. The zenith and azimuth indicate the angles from the station center to the event vertex, while the color scale is the cross-correlation in arbitrary units.

3.2 Online Event Filter

The cross-correlation technique can also be used as a filter to separate causal signals from background thermal noise triggers; the quality parameter is simply the maximum value over all directions $C(\theta,\phi)_{\max}$. With the acceleration described above, and given a station trigger rate of $O(10)$ Hz, this filter is potentially fast enough to be used online at the South Pole, where separation of "interesting" events in near-real time is important given the limited satellite bandwidth for data transfer.

The complementary cumulative distributions of $C_{\max}$ for simulated neutrino events of various energies compared with a representative sample of ARA data are shown in Fig. 6. The ARA data at trigger level consists largely of thermal noise, so the cross-correlation maximum is low, whereas

neutrino events clearly separate from the noise population. Higher-energy neutrino events tend to produce larger signals in the antennas and so typically have a higher cross-correlation due to a larger signal-to-noise ratio. The falloff in the distributions for data and simulation are at slightly different values since simulated noise is not identical to the noise of this particular ARA station. $C_{\max}$ can either be used as a quality criterion with a strict cut, or as an ordering parameter to select the highest quality events. The latter ordering technique is currently used in the ARA online system.

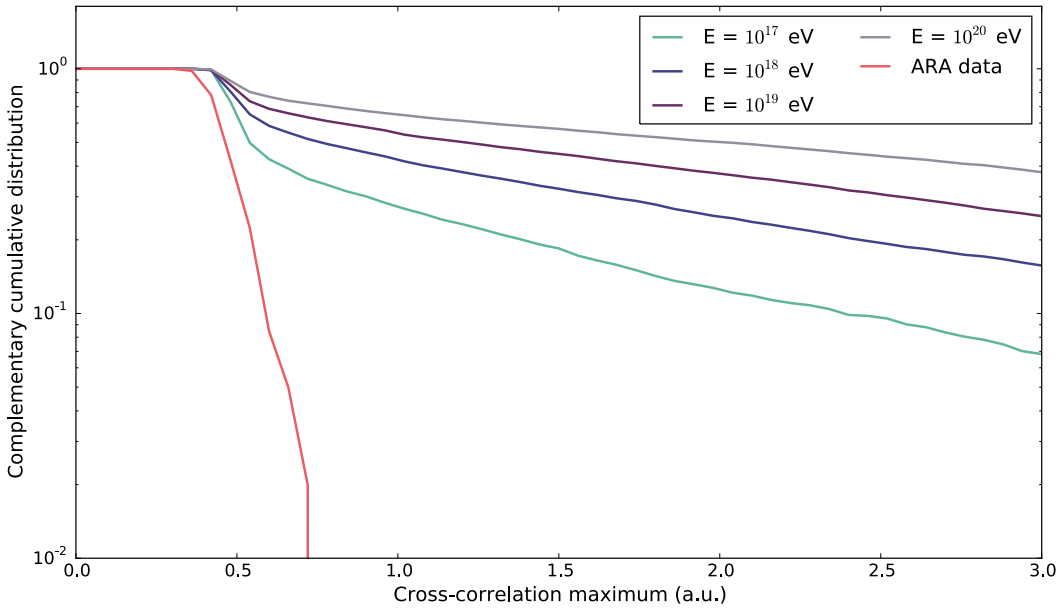

Figure 6: Complementary cumulative distributions of maximum cross-correlation values for simulated neutrino events of various energies and a sample of ARA trigger-level data. The data consist largely of thermal noise and therefore have low cross-correlations.

4. Outlook

The fast approximation of raytrace results with `radiospline` and the OpenCL acceleration of the interferometric cross-correlation have enabled advanced reconstruction of ARA events on a larger scale than previously possible. Spline fits of other raytracing results beyond the propagation delay of the direct ray, such as launch and receipt angles or the propagation delay of the secondary reflected solution, may enable new reconstruction techniques. For example, the time differences between the direct ray and the secondary reflected/refracted ray may improve vertex distance reconstruction beyond what is possible with just the direct rays [10].

The cross-correlation filter of Sec. 3.2 can also be expanded by including antennas of both polarizations. As ARA grows, in order to run online in near-real time and keep up with the aggregate detector trigger rate, the algorithm must be scalable. This is likely possible by simply adding more GPUs and/or reducing the resolution of the skymap, but this has yet to be verified.

5. Acknowledgments

The authors acknowledge the support of the National Science Foundation via grant BIGDATA #1250720.

References

[1] K. Greisen, Phys. Rev. Lett. **16**, 748 (1966).

[2] G. T. Zatsepin and V. A. Kuzmin, JETP Lett. **4**, 78 (1966).

[3] P. Allison *et al.* (ARA Collaboration), Phys. Rev. D **93**, 082003 (2016).

[4] I. Kravchenko, D. Besson, and J. Meyers, J. Glaciol. **50**, 522 (2004).

[5] P. Allison *et al.* (ARA Collaboration), Astropart. Phys. **70**, 62 (2015).

[6] N. Whitehorn, J. van Santen, and S. Lafebre, Comp. Phys. Comm. **184**, 2214 (2013).

[7] A. Romero-Wolf *et al.*, Astropart. Phys. **60**, 72 (2015).

[8] M.-Y. Lu *et al.* (ARA Collaboration), Proc. 35$^{\text{th}}$ ICRC (2017), PoS(ICRC2017)966.

[9] P. Schellart *et al.*, A&A **560**, A98 (2013).

[10] D. Seckel *et al.* (ARA Collaboration), Proc. 35$^{\text{th}}$ ICRC (2017).