Observation of Reactor Antineutrinos with a Rapidly Deployable Surface-Level Detector

Alireza Haghighat,[†] Patrick Huber,[*] Shengchao Li,[‡] Jonathan M. Link,[§] Camillo Mariani,[¶]
Jaewon Park,[**] and Tulasi Subedi[††]

Center for Neutrino Physics, Department of Physics, Virginia Tech, Blacksburg, Virginia, USA

We deploy a small 80 kg antineutrino detector, called MiniCHANDLER, based on a solid plastic scintillator, for nearly three months at a distance of 25 m from a 2.9 GW thermal-power-reactor core at the North Anna Nuclear Generating Station. We report the detection of an antineutrino signal resulting from inverse beta decay at 5.5σ significance with no overburden and minimal shielding. This result also demonstrates that three-dimensional segmentation can be used to significantly improve the signal-to-noise ratio, in this case by a factor of 4. In addition, this measurement represents an observation of a positron spectrum in a small surface-deployed detector; this observation of reactor antineutrinos is achieved with a mobile neutrino detector mounted in an ordinary small trailer.

I. INTRODUCTION

Nuclear reactors have long been known to be a copious source of electron antineutrinos, which are emitted as a by-product of nuclear fission. It is not surprising, therefore, that neutrinos were proposed as a method for monitoring nuclear-reactor operations more than 40 years ago [1]. Neutrino reactor monitoring is nonintrusive, since it can be performed from outside the reactor building. The reactor neutrino signal depends on both the reactor power and the composition of the reactor core. In particular, a core that is rich in plutonium will produce a neutrino spectrum of lower average energy than a reactor that is rich in uranium [2]. These two signatures can be effectively disentangled by simultaneously measuring the neutrino rate and the energy spectrum. Case studies [3,4] have revealed an important advantage of neutrino monitoring compared with the usual nonproliferation safeguards, which rely on a continuous history of reactor operations and refuelings: should this continuity of knowledge be lost for a reactor, it is extremely difficult to restore. Neutrino reactor monitoring does not rely on a detailed knowledge of the reactor's operational history, and thus the continuity-of-knowledge issue is avoided altogether.

There are a number of detailed case studies in the literature highlighting specific potential applications of small above-ground detectors. These applications include power monitoring of reactors [5,6], and monitoring of spent nuclear fuel [7], plutonium disposition, and mixed-oxide fuel usage [8,9]. Recently, a detailed study has been presented [10,11] of how these capabilities could be applied in a future agreement to denuclearize the Korean peninsula.

Reines and Cowan used a reactor as the source for their 1956 neutrino-discovery experiment [12], and, since then, many generations of reactor neutrino experiments have followed, with a reliance on overburden to shield cosmic rays being an element common to all. There have also been a number of successful safeguards-oriented reactor experiments [13–17], starting in the mid-1980s. Here again, these experiments all relied on significant overburden. For real-world applications, such as in nuclear nonproliferation safeguards, it is exactly this dependence on overburden that has prevented the adoption of neutrino technologies. Practical applications require detectors that can operate with minimal shielding. In this paper, we describe such a detector technology and report on the observation of reactor neutrinos in a small-scale prototype; for other similar measurements, see Refs. [18,19]. Here we demonstrate a detector technology that can operate with minimal shielding, has a small detector volume, and requires no liquid scintillator; in particular, the neutrino spectrum is measured over a broad range of energies, including low energies, with high efficiency and good precision. This combination of features has been previously identified as crucial for real-world application of neutrino reactor monitoring [20].

---

[*]pahuber@vt.edu

[†]haghigha@vt.edu

[‡]scli@vt.edu

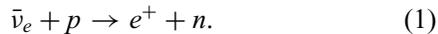[§]jmlink@vt.edu

[¶]camillo@vt.edu

[**]jaewon.park@vt.edu

[††]tpsubedi@vt.edu

In a typical reactor neutrino detector, electron antineutrinos are observed via the process of inverse beta decay (IBD), in which the neutrino interacts with a hydrogen nucleus in an organic scintillator, producing a positron and a neutron:

$$\bar{\nu}_e + p \rightarrow e^+ + n. \tag{1}$$

The positron deposits its kinetic energy in the scintillator and annihilates, resulting in a prompt (or primary) flash of light, while the neutron thermalizes and is captured by a nucleus, producing a delayed (or secondary) signal. The signature of the IBD interaction is the coincidence, in space and time, of positronlike and neutronlike events. This compares favorably with the two largest backgrounds, which are (1) fast neutrons from the cosmic ray flux that recoil off of a proton in the scintillator and are captured, and (2) random coincidences between unrelated positron- and neutronlike events. The random-coincidence events have no correlation in space or time, while the fast-neutron events generally share the temporal correlation of the IBD events but have a larger mean spatial separation due to the greater initial neutron energy and hence speed.

The CHANDLER (Carbon Hydrogen Anti-Neutrino Detector with a Lithium Enhanced Raghavan optical lattice) detector technology is designed for the detection and precision spectral measurement of reactor electron antineutrinos in a high-background surface-level environment. It also allows portable detectors, which are easy to assemble and easy to maintain, while eliminating the complications and hazards associated with a liquid scintillator. The CHANDLER design is based on an optical lattice, invented by Raju Raghavan as a part of the LENS research and development (R&D) program [21]. The Raghavan optical lattice (ROL) transports light by total internal reflection along rows and columns of cubes, as shown in yellow in

the left-hand panel of Fig. 1. The red cube represents the location of the original energy deposition in this example. This gives the spatial resolution of the detector at the level of a single cube, while at the same time maximizing the light-collection efficiency. In the CHANDLER detector, the ROL is formed from layers of plastic scintillating cubes with a size of 6.2 cm in a tightly packed rectangular array. These, in turn, are stacked in alternating layers with thin neutron-detection sheets in between, as shown in Fig. 1. In each layer, light is transported along the rows and columns of the cube lattice by total internal reflection from the inner cube surfaces. This allows determination of the *x*-*y* position of the cube where the original energy deposition occurs (red cube), and the *z* position is obtained as well, since light is largely confined to the layer in which it is generated. The neutron-detection sheets are semitranslucent, i.e., enough of the light produced in these sheets will propagate into the adjacent cubes. However, the light produced in the cubes has only a small chance of traversing a sheet, resulting in light leakage between the layers. The plastics used in the detector naturally maintain a thin cube-to-cube or cube-to-sheet air gap, which is required for total internal reflection. The plastic scintillator cubes are doped with a wavelength-shifting compound so that the light from the neutron sheets can be absorbed in the cubes, reemitted, and then transmitted by total internal reflection. The key to this pairing of a plastic scintillator with neutron-detection sheets is that the scintillator used in the neutron sheets releases its light much more slowly than does the plastic, and this results in a clean neutron signature. Pairing neutron sheets with scintillator cubes was first implemented by the SoLid Collaboration [22], where optically isolated cubes are read out by wavelength-shifting fibers running along the edges of rows and columns of cubes. Replacing the fiber readout with a ROL allows one to increase the photocathode coverage and thus increase light collection.
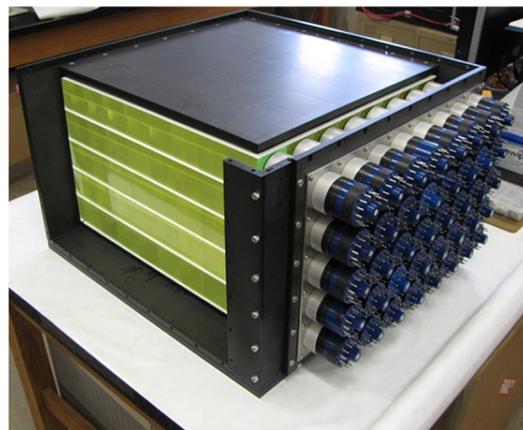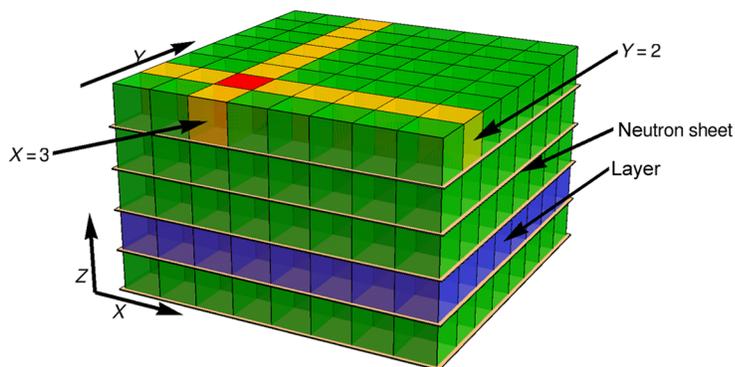

FIG. 1.   Left: schematic illustration of the MiniCHANDLER detector, with the top neutron sheet and photomultiplier tubes (PMTs) not shown. Right: the MiniCHANDLER detector during assembly, with one side open, showing the alternating layers of wavelength-shifting plastic scintillator cubes and neutron sheets.

As a result, the energy resolution increases, while the high spatial resolution and clean neutron tag of the SoLid design are maintained. When combined, these properties have significant advantages in terms of the rejection of backgrounds that could otherwise overwhelm the neutrino signal in a surface-level detector. For details of other ongoing simulation studies and R&D for two-dimensional (2D) segmented plastic detectors, see for instance Refs. [23–26].

The neutron-detection sheets and plastic scintillator used in the CHANDLER detector are sold commercially by Eljen Technology as EJ-426 and EJ-260, respectively. EJ-426 is composed of microparticles of lithium-6 fluoride (^{6}LiF) mixed with microcrystals of silver-activated zinc sulfide (ZnS:Ag) scintillator. Thermal neutrons are captured by a ^{6}Li nucleus, resulting in an alpha particle and a triton, which, due to their high specific energy loss, deposit their energy very locally in the ZnS:Ag scintillator. ZnS:Ag has a scintillation decay constant of about 200 ns, which is about 20 times longer than the decay time of the EJ-260 scintillator used in the cubes. This large difference in the decay times of the scintillation light is used to identify neutron captures and separate them from signals originating in the plastic scintillator. The cube segmentation in the CHANDLER detector makes it possible to apply an unbiased prompt-delayed spatial-separation cut that is well matched to the typical positron-neutron separation of an IBD event. Compared with the standard Daya Bay analysis [27], which uses no spatial-separation cut, the coincidence volume in the CHANDLER detector is reduced by a factor of more than 2000. In addition, this segmentation can be used to veto fast-neutron events with associated proton recoils in more than one cube, and to tag the 511 keV gamma rays from positron annihilation in an IBD event. Together, these topological selections enable us to identify IBD events in a surface-level detector where correlated background events outnumber the true IBD events by more than 400 to 1.

The MiniCHANDLER detector is a 80 kg prototype of the full CHANDLER detector. MiniCHANDLER is designed to maximize the detector mass within our limited project budget, with a detector that replicates light transportation from the middle of the envisioned full-scale detector. MiniCHANDLER consists of five layers of an 8×8 cube array read out by PMTs on only one end of each cube row and column, so that two faces of the detectors are instrumented with PMTs. MiniCHANDLER has six neutron-sheet layers, above and below each cube layer. The sheets are optically connected to the cube layers on both sides; see Fig. 1.

The PMTs used in MiniCHANDLER are Amperex XP2202 tubes, with a custom-built resistive-divider base. The PMTs are operated with a negative high voltage, supplied by a CAEN mainframe, with each channel individually tunable. The PMT signals are read out by a CAEN V1740 waveform digitizer with a 62.5 MHz sample rate, a 12-bit analogue-to-digital converter (ADC), and 64 channels per card. To ensure high fidelity with this relatively sparse sample rate, the PMT signals are first passed through a preamplifier to shape the signal with a 25 ns time constant. The V1470 digitizer is internally triggered on every instance of a channel at or above 14 ADC counts (ADCC) over the baseline. Each trigger leads to a 129-sample readout of all channels in the module, starting about 35 samples before the trigger. Two independently triggered V1470 modules are used to read out the full detector.

Data from the waveform digitizers are sent to a data-acquisition (DAQ) computer over an optical link, where they are processed with a zero-suppression algorithm to suppress data from channels in which the waveform deviates by only 12 ADC counts or fewer from the baseline. Only after this zero suppression are the data written to disk, as separate files for each module. The two data streams are merged off line using events from an external strobe trigger (with a rate of slightly less than 1 Hz) to continuously synchronize the merging based on time stamps from the modules' internal clocks. The strobe-trigger merging is used to estimate the DAQ efficiency, which we find to be greater than 99.5%.

II. REACTOR AND DEPLOYMENT

The MiniCHANDLER detector, electronics, and DAQ computer are loaded into a 14 ft trailer, dubbed the Mobile Neutrino Lab, which is equipped with a carefully designed quiet power supply, Wi-Fi connectivity, and air conditioning, allowing fully remote operation. On June 15, 2017, after several weeks of commissioning and testing at Virginia Tech, the trailer is moved to the North Anna Nuclear Generating Station in Mineral, Virginia. The North Anna plant consists of two pressurized water reactors, each with a licensed thermal power of 2940 MW [28]. The Mobile Neutrino Lab is deployed next to Reactor 2, at a distance of about 25 m from the center of the core. At this location, it is approximately 90 m from the core of Reactor 1, which is therefore responsible for about 7% of the neutrino interactions in the detector [29]. The detector and DAQ computer are up and running in less than one day, which marks the start of the site-specific commissioning. To combat the increased thermal-neutron rate from the reactor, the detector is surrounded by a layer of 1-inch-thick boron-loaded polyethylene with holes for the PMTs. The natural gamma rate at the reactor site is higher than that experienced at Virginia Tech. To combat this, we add 1 inch of lead shielding below the detector, and on the two sides closest to the containment building.

With commissioning complete, the data run begins on August 9, 2017 and lasts until November 2, 2017. During this time we take 1133.6 h of usable reactor-on data

TABLE I. Description of the periods of operational runs. Each run corresponds to 1 h.

Period	Runs	Reason for new period
1	258	10 ADC threshold
2	164	Streamlined disk I/O
3	255	Change to 14 ADC threshold
4	5	Reactor ramp-down
5	569	Reactor off
6	118	High-voltage retune
7	49	Reactor ramp-up
8	476	Reactor at full power

and 675.4 h of reactor-off data. The data are divided into eight periods, where the transitions between periods correspond to changes in the operational state of either the detector system or the reactor. Table I describes the different periods, and lists the reasons for the start of each new operational period. Of particular note is the transition from period 2 to period 3, which corresponds to a shift in the trigger threshold from 10 to 14 ADCC. This becomes necessary when the rate of low-energy gamma rays increases due to the arrival of shipping containers of mildly activated equipment that are parked next to the Mobile Neutrino Lab in preparation for the refueling of Reactor 2.

III. CALIBRATION

For the study described here, a highly accurate energy model and reconstruction are neither a requirement nor an objective. Nevertheless, matching the known energy dependence of reactor neutrinos in an observed reactor-on excess is an essential confirmation of neutrino detection. In addition, we are motivated to test an alternative energy-calibration source made possible by the high segmentation of the ROL. Specifically, in polyvinyltoluene, a minimum-ionizing particle has a dE/dx of about 2 MeV/cm [30], which means that a muon passing vertically through a 6 cm cube deposits an average energy of around 12 MeV. In the following section, we describe how we use vertical muons to measure the light pattern from every cube location in the detector, and how this allows us to fix the energy scale at around 12 MeV. Here, we assume a proportional energy response for energies below 12 MeV.

The PMT high voltage is initially tuned to align the muon peaks across all channels to 1500 ADC counts. To account for gain fluctuations, the muon peak is measured in each channel for each run, and the measured ADC values are scaled to realign the muon peaks. In this context, the muon peaks are not limited to vertical muons, which have limited statistics in a single run, but include all triggers across all cube positions.

IV. EVENT RECONSTRUCTION

Neutron identification in MiniCHANDLER is based on pulse-shape discrimination, using the factor-of-20 difference in the scintillation-light decay times between the neutron sheets and the scintillator cubes. A naive particle-identification (PID) variable can be formed from the ratio of the area under the waveform to the peak value. Large values of this variable correspond to neutronlike events, while small values correspond to signals generated in the plastic scintillator. Large signals, with peak values greater than 1000 ADCC, are eliminated from consideration. If a signal satisfies our neutron PID criterion in at least one PMT channel, the whole event becomes a neutron candidate.

Instrumental effects in MiniCHANDLER, such as PMT flashes and analog overshoot from an earlier large pulse, can generate signals that satisfy this naive neutron PID selection; fortunately, these effects almost never replicate the decaying light pattern of an energy deposition in the neutron sheets. We use a template-based χ^2 criterion to reject these instrumental backgrounds from the neutron candidate list. To obtain the neutron template, we start with a sample of 100 hand-selected neutron-capture waveforms. Each waveform is divided into eight regions. In each region, the ADC counts over baseline are summed, and these sums are divided by the total over all regions to form normalized amplitudes. Then these normalized amplitudes are averaged over the 100 hand-selected waveforms to form the neutron template. Since the events in the plastic scintillator have short pulses, which are contained entirely in the first region, the gamma template is trivial. With these templates, the neutron selection proceeds as follows.

Within each view of each layer, we select the channel with the highest-amplitude signal, compute its normalized amplitudes and uncertainties in the eight regions, and compute the χ^2 values relative to both the neutron (χ_n^2) and the gamma (χ_γ^2) template. The reduced χ^2 values from both the x and the y view are summed, and we select good neutrons satisfying the criteria $\sum_{x,y} \chi_{n_i}^2/\nu_i < 8$ and $\sum_{x,y} \chi_{\gamma_i}^2/\nu_i > 150$, where ν_i is the number of time bins in the template, effectively the number of degrees of freedom. This xy matching fixes the positions of the neutron candidates.

Once neutron identification in a layer is done, we check the consistency of xy-matched neutrons from different layers. Neutron candidate events generally exhibit low occupancy in the detector. Therefore, the xy position of an event is given simply by the location of those PMTs which see the most light. For about half of all neutron-capture events, we see light on only one side of the neutron sheet. We call these events "cube" neutrons, since we cannot distinguish whether the capture happens in the sheet above or below the cube. For these events, the z position of the

neutron is assigned to the middle of the cube. For the remainder of the events, the neutron capture is seen on both sides, and the neutron-capture position is known at the sheet level; we call these events "sheet" neutrons. Any event with more than one neutron candidate among the five layers is rejected. Tests with Li-free neutron sheets in our MicroCHANDLER prototype show that in the absence of ^{6}Li there are practically no neutronlike signals in the detector. Therefore, for the purpose of this analysis, we can treat all neutronlike events as neutrons without introducing any bias.

Event reconstruction for prompt events is somewhat more complicated than for neutrons because the number of active cubes in the detector is often greater than one. This is due to the Compton scattering of positron-annihilation gamma rays in IBD events, and to the possibility of multiple proton recoils in fast-neutron background events. In order to use this topological information, we need a reconstruction that is capable of evaluating energy depositions in multiple active cubes. Here, a challenge arises when there is more than one active cube in a single detector layer. This is a nontrivial problem, because in each detector layer we have 2×8 observed PMT signals, but there are 8×8 unique cube locations. If we knew the true energy deposition in each cube in a layer, expressed as a 64-component vector $\mathbf{e}$, then we could write an expression for the expected PMT responses as a 16-component vector $\mathbf{p}$. This forward problem is represented by

$$\mathbf{p} = \mathbf{M} \cdot \mathbf{e}, \quad (2)$$

where $\mathbf{M}$ is the 16×64 transfer matrix. Each element of the transfer matrix, M_{ij}, describes the size of the signal in PMT j arising from a 1 MeV energy deposition in cube i. This transfer matrix includes all effects arising from light propagation, including attenuation and scattering, and electronic crosstalk. Although about 80% of the light detected is observed in the PMTs at the ends of the row and column centered on the emitting cube, the remaining 20% of light is spread out across the other PMTs in the plane. This unchanneled light is due to tiny imperfections in the ROL and to scattering in the bulk of the plastic cubes. In addition to the unchanneled light, there is a bipolar inductive-pickup crosstalk which is observed in channels neighboring a channel with a large-amplitude pulse.

Our objective is to invert this matrix equation to solve for $\mathbf{e}$, the vector of cube energies, but first we have to determine the transfer matrix $\mathbf{M}$. Even then, there is no exact solution to Eq. (2), since $\mathbf{M}$ has no inverse.

A data-driven approach is used to determine the elements of the transfer matrix. This is the best way ensure that all effects are properly accounted for. We use vertical muons, which are easily identified in our detector by requiring that the observed light be consistent with coming from the same single-cube position in each plane. By definition, a vertical muon produces light in only one cube per plane, and that cube's position is well known from the xy coordinates of the vertical muon. Unchanneled light and electronic crosstalk spread this signal over all channels in the plane. By collecting a large sample of vertical muons, which occur at a rate of 0.7 Hz across the detector, we measure the response of each PMT in each plane to energy depositions from every cube position in that plane. According to our simulation, the most probable energy deposition for a muon that satisfies the vertical selection is 11.42 MeV/cube. The transfer matrix elements are scaled to an equivalent energy of 1 MeV. In constructing the final transfer matrix, which is applied to all layers, we average the elements from the matrices measured in just the middle three detector layers; we do this because we can only be certain that a "vertical" muon's path through a layer is fully contained in a single cube when there are confirming hits above and below that layer. In the case of the top and bottom layers, one of these confirming hits is missing.

A sample of the vertical-muon spectra from cubes at three different distances from the PMTs is shown in Fig. 2. The width of these energy distributions comes from the natural Landau distribution of dE/dx, the geometrical acceptance for muons that are not perfectly vertical, and the intrinsic resolution of the detector. The shift in the peak position, as a function of distance from the PMT, shows the spatial dependence of the detector response function. This effect is explicitly accounted for in the transfer matrix. Throughout the region of interest, this spatial dependence is independent of the deposited energy. This is illustrated
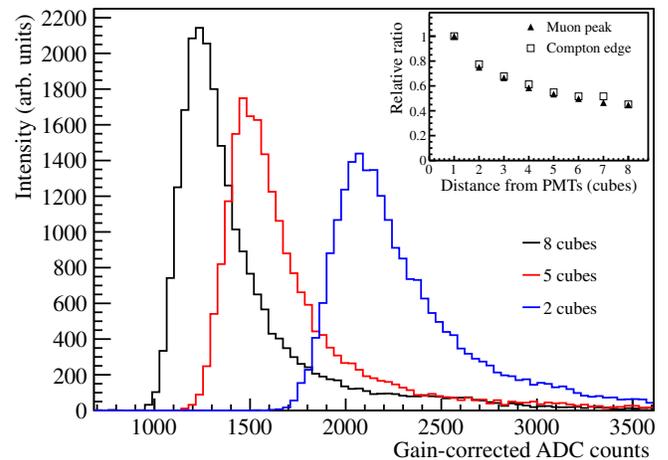

FIG. 2. Vertical-muon energy spectra at three distances from the PMTs. The inset shows the effective light attenuation in the ROL as determined from vertical muons at 11.42 MeV, compared with the attenuation measured with the Compton edge of 1.275 MeV gamma rays from ^{22}Na at around 1 MeV.

in the inset of Fig. 2, which compares the effective attenuation curve for vertical muons with one derived from the approximately 1 MeV Compton edge of 1273 keV gamma rays from a ^{22}Na source.

Both the unchanneled light and the electronic crosstalk scale with the amount of light detected in the primary channel, but the variances of the unchanneled-light and electronic-crosstalk components do not. For unchanneled light, the variance scales with the Poisson statistics of the photons at the PMT cathode, while for crosstalk the variance scales with the number of electrons at the PMT anode. The future full detector will use electronics without crosstalk.

GEANT4 [31] is used to compute the true cube-by-cube energy deposition for a set of simulated gamma and IBD events. Using the transfer matrix **M** and random fluctuations drawn from a Poisson distribution with appropriately scaled variances, this truth information is propagated to create a Monte Carlo realization of the measured PMT signals. This sample is used to test and tune the event reconstruction. As previously stated, Eq. (2) has no exact solution; any approximate solution must compensate for the lack of observables by some regularization scheme. This essentially amounts to using a Bayesian prior to select from among the possible solutions. For our analysis, the goal is to correctly reconstruct the number of cubes with a nonzero energy deposition, with a preference for suitable solutions with the fewest active cubes. This matches our expectation for IBD events, which the Monte Carlo calculation shows will almost never have more than five cubes with true energy depositions above the detection threshold in the MiniCHANDLER detector.

Using the variance found from the data, we construct a suitable likelihood function L to measure how well a given set of reconstructed cube energy depositions $\mathbf{e}_r$ corresponds to the measured PMT signals $\mathbf{p}_s$. To minimize $\hat{L} \equiv -\log L$, we use the following algorithm: First, we set all $e_r(x, y) = 0$ (i.e., all positions have an initial energy deposition of zero), such that the set of cubes with nonzero energy deposition, λ, is initially empty. Then we perform the following steps.

(1) Find the additional cube (x, y) that yields the smallest $\hat{L}$ when $\hat{L}$ is minimized by varying e_r for the cubes in the set λ plus the new cube (x, y).

(2) If $\min[\hat{L}(\lambda)] - \min\{\hat{L}[\lambda + (x, y)]\} < L_c$, go to step 5.

(3) Add cube (x, y) to the set λ.

(4) If λ has fewer than five cubes, return to step 1.

(5) If $\hat{L} < L_g$, accept the event as reconstructed, otherwise declare the reconstruction as failed.

This algorithm allows the reconstruction to assign energy depositions to additional cubes as long as the improvement

in $\hat{L}$ is sufficiently large (greater than L_c). This cutoff prevents overfitting, since adding a cube always will decrease $\hat{L}$. The reconstruction is limited to no more than five cubes in a layer, which is a conservative upper limit relative to the observation in our Monte Carlo calculation that IBD events have no more than five active cubes in the whole detector. Step 5 ensures that the fit is a good match to the data by requiring the final $\hat{L}$ to satisfy a quality criterion (less than L_g). This is rarely violated by IBD events in the Monte Carlo calculation, but in the data, where the true composition of the event types is unknown, we find that about 7% of events fail this criterion in at least one layer. We thus quote a reconstruction efficiency of 93%, but, since the IBD events generally do not share the typical characteristics of the events that fail the reconstruction, we presume it to be much higher for true IBD events. The fit cutoff L_c was tuned on Monte Carlo IBD events, and the quality parameter L_g was tuned on background data samples. This maximizes the fidelity of the reconstruction to the true cube positions and energy depositions, and minimizes reconstruction failures.

As a test, the reconstruction is applied to a sample of vertical muons from across the whole detector. The resulting energy spectrum is fitted to a convolution of Landau and Gaussian distributions. The fitted peak value is in good agreement with the most probable energy deposition obtained from the simulation. We interpret the fitted Gaussian σ as the average energy resolution at 12 MeV, which is found to be 2.6%. If the resolution scaling is purely stochastic, this corresponds to an average resolution of approximately $10\%/\sqrt{E(\text{MeV})}$.

V. IBD ANALYSIS

To compute the expected IBD spectrum and number of events, we use the Huber-Mueller reactor flux model [2,32], and the IBD cross section from Ref. [33] with a neutron lifetime of 878.5 s. The thermal reactor power is taken to be 2.94 GW, and the core-detector distance is 25 m. The detector mass is 80 kg, comprising 4×10^{27} target protons. From the simulation, we compute that 46% of all IBD neutrons in the detector are captured on ^{6}Li. Of these, 34% are lost when we discard the first 40 μs of Δt. The entire reactor-on data set comprises 1133.6 h of good data. Under these assumptions, we expect about 3500 Li-tagged IBD events. Given the uncertainties in the Monte Carlo calculation, the reactor distance, and spill-in and spill-out effects, it is difficult to assign a firm error, but 10%–20% appears reasonable. IBD neutrons created outside the detector and IBD neutrons reflected back into the detector are difficult to simulate precisely, and Monte Carlo tests indicate that this effect is below 10%. The reactor-detector distance is known to within 1 m, translating into about 8% uncertainty. The reactor power is known within 1%, and fuel burnup is not corrected for, but this effect does not exceed 5%.

In quadrature, these errors would add up to 12%–14%, but, for instance, the distance uncertainty is non-Gaussian, i.e., the distance could be between 24 and 26 m but it certainly is not 26.5 m; the same holds for the neutron-capture efficiency and the burnup effect. So, these errors can *not* be added up in quadrature. The exact value of the systematic uncertainty in the number of expected IBD events has no impact on the statistical significance of the IBD signal, since the IDB signal is derived entirely from a comparison of reactor-on and reactor-off data without recourse to the expected number of IBD events. The reason for calculating the expected number of IBD events is to check whether the number of expected events is consistent with the observed events, which it is found to be.

GEANT4 is used to simulate the cube-level energy deposition from IBD events, but we do not use it to propagate photons through the ROL. Instead, we generate the PMT signals in units of ADCC using the forward transfer matrix derived from vertical muons, followed by a Poisson smearing based on the observed and scaled variances. The simulated PMT signals are run through the reconstruction and event selection just like the data. Therefore, any nonlinearity in the reconstructed energy spectrum should be common to both the data and the Monte Carlo calculation, at least to within the precision of this analysis.

To form IBD event candidates, we begin by matching each neutron-capture candidate with all non-neutron events that have a successful reconstruction from the preceding 1000 μs. Next, we apply a prompt-delayed spatial-separation cut. The prompt-event position is assigned to the center of the most energetic cube of the primary event. To assign the position of the delayed event, we distinguish between sheet and cube neutrons, as explained previously. As we expand the allowed separation, more correlated events are included in the sample. At short distances, we find the largest enrichment of true IBD events, but as the separation grows, fast-neutron events start to dominate. To select the optimal separation cut, we study the significance of the IBD signal as a function of the separation cut. Figure 3 shows $\Delta\chi^2$ relative to the null hypothesis, plotted as a function of the maximum allowed prompt-delayed separation. The stepped nature of this plot is due to the quantization of separation distances inherent in our assignment of event positions in the cube structure. The significance peaks at a separation of 1.5 cube lengths, or 9.3 cm. At this distance, the cut includes the 19 nearest cube positions and the 20 nearest sheet positions. From our IBD Monte Carlo calculation, we estimate that 67.3% of the true IBD neutrons are captured within this region. As the fast-neutron rejection improves in future incarnations of the detector, this cut can be opened up to improve the IBD efficiency while maintaining the maximum significance.

The three-dimensional (3D) segmentation of MiniCHANDLER allows us to select events further based on the topology of the event. Under perfect conditions, one would
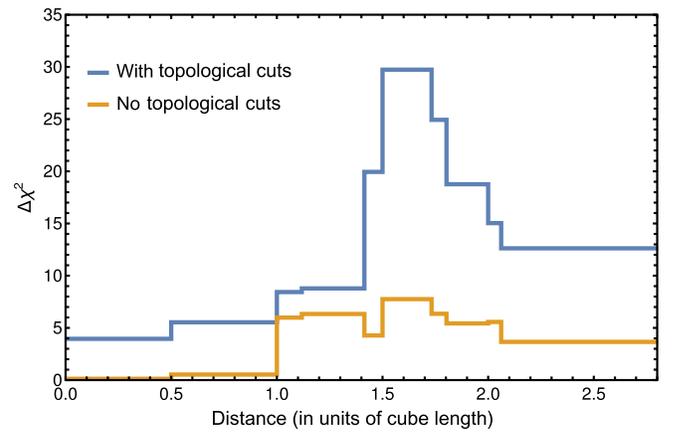

FIG. 3. Significance of IBD signal, expressed as $\Delta\chi^2$ relative to the null hypothesis, plotted as a function of the maximum allowed distance between the cube that contains the prompt signal and the one that contains the delayed signal, in units of cube length. The significance is determined by varying the cuts in the data and running through the full analysis.

design cuts to specifically tag the two 511 keV positron-annihilation gamma rays. In the current MiniCHANDLER detector, this is not practical for two reasons: First, the detector is too small to efficiently contain the first Compton scatter from the two annihilation gamma rays. Second, with the current light-collection scheme, the detector's energy threshold is about 50 keV, and at that level, many of the annihilation-gamma Compton scatters are unseen in the detector. Therefore, we implement a set of cuts to retain events with any hint of positron-annihilation gamma rays, while rejecting events that are clearly inconsistent with their presence. Specifically, we require there to be least one cube, beyond the primary (or highest-energy) cube, with an energy deposition in the range of 50 keV $\leq e_r \leq$ 511 keV. Further, we require that the sum of the energies in all cubes, excluding the primary cube and its most energetic immediate neighbor, be no more than 1022 keV, and that outside of those two cubes there is no single cube energy above 511 keV. These last cuts are designed to remove fast neutrons with multiple proton recoils. As can be seen by comparing the blue and orange lines in Fig. 3, these topological cuts improve the significance of the signal from $\Delta\chi^2 = 7.7$ to $\Delta\chi^2 = 29.7$, or, equivalently, the signal-to-noise ratio is improved by a factor of about 4. This demonstrates that the fine-grained 3D segmentation at the core of the CHANDLER technology adds considerable value relative to the coarser-grained 2D segmentation used in other contemporary detectors [34,35]. With anticipated improvements in the light collection and a larger detector to better contain the annihilation gamma rays, the efficacy of these topological cuts should be significantly enhanced.

The surviving events are split into reactor-on and reactor-off samples (see Table I). In each sample, they are

sorted by their reconstructed prompt energy into 20 bins from 0.5 to 20 MeV, with the lowest-energy bin being 0.5 MeV wide and all other bins being 1 MeV wide. In each energy bin, the distribution of the prompt-delayed Δt is fitted with an exponential plus a flat function. The exponential time constant τ is fixed at 94 μs, determined from a single Δt fit for the data from all energy bins and reactor periods (see Fig. 4). The Δt fits are used to statistically separate the time-correlated events (the exponential component) from the random-coincidence events (the flat component). Using all positron candidate events in the 1000 μs preceding a neutron—as opposed to just using the first event, or vetoing all events when two or more positron candidates are observed—ensures that the Δt distribution for the random-coincidence contribution is flat over all times. Then, by fitting this distribution out to more than ten neutron-capture lifetimes, we get a high-fidelity high-statistics measure of the random component, which we then subtract to get the correlated rates. A sample Δt distribution, with a fit, is shown in Fig. 4. Due to effects related to the analog side of our signal-processing chain, we exclude the first 40 μs from the fit. In the subsequent analysis, this results in a loss of 34% of all true IBD events.

In the final step of the analysis, we perform a background subtraction by taking the difference between the correlated events in the reactor-on periods and those in the reactor-off periods. In this step, there is a danger of introducing structure into the energy spectrum if the detector operation is not stable over time. Figure 5 shows the rates of correlated (red) and random-coincidence (blue) events as a function of time, as extracted from the Δt-distribution fits but without topological cuts applied, to enhance the fast-neutron events relative to the IBD events.
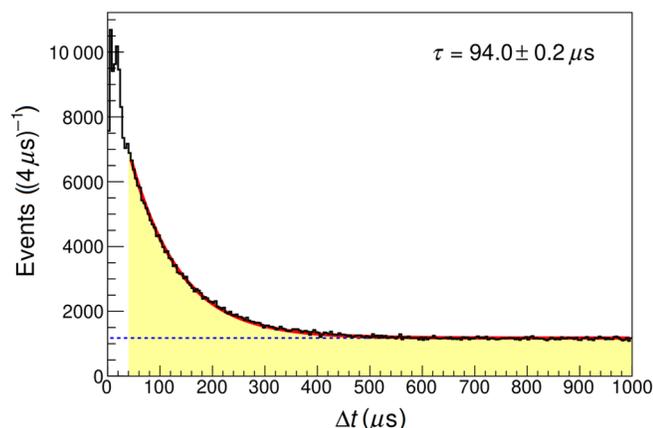

FIG. 4.   Histogram of Δt for all events in the energy range of 0.5–20 MeV. This distribution is fitted with an exponential plus constant to extract the true correlated events. The fit begins above 40 μs (shown in yellow) to bypass the low-Δt region, where instrumental effects distort the distribution.

The random-coincidence rate shows large variations between periods, which are linked to specific operational events at the plant. For example, during the shutdown, when the thermal-neutron rate from the reactor is essentially zero, the random-coincidence rate is cut by a half. Similarly, at the start of period 3 we see a slightly higher random-coincidence rate, which corresponds to the arrival of several shipping containers as discussed earlier. The increased trigger threshold from 10 to 14 ADCC that followed this event is applied in software after the fact to the data from periods 1 and 2 to ensure uniformity across the periods.

On the other hand, the period-to-period jumps observed in the random-coincidence rate are not seen in the correlated-event rate. Instead, we see smaller undulations that are anticorrelated with the atmospheric pressure. This is exactly what one would expect if the correlated rate is dominated by fast neutrons in the cosmic ray flux, as should be the case here. It is well known that the cosmic neutron rate is related to the atmospheric pressure, which is a measure of the mass of the atmosphere above. The air pressure shown in the middle panel of Fig. 5 is measured at Louisa County Airport, located 16.7 km from the North Anna Nuclear Generating Station, and is obtained from the NOAA website [36]. Using these data, we compute a correction factor for the measured pressure P relative to the average pressure P_0, which is equal to $e^{-\alpha(P-P_0)}$ with $\alpha = 7.3$ atm^{-1} [37]. In the bottom panel of Fig. 5, this correction factor is applied to the measured correlated-event rates, which, once corrected, are stable across all data-taking periods. The orange band represents the average statistical error of the correlated-event rate as measured in 8 h blocks.

While the impact of air pressure on the fast-neutron rate is a well-understood phenomenon that can be compensated for in the overall rate, it is not immediately clear whether differences in the average air pressure between the reactor-on and reactor-off periods could introduce an energy dependence into the correlated rate that could mimic an IBD signal. To test this hypothesis, the reactor-on data are split evenly into high-pressure and low-pressure sets, and the analysis is run on both halves. The IBD excesses measured in the two subsamples agree to within 1σ.

An analysis based purely on the total correlated-event rates is conducted as a cross-check of the spectral analysis. In this case, the air-pressure correction is applied run by run. This incorrectly rescales the IBD-signal events; however, since the correction is at most 5% in any given run, this is of little consequence for the current purpose. Also, the DAQ live-time efficiency has a slight systematic difference between reactor-on and reactor-off runs due to the extra thermal-neutron triggers that occur when the reactor is on. We measure this efficiency in each run by comparing the number of recorded strobe triggers with
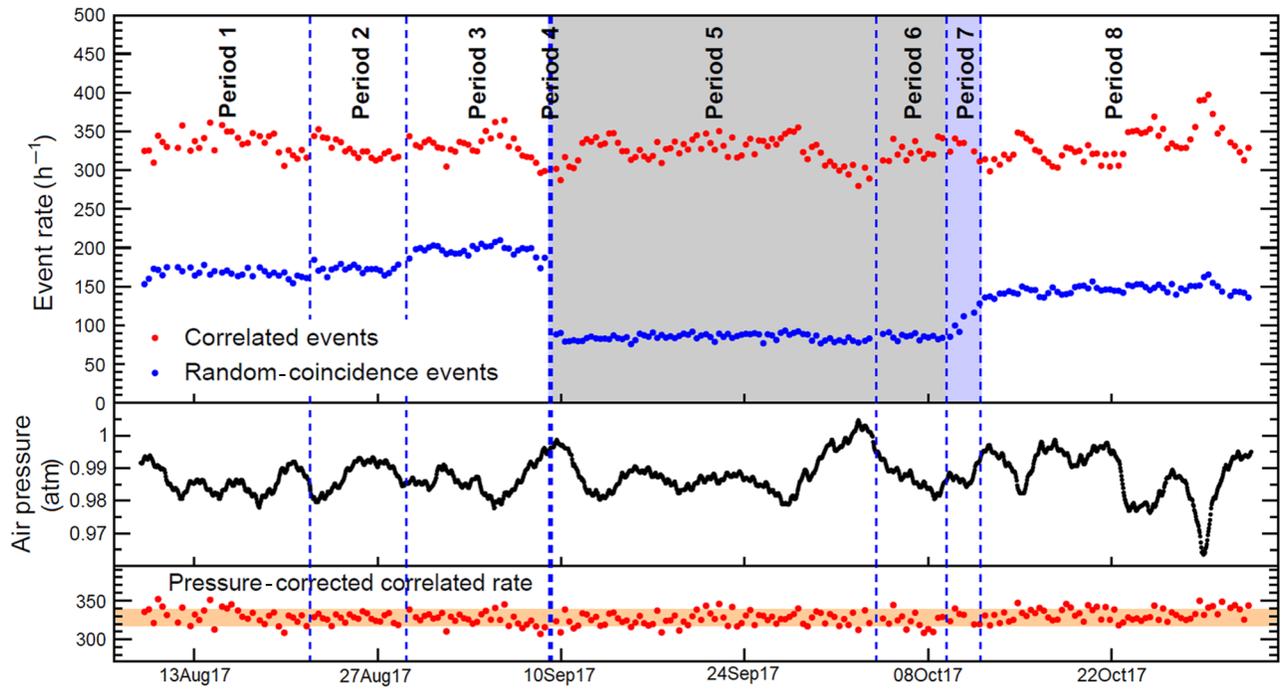
FIG. 5.   The top panel shows the rates, as a function of time, of correlated (red) and random-coincidence (blue) events as extracted from a fit to the Δt distribution for each 8 h period. These events are selected without topological cuts applied, in order to enhance the fast-neutron events relative to the IBD events. The middle panel shows the time dependence of the atmospheric pressure, which is anticorrelated with the variations of the correlated-event rate. The bottom panel shows the correlated-event rate, corrected for atmospheric pressure as described in the text. The orange band is the average 1σ uncertainty in the correlated-event rate. The run periods are described in Table I. The gray-shaded periods (periods 5 and 6) correspond to reactor-off periods. The blue-shaded periods (periods 4 and 7) correspond to reactor-power ramping, and are not used in the IBD analysis.

the number that are sent. The statistical significance of this counting analysis is a strong function of the signal-to-noise ratio, and so we apply a 3–8 MeV energy cut, which should retain 58% of IBD events while reducing the background by a factor of 3.6. To further enhance the signal-to-noise ratio, the topological cuts are applied. The data are divided into three time periods: before, during, and after reactor shutdown. These correspond to the data-taking periods in Table I as periods 1–3, periods 5–6, and period 8, respectively. The results are shown in Fig. 6.

This analysis finds an on:off excess of 1.22 ± 0.35 events/h, which corresponds to a 3.5σ significance. This compares well with the expected IBD rate in the 3–8 MeV range of 1.27 events/h, which demonstrates that the topological cuts are highly efficient for true IBD events. If we multiply the observed excess by the total reactor-on time and correct for the 58% efficiency of the 3–8 MeV cut, this corresponds to 2418 ± 700 events, which is entirely consistent with the result of the spectral analysis. The lower significance of the rate-only analysis is partly due to the inefficiency of the energy cut and partly to the absence of information from the IBD spectrum. If we use the full 0.5–20 MeV range, we find an on:off excess of 2.48 ± 0.94 events/h compared with an expectation of 2.17

IBD events/h. This corresponds to only 2.6σ significance, demonstrating a well-known feature of rate-only analyses: tight cuts must be applied to obtain a suitable signal-to-noise ratio, resulting in a lower overall IBD efficiency.

It is worth noting that the pressure and DAQ-efficiency corrections are comparable in size to the expected
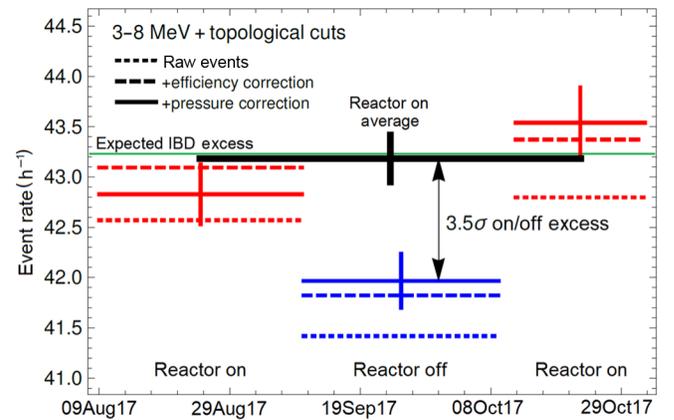

FIG. 6.   Results of the total correlated-rate analysis, which are consistent with the conclusions of the spectral analysis.

IBD excess in this rate analysis. Therefore, when the signal-to-background ratio is low, this type of analysis may be less reliable than desired. Also, the random-coincidence rates can be strongly correlated with the reactor power (see Fig. 5), and therefore must be separated in a robust way, such as by the Δt-fit method illustrated in Fig. 4.

In a spectrum-based analysis, we can exploit the fact that no IBD events are expected above 8 MeV, and thus we can use this part of the spectrum as a side band to normalize the reactor-on and reactor-off periods relative to each other. We calculate a scale factor by taking the ratio of reactor-on and reactor-off correlated events in the 8–20 MeV region. Our computed scale factor is 1.666 ± 0.013, which turns out to be very close to the factor we get from a dead reckoning of the relative reactor-on:reactor-off live time, 1.673 ± 0.005. That we get this good agreement in spite of the approximately 4% rms variation in the atmospheric-pressure correction is due to the fact that the difference between the average pressure corrections in the reactor-on and reactor-off periods is, by chance, quite low (less than 0.6%). This scale factor is applied to the correlated-event numbers in all energy bins of the reactor-off spectrum, and then we perform the reactor-off subtraction. The resulting spectrum is shown in Fig. 7. The error bars are obtained by propagating the error in the correlated-event rate obtained from the Δt fit in each bin from both the reactor-on and the reactor-off periods. The bin-to-bin correlated error arising from the scale factor is not shown in the plot, but it is included in the computation of the signal significance.

We perform a one-parameter fit of the observed reactor-on-off difference to the predicted signal spectrum. In this fit, we fully account for the statistical uncertainty of the normalization between the reactor-on and reactor-off data sets. The result of the fit is the best-fit value $\hat{a}$ of the amplitude a. This is compared with the null hypothesis, where $a = 0$, and the signal significance is

$$\sqrt{\chi^2(\hat{a}) - \chi^2(a = 0)}.$$

Our best-fit value corresponds to 2880 ± 528 IBD events, for a ratio of observed to expected events of $82\% \pm 15\%$. Given that the distance cut has a simulated efficiency of 67% and that the topological cuts are expected to be very efficient for true IBD events, this is within expectations. Overall, this constitutes a 5.5σ detection of reactor neutrinos, in a detector with no overburden. Our signal has the temporal, spatial, and energy signature expected for true IBD events.

Not surprisingly, we find the efficiency of the spectral analysis to be more than twice that of the most sensitive rate analysis. In the rate analysis, the low-energy bins, where the signal-to-noise ratio is poor, can still contribute to the signal significance because their large uncertainties are contained bin by bin, such that they do not dilute the significance of the higher-energy bins where the signal-to-noise ratio is much better. Furthermore, the spectral analysis does not rely on sizable corrections for the DAQ efficiency and atmospheric pressure, which makes it inherently more robust.

VI. OUTLOOK

An 80 kg prototype, as presented in this paper, is sufficient to demonstrate reactor-on-off detection for a multi-gigawatt reactor over a period of a few weeks. However, in a safeguards context, there are numerous ways in which this information can be obtained much more easily, without recourse to neutrinos.

The unique capability offered by neutrino reactor monitoring is an *in-situ* quasi-real-time determination of the
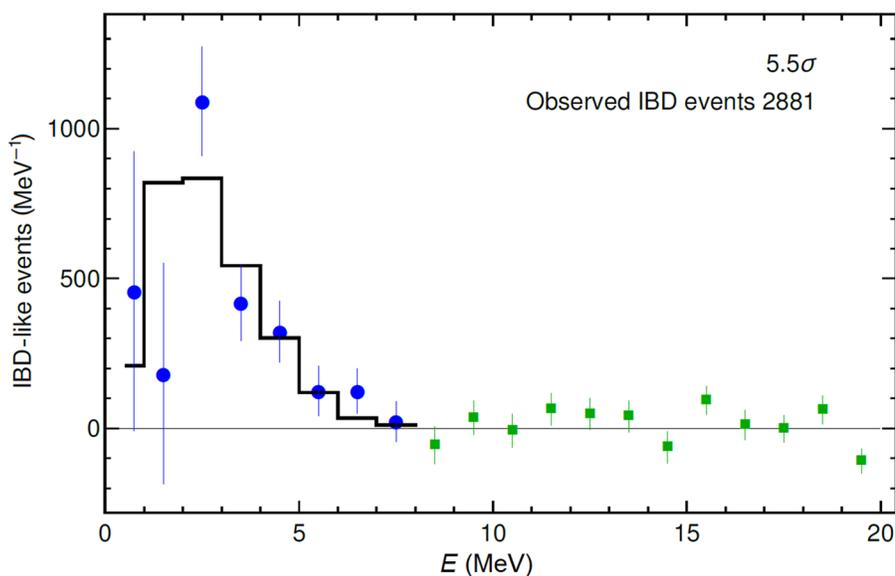

FIG. 7. Difference between reactor-on correlated events and reactor-off correlated events. The data points in green are used to determine the reactor-off normalization in this subtraction. The blue data points are in the IBD-signal region, and the histogram is the best-fit Monte Carlo IBD spectrum.

core inventory of plutonium isotopes. All use-case scenarios of reactor neutrino monitoring that go beyond mere reactor-on-off detection require a high-statistics measurement of the neutrino spectrum. Plutonium-production reactors typically have a thermal power in the $20 - 200$ MW$_{th}$ range, thus requiring a fairly sizable active detector mass. For instance, the case studies presented in Refs. [3,4] are based on a notional 5-ton 100%-efficient detector. In order to stay within the weight limits of a typical shipping container, a detector module should not exceed 20 metric tons, translating to a required overall neutrino detection efficiency of about 25%. Furthermore, while a liquid scintillator may not be a technical impossibility for safeguards applications, it would require significant engineering controls to be practical, making this technology easier for a host country to reject. In summary, the results presented here establish a highly efficient liquid-free unshielded detector with full spectral-measurement capabilities, in other words, a real step in the direction of a practical safeguards detector.

The MiniCHANDLER project is undertaken with the singular goal of demonstrating the detection of reactor neutrinos and their energy spectrum with this alternative technology. Bench tests with our MicroCHANDLER prototype show that a combination of new PMTs (Hamamatsu R6321-100) and light guides improves the energy resolution by a factor of two over the use of Amperex XP2202 PMTs alone, as implemented in this version of MiniCHANDLER. Critically, the proposed alternative optics provide a clean resolution of the Compton edge of the 511 keV gamma rays, which will allow us to implement topological selections with greatly improved fast-neutron rejection efficiency.

Other future improvements include an upgrade of the electronics, based on the readout of the SoLid detector [38]. This will have at least three known benefits: (1) increasing the dynamic range by a factor of 4, (2) repairing an undershoot or overshoot in the analog signal affecting high-primary-energy event pairs with $\Delta t < 40\ \mu$s, and (3) eliminating electronic crosstalk. Additionally, we will double the ^{6}Li concentration by putting a neutron sheet in the middle of each cube layer. Simulations show that this so called "half-cube" modification should increase the capture efficiency of the ^{6}Li by 35%, while decreasing the capture time by 48% [39]. After returning from North Anna, we test this configuration by modifying a single layer of the MiniCHANDLER detector. We find that it reduces the capture time and increases the ^{6}Li capture rate in agreement with the simulations, while having no measurable effect on the light collection. Finally, simulations show that adding just a meter of water-equivalent shielding would reduce the fast-neutron background by an order of magnitude [39]. Future deployments of CHANDLER detectors will likely be accompanied by a water

tank, which can be filled on site to provide an overburden of up to 1 m.

With the aforementioned improvements, we expect to achieve a signal-to-noise ratio of better than 1:1, which is essential for the safeguards goal of determining the plutonium content from distortions in the neutrino spectrum.

[1] A. A. Borovoi and L. A. Mikaelyan, Possibilities of the practical use of neutrinos, Sov. At. Energy **44**, 589 (1978).

[2] Patrick Huber, On the determination of anti-neutrino spectra from nuclear reactors, Phys. Rev. **C84**, 024617 (2011).

[3] Eric Christensen, Patrick Huber, and Patrick Jaffke, Antineutrino reactor safeguards – A case study, Sci. Global Secur. **23**, 20 (2015).

[4] Eric Christensen, Patrick Huber, Patrick Jaffke, and Thomas E. Shea, Antineutrino Monitoring for Heavy Water Reactors, Phys. Rev. Lett. **113**, 042503 (2014).

[5] Adam Bernstein, Yi-fang Wang, Giorgio Gratta, and Todd West, Nuclear reactor safeguards and monitoring with antineutrino detectors, J. Appl. Phys. **91**, 4672 (2002).

[6] A. Bernstein, N. S. Bowden, A. Misner, and T. Palmer, Monitoring the thermal power of nuclear reactors with a prototype cubic meter antineutrino detector, J. Appl. Phys. **103**, 074905 (2008).

[7] Vedran Brdar, Patrick Huber, and Joachim Kopp, Antineutrino Monitoring of Spent Nuclear Fuel, Phys. Rev. Appl. **8**, 054050 (2017).

[8] Adam Bernstein, Nathaniel S. Bowden, and Anna S. Erickson, Reactors as a Source of Antineutrinos: The Effect of Fuel Loading and Burnup for Mixed Oxide Fuels, Phys. Rev. Appl. **9**, 014003 (2018).

[9] Patrick Jaffke and Patrick Huber, Determining Reactor Fuel Type From Continuous Antineutrino Monitoring, Phys. Rev. Appl. **8**, 034005 (2017).

[10] Rachel Carr *et al.*, Neutrino physics for Korean diplomacy, Science **362**, 649 (2018).

[11] Rachel Carr *et al.*, Neutrino-based tools for nuclear verification and diplomacy in North Korea, Sci. Global Secur. **27**, 15 (2019).

[12] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Detection of the free neutrino: A confirmation, Science **124**, 103 (1956).

[13] V. A. Korovkin, S. A. Kodanev, N. S. Panashchenko, D. A. Sokolov, O. M. Solovyanov, N. D. Tverdovskii, A. D. Yarichin, S. N. Ketov, V. I. Kopeikin, I. N. Machulin, L. A. Mikaelyan, and V. V. Sinev, Measuring nuclear plant power output by neutrino detection, Sov. At. Energy **65**, 712 (1988).

[14] Yu. V. Klimov, V. I. Kopeikin, L. A. Mikaelyan, K. V. Ozerov, and V. V. Sinev, Measurement of variations of the cross section of the reaction $\bar{\nu}_e + p \rightarrow e^+ + n$ in the $\bar{\nu}_e$ flux from a reactor, Sov. J. Nucl. Phys. **51**, 225 (1990).

[15] Yu. V. Klimov, V. I. Kopeikin, L. A. Mikaelyan, K. V. Ozerov, and V. V. Sinev, Neutrino method remote measurement of reactor power and power output, At. Energy **76**, 123 (1994).

[16] N. S. Bowden, A. Bernstein, M. Allen, J. S. Brennan, M. Cunningham, J. K. Estrada, C. M. R. Greaves, C. Hagmann, J. Lund, W. Mengesha, T. D. Weinbeck, and C. D. Winant, Experimental results from an antineutrino detector for cooperative monitoring of nuclear reactors, Nucl. Instrum. Methods **A572**, 985 (2007).

[17] G. Boireau *et al.* (NUCIFER), Online monitoring of the osiris reactor with the nucifer neutrino detector, Phys. Rev. **D93**, 112006 (2016).

[18] S. Oguri, Y. Kuroda, Y. Kato, R. Nakata, Y. Inoue, C. Ito, and M. Minowa, Reactor antineutrino monitoring with a plastic scintillator array as a new safeguards method, Nucl. Instrum. Methods **A757**, 33 (2014).

[19] J. Carroll, J. Coleman, M. Lockwood, C. Metelko, M. Murdoch, Y. Schnellbach, C. Touramanis, R. Mills, G. Davies, and A. Roberts, Monitoring Reactor Anti-Neutrinos Using a Plastic Scintillator Detector in a Mobile Laboratory, arXiv:1811.01006 (2018).

[20] "Final Report: Focused Workshop on Antineutrino Detection for Safeguards Applications," IAEA Workshop, IAEA Headquarters, Vienna, Austria (2008).

[21] Christian Grieb, J. M. Link, and R. S. Raghavan, Probing active to sterile neutrino oscillations in the LENS detector, Phys. Rev. **D75**, 093006 (2007).

[22] Y. Abreu *et al.* (SoLid), A novel segmented-scintillator antineutrino detector, JINST **12**, P04024 (2017).

[23] M. Battaglieri, R. De Vita, G. Firpo, P. Neuhold, M. Osipenko, D. Piombo, G. Ricco, M. Ripani, and M. Taiuti, An anti-neutrino detector to monitor nuclear reactor's power and fuel composition, Nucl. Instrum. Methods **A617**, 209 (2010).

[24] V. K. S. Kashyap, L. M. Pant, A. K. Mohanty, and V. M. Datar, Simulation results of liquid and plastic scintillator detectors for reactor antineutrino detection – A comparison, JINST **11**, P03005 (2016).

[25] D. Mulmule, S. P. Behera, P. K. Netrakanti, D. K. Mishra, V. K. S. Kashyap, V. Jha, L. M. Pant, B. K. Nayak, and A. Saxena, A plastic scintillator array for reactor based anti-neutrino studies, Nucl. Instrum. Methods **A911**, 104 (2018).

[26] Mustafa Kandemir and Altan Cakir, Comparison of plastic antineutrino detector designs in the context of near field reactor monitoring, Nucl. Instrum. Methods **A927**, 353 (2019).

[27] F. P. An *et al.* (Daya Bay), New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay, Phys. Rev. Lett. **115**, 111802 (2015).

[28] https://www.nrc.gov/info-finder/reactors/na2.html.

[29] During the Reactor 2 shutdown, Reactor 1 is operating at full power, so, in the reactor-off subtraction that is performed, its events are counted as part of the background. Therefore, we calculate the expected rate relative to Reactor 2 alone.

[30] http://pdg.lbl.gov/2018/AtomicNuclearProperties/HTML/polyvinyltoluene.html.

[31] A. Dellacqua *et al.*, GEANT-4: An Object oriented toolkit for simulation in HEP, Tech. Rep. (CERN, 1994).

[32] Th. A. Mueller, D. Lhuillier, M. Fallot, A. Letourneau, S. Cormon, M. Fechner, L. Giot, T. Lasserre, J. Martino, G. Mention, A. Porta, and F. Yermia, Improved predictions of reactor antineutrino spectra, Phys. Rev. **C83**, 054615 (2011).

[33] P. Vogel and John F. Beacom, Angular distribution of neutron inverse beta decay, anti-neutrino(e) + p → e+ + n, Phys. Rev. **D60**, 053003 (1999).

[34] J. Ashenfelter *et al.* (PROSPECT), Background radiation measurements at high power research reactors, Nucl. Instrum. Methods **A806**, 401 (2016).

[35] N. Allemandou *et al.* (STEREO), The STEREO experiment, JINST **13**, 07 (2018).

[36] https://www.ncdc.noaa.gov/cdo-web/datatools/lcd, Station ID WBAN:03715.

[37] S. Lindgren, On the pressure dependence of the cosmic ray intensity recorded by the standard neutron monitor, Tellus **14**, 44 (1962).

[38] L. Arnold, W. Beaumont, D. Cussans, D. Newbold, N. Ryder, A. Weber (SoLid), The SoLid anti-neutrino detector's readout system, JINST **12**, C02012 (2017).

[39] W. Walters and A. Haghighat, in *Advances in Nuclear Nonproliferation Technology and Policy Conference (ANTPC), September 25–30* (American Nuclear Society, Santa Fe, NM, 2016).