AIP | **Review of Scientific Instruments**

PRECISION IMAGING OF PARTICLE INTERACTIONS

MUON ENTRY

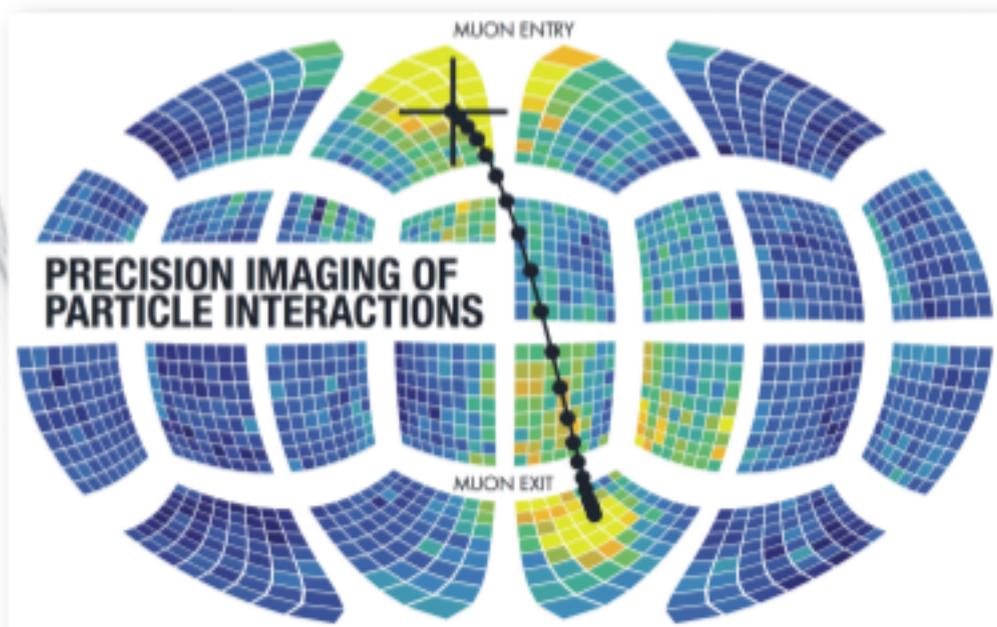MUON EXIT

INVITED ARTICLE:
miniTimeCube
by V. A. Li, R. Dorrill, M. J. Duvall, J. Koblanski,
S. Negrashov, M. Sakai, S. A. Wipperfurth, K. Engel,
G. R. Jocher, J. G. Learned, L. Macchiarulo, S. Matsuno,
W. F. McDonough, H. P. Mumm, J. Murillo, K. Nishimura,
M. Rosen, S. M. Usman, and G. S. Varner

Invited Article: miniTimeCube

V. A. Li,[1,a)] R. Dorrill,[1] M. J. Duvall,[1] J. Koblanski,[1] S. Negrashov,[1,2] M. Sakai,[1]
S. A. Wipperfurth,[3] K. Engel,[3] G. R. Jocher,[4] J. G. Learned,[1] L. Macchiarulo,[1] S. Matsuno,[1]
W. F. McDonough,[3] H. P. Mumm,[5] J. Murillo,[1] K. Nishimura,[4] M. Rosen,[1] S. M. Usman,[6,7]
and G. S. Varner[1]

[1]*Department of Physics and Astronomy, University of Hawaii at Manoa, Honolulu, Hawaii 96822, USA*
[2]*Department of Information and Computer Sciences, University of Hawaii at Manoa,*
Honolulu, Hawaii 96822, USA
[3]*Department of Geology, University of Maryland, College Park, Maryland 20742, USA*
[4]*Ultralytics LLC, Arlington, Virginia 22203, USA*
[5]*National Institute of Standards and Technology, 100 Bureau Dr., Gaithersburg, Maryland 20899, USA*
[6]*Exploratory Science and Technology Branch, National Geospatial-Intelligence Agency,*
Springfield, Virginia 22150, USA
[7]*Department of Geography and Geoinformation Science, George Mason University,*
Fairfax, Virginia 22030, USA

We present the development of the miniTimeCube (mTC), a novel compact neutrino detector. The mTC is a multipurpose detector, aiming to detect not only neutrinos but also fast/thermal neutrons. Potential applications include the counterproliferation of nuclear materials and the investigation of antineutrino short-baseline effects. The mTC is a plastic 0.2% ^{10}B–doped scintillator (13 cm)^3 cube surrounded by 24 Micro-Channel Plate (MCP) photon detectors, each with an 8×8 anode totaling 1536 individual channels/pixels viewing the scintillator. It uses custom-made electronics modules which mount on top of the MCPs, making our detector compact and able to both distinguish different types of events and reject noise in real time. The detector is currently deployed and being tested at the National Institute of Standards and Technology Center for Neutron Research nuclear reactor (20 MW$_{th}$) in Gaithersburg MD. A shield for further tests is being constructed, and calibration and upgrades are ongoing. The mTC's improved spatiotemporal resolution will allow for determination of incident particle directions beyond previous capabilities. © *2016 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).* [http://dx.doi.org/10.1063/1.4942243]

I. INTRODUCTION: THE MOTIVATION BEHIND COMPACT NEUTRINO DETECTORS AND THE miniTimeCube (mTC)

A number of fundamental mysteries remain in the field of neutrino physics, for instance the structure of the mass hierarchy of the three known neutrinos, and the possible existence of sterile neutrinos that interact only through mixing. Further, their ultimate nature as Majorana or Dirac fermions has yet to be determined. At the same time, our understanding of neutrinos has reached a turning point where practical applications of neutrino detection are becoming increasingly feasible. This understanding, combined with recent developments in the areas of fast photodetectors, high-quality doped scintillators, electronics, and computing, has led to the possibility of a new generation of compact, highly instrumented neutrino detectors that were previously impractical and unaffordable. These detectors will allow exploration of fundamental neutrino properties, as well as practical applications in the fields of reactor safety and nuclear security.

Specifically, measurements of reactor antineutrinos provide a number of scientific opportunities, such as the detailed study of neutrino oscillations at very short baselines, and investigation of the reactor antineutrino anomaly, which may be connected to the existence of sterile neutrinos.[1] The mTC, shown in simulation in Fig. 1, represents a new step in this direction. The mTC is a compact (∼2200 cm^3 active volume), densely instrumented, fast timing plastic-scintillator detector designed as a proof-of-concept for future reactor antineutrino detectors. In addition to these scientific studies, the mTC is also designed for practical applications, such as directional neutrino detection and reactor monitoring for nonproliferation.

A. The history and inception of the mTC

The motivation for a compact neutrino detector began with a study involving National Geospatial-Intelligence Agency (NGA), Integrity Applications Incorporated (IAI), and UH personnel in 2011. It is an evolution of a CCD-based detection concept,[2,3] which was found to have issues with scalability to large detection volumes. To avoid similar problems, the mTC utilizes time as an extra dimension to reconstruct the event kinematics. In the mTC concept, a Fermat surface is defined by the first light arrival, leading to spatial and angular resolutions well below what one would expect from the scintillator decay
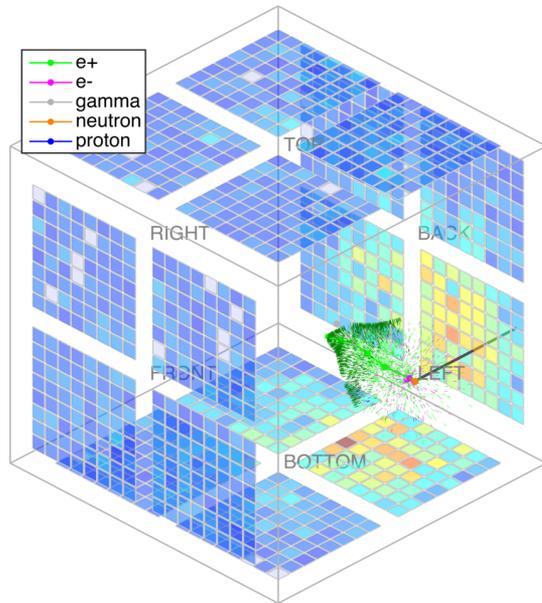

87, 021301-1

FIG. 1. GEANT and MATLAB simulation of a 10 MeV $\bar{\nu}_e$ interaction in the 13 cm cubical mTC. Photons colored identically to parent particles.

times.[4] This leads to particle location resolutions on the order of millimeters instead of the meter scale one would naively expect from scintillator decay time constants.

B. Technological context

The mTC concept requires excellent single photon timing resolution, which is achieved using commercial micro-channel plate (MCP) photomultiplier tubes with excellent intrinsic timing (~50 ps). Combined with readout electronics we expect single photon timing resolutions of 100 ps or better, corresponding to about 2 cm spatial resolution in the scintillator. Further improvement is achieved by multiple pixel constraints, roughly scaling as $1/\sqrt{N_{\text{pe}}}$.

The mTC's state-of-the-art fast-timing and pixelization allow many novel measurements. Although its small size may prohibit full investigation of some of the proposed applications, it serves as a proof-of-concept and model for future detectors such as NuLat.[5]

The preliminary design of the detector and initial performance simulations was conducted in 2011, with construction starting in the same year. The initial version of the detector, shown in Fig. 2, was completed at the end of 2013. In January 2014, we started testing and calibrating the detector at National Institute of Standards and Technology (NIST). A number of upgrades have been performed or are underway as a result of lessons learned from these initial studies. We expect to begin operation at the NIST reactor, pending installation of a shielding cave to reduce neutron backgrounds, in late 2015.

C. Design of the mTC

The core detection volume of the mTC is a $(13 \text{ cm})^3$ cube of plastic scintillator (Eljen Technology EJ-254), doped with 1% natural boron (0.2% ^{10}B).[6] The scintillator decay constant is 2.2 ns.
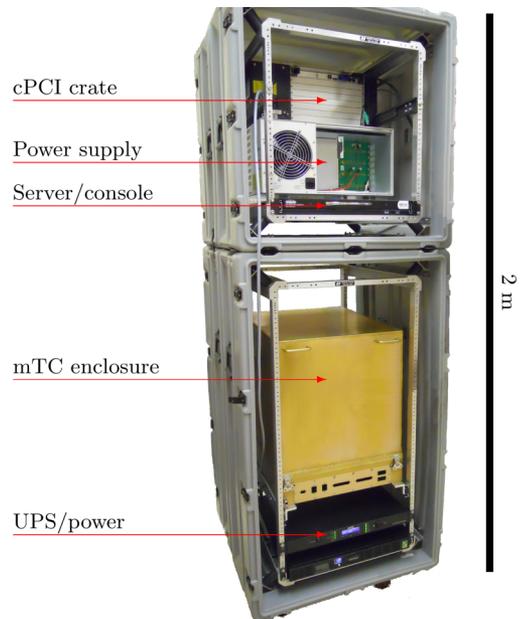

FIG. 2. Photograph of mTC's mount-racks, light-tight aluminum enclosure, data acquisition system, and power supplies.

A total of 24 PLANACON MCP-Photo-Multiplier Tubes (PMTs) (PHOTONIS XP85012), hereafter referred to as simply "MCPs," shown in Fig. 3,[7] are used to detect photons from the scintillator volume. They are coupled to the scintillator cube using optical grease (Eljen Technology EJ-550) and clamped in place to the cube for mechanical stability. The anode plane of each MCP is segmented into 64 pixels, leading to a total of 1536 readout channels. The scintillation and Cherenkov spectra expected for EJ-254 are shown in Fig. 4, along with the typical quantum efficiency (QE) curve of the MCP, showing matching of the QE to the scintillation spectrum. The sensitivity of the scintillator, including coverage factors and detection efficiency of the MCPs, is ~1000 photoelectrons/MeV.

The 1536 MCP readout channels are instrumented with custom electronics developed at the University of Hawaii. These electronics mount directly to the MCPs, providing multi-gigasample per second sampling and on-board digitization of the MCP signals with a timing resolution of <100 ps. By preserving the excellent timing resolution of the MCPs, we retain the ability to study advanced reconstruction techniques
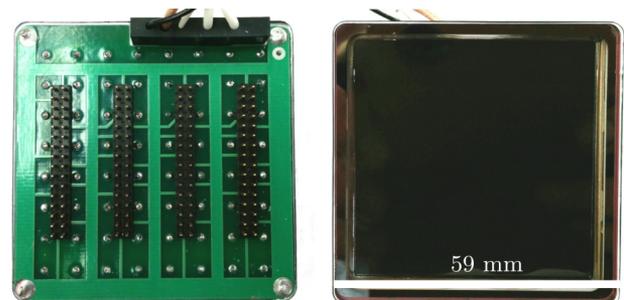

FIG. 3. Photograph of a PHOTONIS PLANACON MCP XP 85012, one of 24 MCPs used in the mTC.
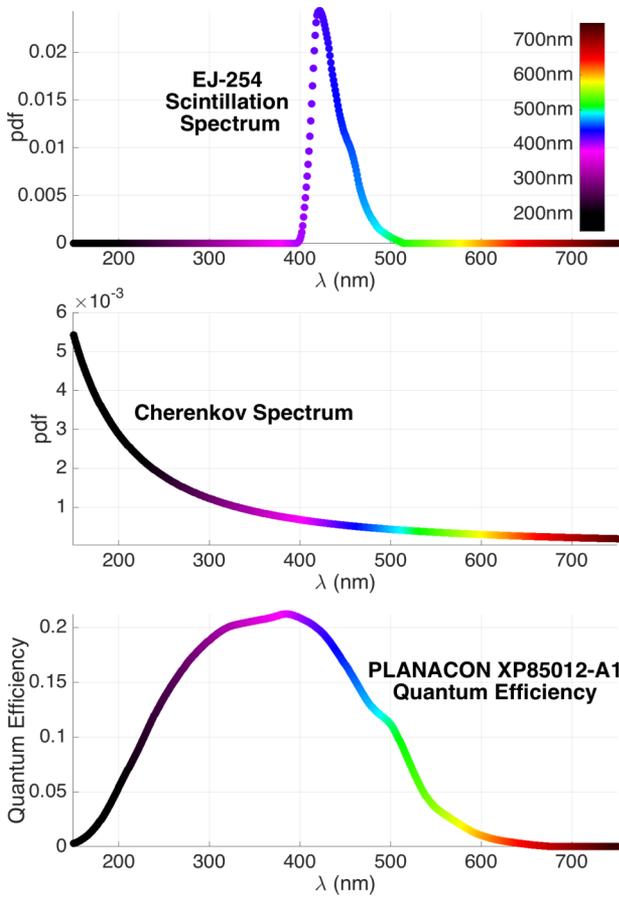
FIG. 4. Scintillation, Cherenkov, and QE spectra for the mTC. GEANT and MATLAB MC models include all effects of chromatic dispersion. Most Cherenkov photons in the UV region attenuate very quickly and are not observed.

(e.g., incorporation of the shape of the scintillator decay time distribution and the fast timing of the Cherenkov photons). A model of the scintillator cube with one face of photodetectors and corresponding readout electronics populated is shown in Fig. 5. The compact nature of the readout electronics keeps the core of the mTC compact. The net dimensions of the cube, MCPs, and electronics fit inside an ~1/8 m^3 volume. The electronics is discussed in more detail in Section IV.
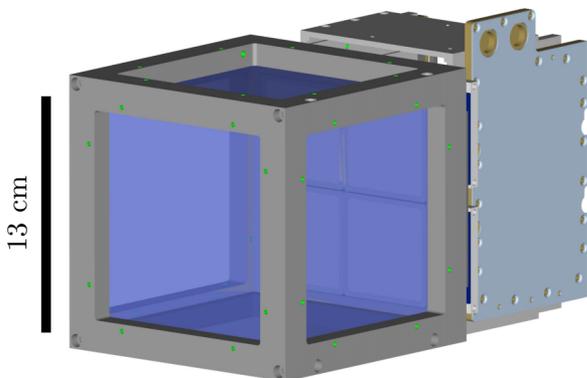

FIG. 5. CAD of the mTC scintillator cube with one face populated with four MCPs and two electronics board stacks connected.

The main detector, ancillary electronics, and power supplies fit in stacked plastic cases, with a clearance footprint of 0.75 m wide by 1.2 m deep by 2.5 m high, and require only 115 V AC and a network connection for remote operation. The assembled and integrated mTC, including associated servers for data acquisition, is shown in Fig. 2. A water-based chiller, with flow around 8 LPM, provides cooling needed for operation in the shielded enclosure. The power consumption is roughly 2 kW, including ~1 kW from the chiller itself. The size and power consumption make this a relatively portable detector, capable of being operated from a truck or a ship.

II. NEUTRINOS IN MTC

The process of identifying a neutrino interaction in the mTC is similar to the one first used in the Reines and Cowan experiment[8,9] in 1956, and used by many experiments since.

An electron antineutrino emitted from a nuclear reactor interacts with a proton inside the plastic scintillator medium, producing a positron and a neutron via inverse beta-decay (IBD),

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

This reaction has a cross section of $\sigma_{tot} \cong 5 \times 10^{-43}$ cm^2 at a neutrino energy $E_\nu = 2$ MeV and an energy threshold of $E_\nu = 1.806$ MeV (in the lab frame, where the proton is at rest). The characteristic time scale between *prompt* (positron annihilation) and *delayed* (neutron capture) signals is used as the primary signature for identifying neutrino events. As outlined below, the positrons scatter nearly isotropically after the neutrino interaction, with the positron taking most of the kinetic energy and the neutron taking most of the momentum. If one records the direction and energy of the positron and the first scatter of the neutron, one can back-reconstruct the incident direction of neutrino. Additionally, further scatters of the neutron can also be used to improve the reconstruction.

A. Prompt signal

The IBD prompt signal generates anywhere from several hundred to several thousand Photo-Electrons (PEs) in the mTC, as shown in Figs. 6 and 7. The energy of this signal is used to reconstruct the incoming $\bar{\nu}_e$ energy, and the location of the signal may be used for directional determination of the $\bar{\nu}_e$ angle, however weakly, by pairing it with the delayed signal location. The prompt signal is composed of a short positron track (~1 cm), any electrons it may interact with, and two equal and opposite 511 keV gammas produced upon positron annihilation.

Simulation projects that the mTC reaches about 10%-15% $\bar{\nu}_e$ energy resolution, ultimately limited by its small size. The positron track at the lower range of the $\bar{\nu}_e$ energy spectrum is on the order of a cm, though higher energy $\bar{\nu}_e$'s will produce longer tracks proportional to their energy above the $E_{\bar{\nu}_e} > 1.8$ MeV threshold. As a result, higher energy $\bar{\nu}_e$'s tend to produce positrons which leave the detector with ever greater likelihood, causing a certain amount of energy under-estimation at higher $\bar{\nu}_e$ energies.
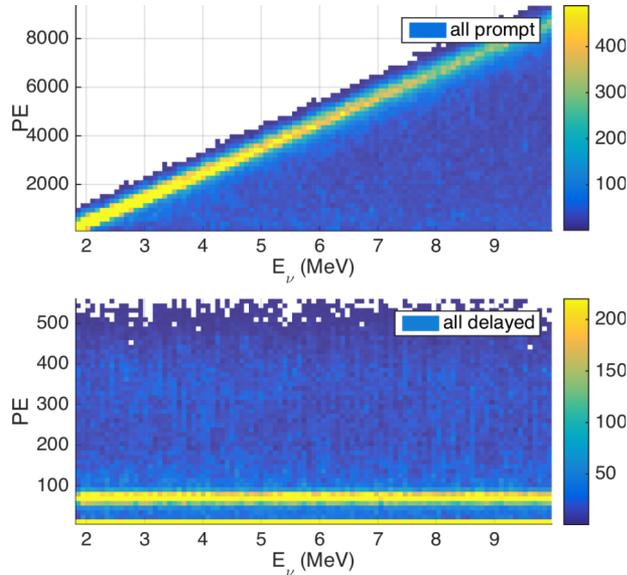
FIG. 6. Simulated distributions of the number of PE produced as a function of $\bar{\nu}_e$ energy for the prompt signal (top) and delayed signal (bottom). In top figure, a long tail of under-estimated energies is produced by longer positrons leaving the detector. In the bottom figure, the long tail of higher energy delayed events is due to the 478 keV gamma produced on neutron capture depositing part of it's energy randomly inside the detector.

In addition to severed positron tracks, a second problem arises at all $\bar{\nu}_e$ energies: uncertainty in prompt signal is introduced via the two 511 keV gammas. In a larger detector such as KamLAND, the annihilation gammas typically deposit their full energy within the scintillation volume. In a very small detector like mTC, these gammas deposit varying amounts of energy from event to event, smearing the prompt energy resolution. On average, the annihilation gammas deposit about 1/3 of their energy in the mTC, but the proportion varies event to event and is impossible to predict *a priori* for a specific IBD event.

More information on expected energy resolutions can be found in Section VI.

B. Delayed signal

The neutron from the neutrino interaction scatters elastically on the scintillator medium and, after thermalizing, captures on the ^{10}B embedded in the scintillator. On average, the neutron travels for a few centimeters before being captured, as shown in Fig. 8.

The primary reaction for neutron capture is[10–12]

$$
\begin{aligned}
n + {}^{10}\text{B} &\rightarrow {}^{7}\text{Li}(1015\text{ keV}) + {}^{4}\text{He}(1775\text{ keV}), &\sim\!6\% \\
&\rightarrow {}^{7}\text{Li}^{*} + {}^{4}\text{He}(1471\text{ keV}), &\sim\!94\% \\
&\hookrightarrow {}^{7}\text{Li}^{*} \rightarrow {}^{7}\text{Li}(839\text{ keV}) + \gamma(478\text{ keV}). &(2)
\end{aligned}
$$

The cross section for neutron capture on ^{10}B as well as the linear attenuation coefficient as a function of neutron energy has been studied.[6,12–14] For a completely thermal neutron ($E_n = 0.025$ eV), the total cross section is equal to 3836 b. The neutron from an IBD reaction can also be captured on a proton in $\sim\!180$ μs, resulting in 2.2 MeV γ's depositing energy via Compton scattering.[15] The fraction of thermal neutrons captured on ^{10}B is 25.6 times larger than on ^{1}H.[6]

To confirm our understanding of the Monte Carlo (MC) results shown in Fig. 8, we can estimate the detector's neutron-capture efficiency analytically. Because the majority of neutrons capturing on ^{10}B is so large, neutron thermalization and capture on ^{10}B are the dominating process in determining the detector's neutron-capture efficiency. We can therefore get a rough estimate of this number by considering a typical neutron undergoing this process. For this calculation, we treat the neutron's path as a random-walk series of elastic scatters on ^{1}H in two parts: (1) production to thermalization, and (2) thermalization to capture. Combining these two results will
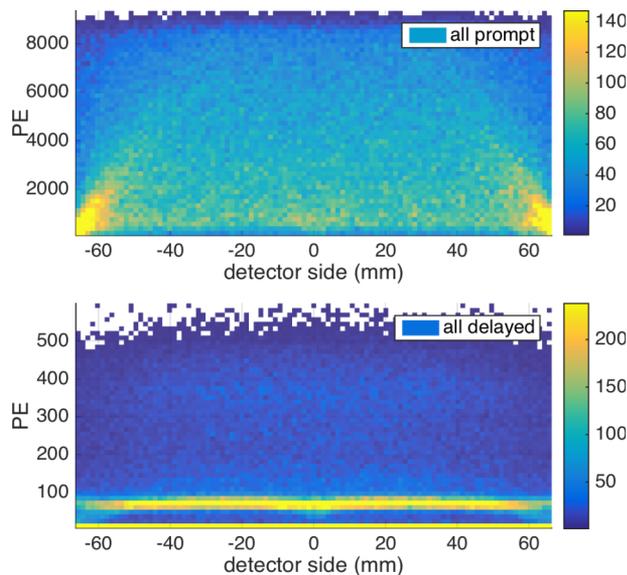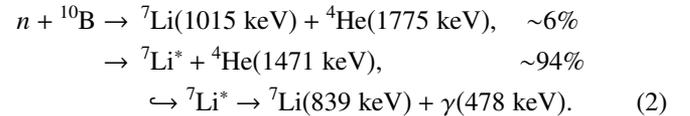

FIG. 7. Simulated distributions of the number of PE produced for the prompt signal (top) and delayed signal (bottom) as a function of vertex location. 0 mm is the center of the detector, and 67 mm is at the edge.
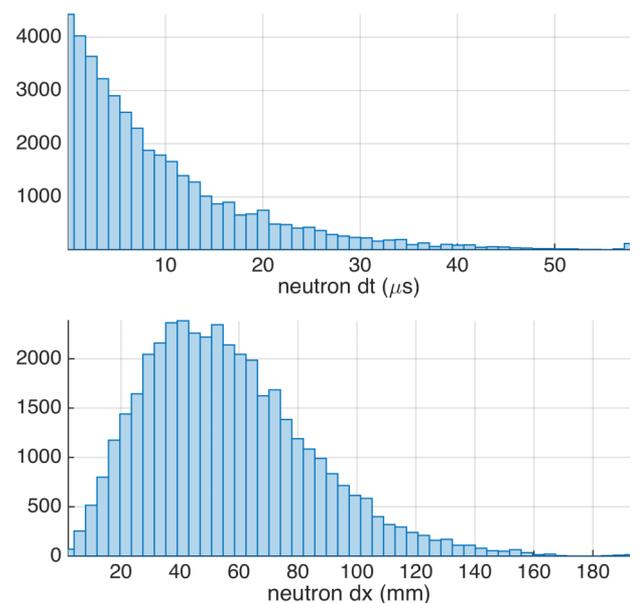

FIG. 8. Simulated frequencies, in arbitrary units, of Monte Carlo generated IBD events as a function of time and distance of neutron production to neutron capture in the mTC scintillator.

give us a general idea of where neutrons are likely to be lost and therefore what fraction of neutrons should capture without escaping the cube.

We will take a typical IBD neutron to have $K_{init} \sim 4$ keV. At these neutron energies, we can safely use nonrelativistic kinematics: $K \sim O(1 \text{ keV}) \ll m_0$.

(1) Production to thermalization, d_{therm}.

$$d_{therm} \sim \lambda_{es} \sqrt{N_{therm}}, \tag{3}$$

where λ_{es} is the mean free path (MFP) for elastic scattering and N_{therm} is the number of steps to thermalization. We can get λ_{es} from the cross section σ_{es} (20 b for hydrogen) and the volume density of targets n_H,

$$\lambda_{es} \sim \frac{1}{n_H \, \sigma_{es}}. \tag{4}$$

To find N_{therm}, we assume that on average the neutron loses half of its excess KE on each collision with a proton. Then N_{therm} is simply,

$$N_{therm} \sim \log_2\left(\frac{K_{init}}{K_{therm}}\right). \tag{5}$$

Combining the above equations and data from the scintillator manufacturer,[6] we get ∼17 steps at ∼1 cm each for a distance of

$$d_{therm} \sim 4 \text{ cm}. \tag{6}$$

(2) Thermalization to capture, d_{cap}.

Once thermalized, the neutron will typically capture after traveling the corresponding MFP, λ_{cap}. Because this distance is longer than λ_{es} above, the neutron will continue its elastic scattering on H during this time, with a number of steps equal to $\lambda_{cap}/\lambda_{es}$. The displacement for this part of the process is therefore

$$d_{cap} \sim \lambda_{es} \sqrt{\frac{\lambda_{cap}}{\lambda_{es}}} \sim \lambda_{es} \sqrt{\frac{1}{\lambda_{es}} \cdot \frac{1}{n_{B10} \, \sigma_{cap}}} \sim 1.5 \text{ cm}. \tag{7}$$

(3) Total (production to capture), d_{tot}.

Finally, taking d_{therm} and d_{cap} to be two steps of a random walk, we have

$$d_{tot} \sim \langle d \rangle \sqrt{N} = \frac{d_{therm} + d_{cap}}{2} \sqrt{2} \sim 4 \text{ cm}, \tag{8}$$

as the typical total distance between production and capture (again, ignoring corrections for effects like capture before thermalization, capture on hydrogen, etc.).

(4) Neutron-capture efficiency, $n_{captured}/n_{total}$.

We can use an imaginary sphere of radius $d_{tot} \sim 4$ cm to roughly estimate the capture rates in various regions of the cube. For example, a neutron produced at the surface of the cube but near the center of a face will generally have ∼1/2 probability to move inward and capture or to move outward and escape; however, a neutron produced deeper than d_{tot} ∼ 4 cm into the face will most likely capture inside the cube. Averaging over the depth indicates that ∼3/4 of the neutrons produced in this region should capture inside the cube. We can make similar estimates for the rates in the other regions of the cube (i.e., edges, corners, and interior) and combine these

estimates to get our final result,

$$\frac{n_{captured}}{n_{total}} \sim \frac{1}{2}. \tag{9}$$

This is in general agreement with simulations indicating that ∼55% of the neutrons produced inside the cube capture without escaping.

Using this same approach, we can get a rough estimate for how long this process might take $t_{tot} = t_{therm} + t_{cap}$. It is relevant to keep in mind that N_{therm} depends on K_{init} as discussed above.

(1) Production to thermalization, t_{therm}.

As in the distance calculation, we will assume that on average, the neutron loses half of its kinetic energy on each collision. (For this calculation, we approximate $K_{therm} \sim 0$.) After n collisions, this becomes $K_n = 2^{-n} K_{init}$. We then immediately have

$$v_n = \sqrt{\frac{2K_n}{m}} = 2^{-n/2} \sqrt{\frac{2K_{init}}{m}} = 2^{-n/2} v_{init}. \tag{10}$$

Since the distance travelled in each step is λ_{es} from above, the time for each step is

$$t_n = \frac{\lambda_{es}}{v_n} = 2^{n/2} \frac{\lambda_{es}}{v_{init}}. \tag{11}$$

The total time to thermalization is then the sum of these steps over N_{therm} terms,

$$t_{therm} \sim \sum t_n = \sum_{n=0}^{N_{therm}} 2^{n/2} \frac{\lambda_{es}}{v_{init}} \sim 10 \ \mu s. \tag{12}$$

(2) Thermalization to capture, t_{cap}.

The average speed after thermalization is constant by definition, so the time to capture will be λ_{cap} from above divided by this speed,

$$t_{cap} \sim \frac{\lambda_{cap}}{v_{N_{therm}}} = 2^{N_{therm}/2} \frac{\lambda_{cap}}{v_{init}} \sim 10 \ \mu s. \tag{13}$$

(3) Total (production to capture), t_{tot}.

Combining these results gives us as follows:

$$t_{tot} = t_{therm} + t_{cap} \sim 20 \ \mu s, \tag{14}$$

which is also in general agreement with the Monte Carlo.

Reconstruction of the neutrino's direction largely depends on the neutron direction reconstruction, improved by the positron direction and energy. Having one or more neutron scatters improves the resolution, but even the neutron capture location after many scatters retains information on the initial neutron direction, as was demonstrated in the CHOOZ experiment.[16] Full reconstruction algorithms, currently under development, will take all the information into account in solving for incoming neutrino direction.

The light yield of these neutron scatters can present some difficulty when detecting and reconstructing events. Ionization density quenching on two charged particles (^{4}He and ^{7}Li) with ≳2.3 MeV kinetic energy in the reaction, Eq. (2), results in a small total light output, about 60 keV electron-equivalent energy deposition.[17] However, due to the small size of the detector and high MCP surface coverage, mTC has the high light collection efficiency crucial to detect the weak light from these
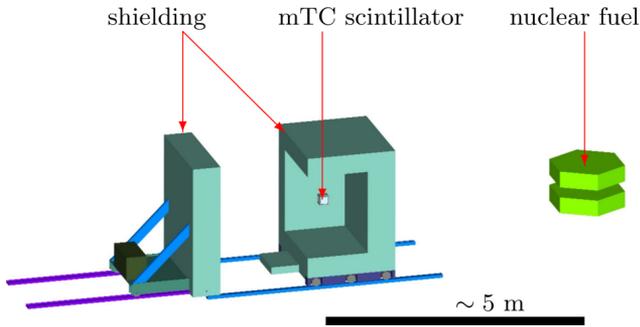
FIG. 9. Relative location of the scintillator cube inside the movable cave (one face made transparent) with respect to the reactor core (upper and lower fuel segments are approximated by two hexagonal prisms).

delayed signals. As a result, only in relatively small-volume (~liter sized), ^{10}B-doped scintillator detectors can incident antineutrino direction currently being reconstructed based on neutron directionality.[17]

C. mTC at NIST reactor

The mTC currently sits on-site at the NIST Center for Neutron Research (NCNR), which houses the NIST 20 MW$_{th}$ split-core research reactor. This reactor has a compact core (Figs. 9 and 10) with 30 fuel elements, each containing 2 segments of highly enriched uranium fuel U_3O_8/Al (^{235}U, 93% enrichment). Fuel elements are submerged in heavy water which serves as a moderator and coolant. The upper and lower fuel segments, each 27.9 cm high, are separated by a 17.8 cm unfueled gap which serves as a "flux trap" to minimize the fast-neutron and gamma backgrounds in the neutron beam lines. The overall dimensions of the core are 1.12 m in diameter by 0.74 m in height. The NIST reactor cycle is 38 days on followed by 10 days off for refueling.

Full Monte Carlo N-particle (MCNP) simulations of the core are available to onsite collaborations.[18,19]
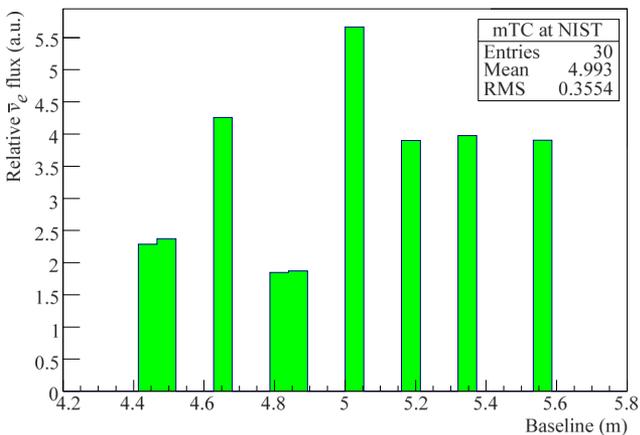

FIG. 10. Relative distribution of $\bar{\nu}_e$ flux as a function of baseline from a nominal mTC position to each fuel element in the core. The mean source location of flux is at ~5 m and the effective spread is 0.36 m, or an inherent smearing of about 7% on the baseline. Specifics of this distribution will vary by fuel loading conditions.

FIG. 11. Photograph of the mTC shielding next to the reactor.

Using a total thermal power of 20 MW$_{th}$, an average number of 6 $\bar{\nu}_e$ produced per fission (with ~1.5 $\bar{\nu}_e$ above IBD threshold), and thermal energies released per fission of ^{235}U and ^{238}U, one can roughly estimate the total number of neutrinos produced at the reactor core to be ~4×10^{18} s^{-1} $\bar{\nu}_e$. This corresponds to a flux of ~1.1×10^{12} cm^{-2} s^{-1} $\bar{\nu}_e$ at the miniTimeCube location ~5 m away from the center of the reactor core, Fig. 11. It further corresponds to a number of antineutrino interactions with ^{1}H via IBD reaction in the plastic scintillator on the order of a few events per day.

More precisely, the total number of expected antineutrinos from the reactor observed in the detector is given by

$$N_{\bar{\nu}_e}^{\text{obs}} = \frac{N_p}{4\pi L^2} \int \epsilon_{det} P(\bar{\nu}_e \to \bar{\nu}_e) \frac{d\sigma(E_{e^+}, E_{\bar{\nu}_e})}{dE_{e^+}}$$
$$\times \frac{d^2 N_{\bar{\nu}_e}(E_{\bar{\nu}_e}, t)}{dE_{\bar{\nu}_e} dt} dE_{e^+} dE_{\bar{\nu}_e} dt, \tag{15}$$

where N_p — number of hydrogen atoms in the scintillator ("free protons"); L — distance between production and detection points of the antineutrino, ϵ_{det} — detector efficiency ~30% based on MC simulations for IBD detection in the mTC; $P(\bar{\nu}_e \to \bar{\nu}_e)$ — survival probability of electron antineutrino;[5,15] and $\frac{d\sigma(E_{e^+}, E_{\bar{\nu}_e})}{dE_{e^+}}$ — differential cross section of the IBD process as a function of positron energy E_{e^+} and antineutrino energy $E_{\bar{\nu}_e}$.[20]

Further details on antineutrino production at reactor facilities can be found in the literature, including fuel time-dependence for the NIST nuclear reactor,[18,19] evaluation of thermal energies released per fission of the four main isotopes,[21] and spectrum of antineutrinos produced from the four main isotopes.[22,23]

In addition to NIST, we have actively considered two other deployment sites: Typical Power Reactors (TPRs) and nuclear-powered ships. Their parameters are listed in Table I.

TABLE I. Approximate parameters at potential mTC deployment sites, including NIST, a typical power reactor, and a nuclear-powered ship. "Compact core" indicates a core where all fuel elements are contained within a few meter radius.

Parameter	NIST	TPR
Power, GW_{th}	0.02	3
⟨Baseline⟩, m	5	25
Fuel	HEU	mixed
Fuel cycle, on/off days	38/10	400/10
Compact core	☑	☐
⟨Event rate⟩, $\bar{\nu}_e$ / day	~1	~10

III. BACKGROUNDS

Backgrounds in the mTC come from several sources. First, there are the "natural sources," most prominently cosmic radiation. Of those, which consist of high energy neutrons, gammas, and muons, along with their collisional products, potentially the most serious for IBD detectors is the nearly irreducible background of some long lived muon-produced isotopes as we discuss in Section III A. The local environmental backgrounds, such as radioactivity, are not as much of a problem as the cosmic ray associated backgrounds. In Section III B we discuss backgrounds relevant to a reactor and specifically the NIST reactor location.

A. Cosmic ray backgrounds

Cosmic rays produce an inescapable background for IBD detectors. Unfortunately, all the reactors to which we may have access are at best a few meters water equivalent (mwe) under the surface. About 2 mwe is enough to shield from extensive air showers, clearing the remnant hadrons and most electromagnetic components. Muons, however, penetrate to the greatest depths, in ever decreasing numbers but increasing mean energies. These muons may generate local particles, and so shielding is somewhat of a double edged sword. Sea-level muons make neutrons and other hadrons in nuclear interactions, though with something on the order of a 2 km radiation length. The mean muon energy at the Earth's surface is about 2 GeV with a penetrating power of about 10 mwe. Muons coming through the mTC (at about 1/s) often (~10%) come with knock-on electrons (Fig. 12). More dangerous are gammas and neutrons, which can fake the prompt signature of a neutrino.

Precise calculations of these rates are difficult because they depend upon details of the overburden, the local geometry, and shielding, in particular. Isotopes and various spallation products of cosmic-ray muons can be a serious background for neutrino signals. Although many of the isotopes can be filtered from analysis using various cuts, long-lived isotopes such as ^{8}He and ^{9}Li may have lifetimes on the order of a second and decay by beta emission into neutron-unstable daughters. These are two backgrounds that can mimic IBD events in the mTC, but are in fact negligible, as we show below.

In order to study this problem in more detail, a GEANT4[24] simulation of the EJ-254 plastic scintillator was conducted. Sea-level spectrum cosmic ray muons were incident on a (10
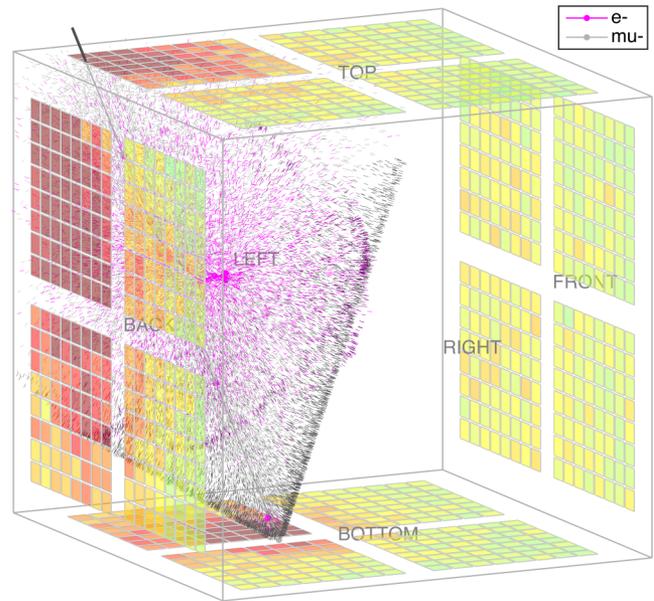

FIG. 12. A simulated muon traversing the mTC, with scintillation photons and Cherenkov cone visible.

$\times 10 \times 10$) m^3 cube of scintillator. The isotope yield per muon event for this simulation run is tallied and shown in Fig. 13.

Figure 14 shows the average secondary particle yield per unit muon track length per unit medium density for simulated cosmic ray muons using a sea-level energy spectrum. The result implies an isotope yield of ~6.86×10^{-10} cm^2/g for ^{9}Li and ~9.79×10^{-11} cm^2/g for ^{8}He. The atmospheric muon
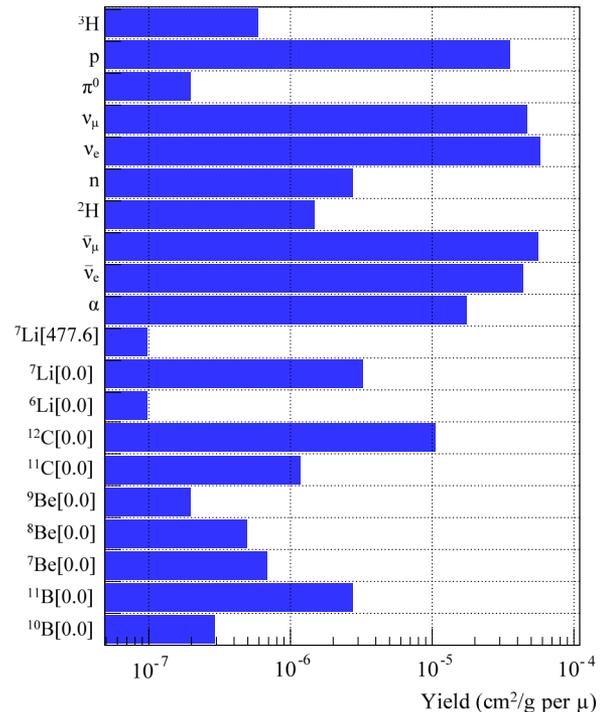

FIG. 13. Cosmogenic isotope production yield due to sea-level spectrum muons passing through 10 m of EJ-254 plastic. 10^4 events were simulated for this result. The number enclosed in brackets in the labels along the y-axis is the excitation energy of the isotope in units of keV. ^{9}Li and ^{8}He were not observed.
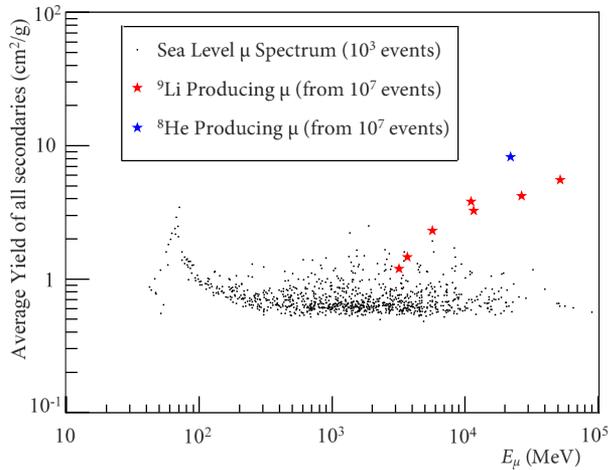
FIG. 14. Average yield per unit muon track length per unit medium density of all non-photon secondary particles versus muon energy. The black points show the result of a run of 10^3 sea-level spectrum muons. Superimposed on the figure are eight specific muon events that had produced a ^{8}He or ^{9}Li daughter depicted by the colored stars. These eight events were extracted from a much larger ensemble with an increased statistics of 10^7 events in order to produce the rare events.

rate traversing the 13 cm cube is about 1/s (depending upon overhead shielding), and so, the rate of these events being produced in the mTC is estimated to be less than 1 event per year. In addition, the general behavior of typical sea-level spectrum muons shown by the black points involves a relatively constant production of secondaries with respect to energies above a few hundred MeV, whereas a trend of increasing daughter production is clearly seen for those producing ^{9}Li. This suggests that the ^{9}Li isotope is most likely produced in showering muon events at high energies, which can be easily vetoed. Rejection of backgrounds associated with ^{8}He will require more statistics and further investigation.

Peripheral geometries of the detector and its *in situ* environment pose a non-negligible contribution to the cosmogenic backgrounds and a more accurate study with these effects fully taken into account will need to be conducted in the future. Design and production of the shielding cave is currently underway and these background studies will be pursued in parallel as development continues.

Finally, with an mTC-type detector, it is straightforward to implement additional vetos to reject backgrounds if needed. This can be accomplished by installing scintillator paddles around the mTC and feeding additional trigger signals to the trigger/clock distribution board.

B. Backgrounds at the NIST reactor location

Detailed background studies must be performed at a particular reactor site, since all venues differ and the backgrounds depend in detail upon local conditions. A group preparing for the PROSPECT experiment carried out a detailed background survey in the mTC location.[25]

Figure 15 shows the gamma spectrum at the proposed mTC location without shielding. The difference between the reactor on/off spectra is readily visible. The "reactor on" spectrum extends to relatively high energies due to prompt gammas
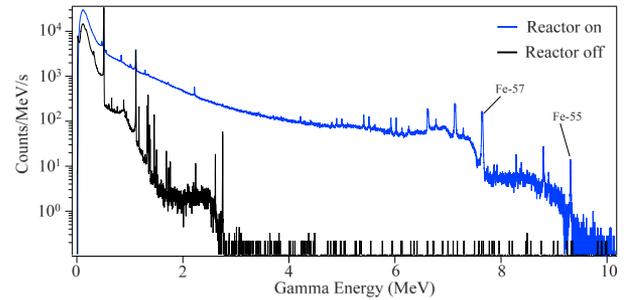

FIG. 15. High Purity Germanium (HPGe) gamma-ray spectrometer response at the mTC location adjacent to the NIST reactor. HPGe spectrometer is 55 mm in length and 62.5 mm in diameter. The observed Fe lines are from neutron capture on surrounding shielding and structural materials.

from neutron capture thus posing additional challenges for shielding.

Deployment of mTC as an antineutrino detector at the NIST reactor, where the backgrounds are particularly high due to adjacent neutron scattering instruments, requires shielding from various background signals that could overwhelm or create false events within the scintillating volume (e.g., high-energy gammas, thermal neutrons, fast neutrons, cosmic ray induced muons, and their decay products). In most other anticipated deployment locations, we do not expect shielding to be critical.

Towards that effort, a multi-layered shielding cave was designed to encase the mTC and most of its associated electronics during the testing at the NIST reactor. The mTC detector will be deployed inside this shielding cave and together they will be placed against the face of the reactor biological shield. The shielding cave is comprised of six nested cubes, with the outermost dimensions yielding a footprint of ~1.8 m × 2 m × 2.7 m and a total internal wall thickness of roughly 0.4 m. From exterior to interior, the layers are as follows (Fig. 16):
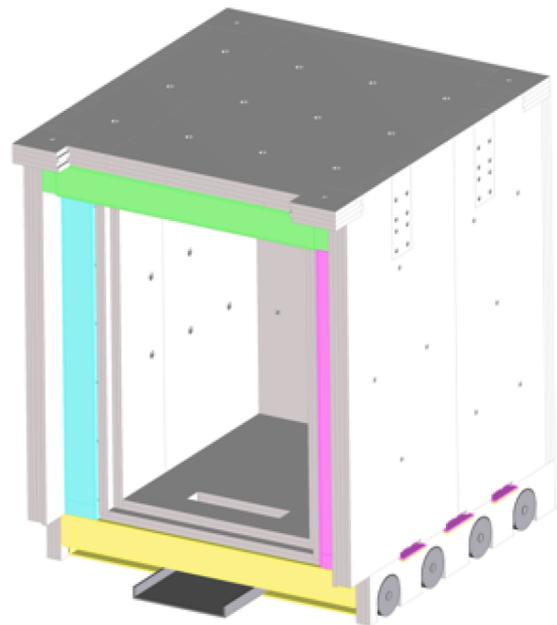

FIG. 16. CAD of the multi-layer shielding for the mTC.

1. 10 cm of 5% borated polyethylene sheet;
2. 1 cm of A36 steel plate;
3. 15 cm of steel shot and paraffin wax mixture;
4. 1 cm of A36 steel plate;
5. 10 cm of 5% borated polyethylene sheet; and
6. interior cavity for housing the mTC and associated electronics (dimensions 1 m × 1.2 m × 1.5 m).

Borated polyethylene was chosen for its neutron absorbing properties while the layer of steel shot and wax acts as both a neutron and gamma absorber. All layers serve to attenuate the muon flux, albeit less efficiently. The overlapping construction removes potential line-of-sight and the hermetic design inhibits the penetration of thermal neutrons, which exhibit gas like properties.

The shielding cave is 20 tons. The entire weight of the cave is supported on rails already laid into the floor at the NCNR, which allows the cave to be moved across different baselines.

The interior of the shielding cave will house the mTC, complete with its electronics rack containing high voltage (HV) power supplies and support electronics. The system requires less than 2 kW of 115 V AC power and has its own uninterruptible power supply system. Accommodations for cooling of the electronics will be used, with access for cooling and electricity through a floor tray. During mTC operation, access to the interior of the cave is expected to be infrequent.

We use GEANT4 to estimate effectiveness of the different shielding layers in attenuating potential backgrounds. The modeled environment includes the shielding cave, a 0.5 m thick concrete roof above the area, and a concrete reactor bio-shield next to the shielding cave. Three primary particles (muons, neutrons, and gammas) and two particle sources (atmospheric and reactor), including their relevant energy spectra and fluxes, were used.

Atmospheric gamma and muon energy spectra were calculated using the Cosmic-RaY (CRY) shower Library.[26] We use the spectrum and flux for atmospheric neutrons from Gordon *et al.*[27] A Maxwell-Boltzmann distribution at 600 K with an integrated flux of 3-4 neutrons cm^{-2} s^{-1} was used to represent reactor neutrons outside the shielding cave. The integrated flux was chosen to match energy-insensitive Bonner ball measurements taken *in situ*. Because the neutron energy spectrum was not measured, a higher characteristic temperature (vs 298 K) allows the simulation to conservatively account for a portion of the reactor neutrons not being thermalized. The reactor gamma spectrum and flux (~200 cm^{-2} s^{-1} above 100 keV) comes from a measurement at NIST adjacent to the proposed mTC site with the reactor on.

For the purposes of preliminary Monte Carlo work, atmospheric particles were assumed to follow a cos$^2 \theta$ angular distribution. Reactor particles were assumed to be isotropic, although significant spatial variation coming from localized source has been measured, and if needed will be incorporated in later work.

Preliminary estimates of the resultant particle fluxes through the mTC volume with and without the cave present are shown in Table II. These include secondaries produced within the shielding material itself. Measurements taken *in situ* useful

TABLE II. Particle fluxes through the mTC volume with and without the shielding cave present. The neutron flux is dominated by near-thermal neutrons hence the large attenuation factor. Muons are incident on the mTC at a rate of less than 3.5 Hz and are not significantly affected by the shielding cave.

Type	Normal	Shielded	Normal	Shielded	Attenuation
	#/mTC/s		#/cm^2/s		%
Neutron	3 391	0.082	4.0	9.7×10^{-5}	99.9%
Gamma	169 015	325	2.0×10^2	3.8×10^{-1}	99.8%

for Monte Carlo validation are planned and will be reported in a future publication.

The Monte Carlo model of the mTC using the shielded fluxes in Table II shows a signal to background (for uncorrelated events only) of roughly 1:1. These uncorrelated events are usually composed of two independent gammas entering the detector within our 12 μs time window, the first creating a false prompt signal and the second a false delayed signal. An order of magnitude less likely is uncorrelated backgrounds in which a neutron creates a false delayed signal instead of a gamma. These simulation results indicate that accidental coincidences from uncorrelated backgrounds will likely not be our dominant background source, and we are beginning to focus more on correlated secondaries originating from high-energy cosmogenic neutrons and muons.

IV. ELECTRONICS

The mTC concept puts stringent requirements on the channel density, timing performance, synchronization, and power consumption of the detector. In order to fully utilize the spatial information provided by the pixelization of the 24 MCP-PMTs, all 1536 channels must be separately instrumented. The readout for each pixel must preserve the O(100 ps) timing provided by the photodetector, and to avoid further timing degradation, the timing of all detector channels must be synchronized to one another at a level significantly below the transit time spread of the MCP-PMT. An online trigger system is required to isolate physics interactions of interest from backgrounds.

A. Front-end electronics

The core of the front-end electronics functionality is provided by the Ice Radio Sampler (IRS),[28] a family of application specific integrated circuits (ASICs) developed at the University of Hawaii. The IRS has been used in a variety of projects that require fast sampling and deep buffering.[29–31]

The IRS ASIC architecture is shown schematically in Fig. 17, and a list of operating parameters can be found in Table III. The ASIC has 8 analog input channels, each with a sampling stage, intermediate and deep storage stages, on-chip digitization, and per-channel threshold triggers. The sampling stage is a multi-GSa/s switched capacitor array (SCA) waveform sampler, similar to other ASICs,[32–34] in which a sampling clock propagates down a delay line, with subsequent delay
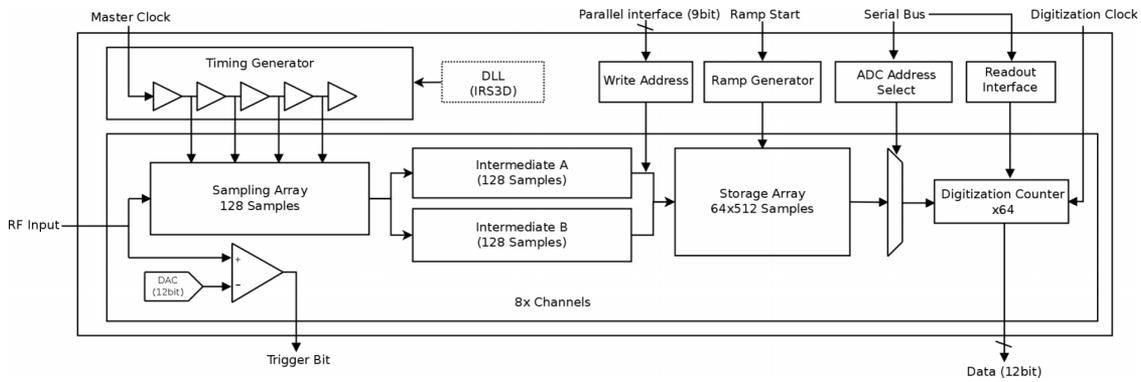
FIG. 17. Block diagram of the IRS ASIC architecture. Eight channels of analog input are received by a set of eight sampling arrays, with sampling timing based on a common timing generator, driven by an external clock. This timing generator also determines timing of transfers from the sampling arrays to intermediate and storage arrays. The target location for the transfer from the intermediate to storage array is controlled by the user with a 9-bit parallel bus. A separate pin is used to start an internal voltage ramp, used to digitize 64-samples of the storage array for all eight channels in parallel. Selection of the storage address to digitize is controlled through a serial interface. A clock for the Wilkinson digitization process is generated internally (IRS3B) or provided externally (IRS3D). Once data are digitized, the channel and sample to readout are controlled by a second independent serial interface. Digitized data are available on a parallel 12-bit bus. A number of DACs and internal timing parameters are controlled by a third serial register interface.

stages utilized to create short, GHz-scale timing intervals to sample the input signal onto capacitors. Unlike other SCA waveform samplers, this sampling array is connected to deeper buffers to allow for higher trigger latencies and larger time records per event. Buffer amplifiers are used to drive the stored voltages from the sampling array into a deep sampling array consisting of 32 768 storage capacitors per channel. This transfer occurs via an intermediate storage array to accommodate the settling time of the buffer amplifiers. Signals required to coordinate the intermediate transfers are provided by an internal timing generator, and the final location of the samples in the deep storage array is provided by a parallel address bus that is driven by the user, allowing for flexible and user-defined memory management schemes. The IRS includes 12-bit Wilkinson ADCs, which digitize 64-sample blocks of the storage array for all 8 channels in parallel. Readout of the digitized data is done one sample at a time through a 12-bit parallel bus. Selection of the channel and sample number is provided by the user via a serial interface. A typical digitized MCP pulse in the mTC system is shown in Fig. 18.

Analog inputs for each channel are also monitored by a comparator, with the digital trigger bits available to the user. These bits can be used to monitor which sections of analog memory have signals above a user-defined threshold, allowing the user to select only those windows which have signals of interest to be read out. Further details on the IRS ASICs and their performance will be presented in an upcoming publication.

One front-end electronics module, or "board stack," shown in Fig. 19, includes 16 IRS ASICs (128 total input channels from 2 MCPs). Each analog input is amplified by an RF amplifier before arriving at the IRS ASIC. The initial version of the mTC was developed with the IRS3B ASIC, and an upgrade is in progress to move to new board stacks using the IRS3D, a new revision of the ASIC that includes improvements to reduce noise and improve timing stability. In both versions of the front-end electronics, a single Field Programmable Gate Array (FPGA) (Xilinx XC6SLX150T)

TABLE III. Operating parameters for the IRS family of ASICs and nominal ASIC operating conditions for the mTC. Full performance parameters will be reported in a future publication.

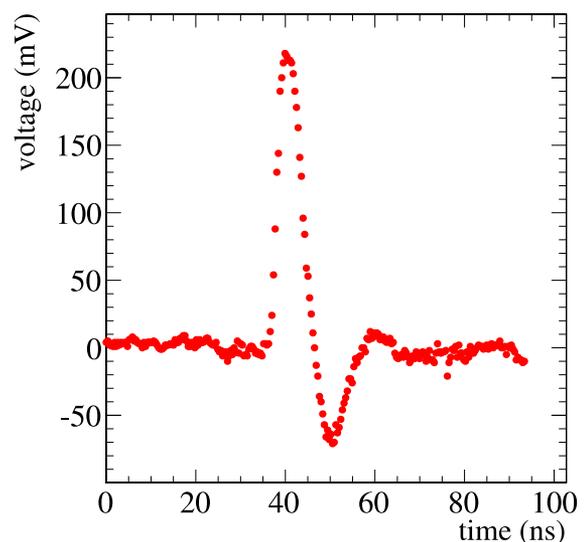| Parameter	IRS range	mTC setting
Channels	8	
Sampling cells	128	
Storage depth	32 768	
Analog bandwidth	>300 MHz	
Digitization	On-chip Wilkinson	
Quantization	12(9)-bits logged(effective)	
Dynamic range	~2 V	
Typical noise	~1 mV$_{RMS}$	
Sampling rate	1–4 GSa/s	2.73 GSa/s
Master clock	8–31 MHz	21.3 MHz
Buffer time	(8–32)μs	12.0 μs
Conversion time	>2 μs	6.2 μs

FIG. 18. Example MCP pulse digitized with the IRS3B. Times for each point are calculated based on a nominal sampling rate of 2.73 GSa/s. Voltages are calculated based on a nominal conversion factor of 0.6 mV/ADC count, and represent the signal after passing through an external amplifier.
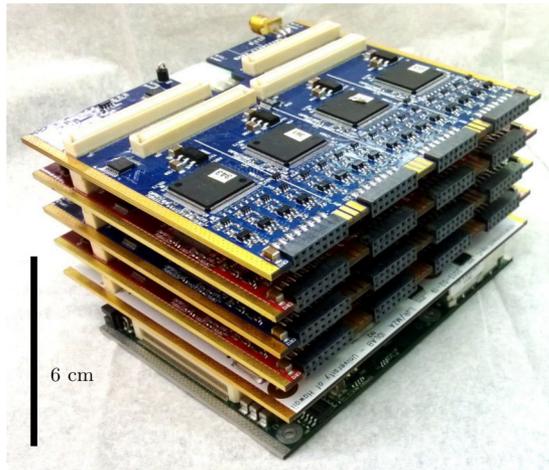
FIG. 19. A photograph of one of the twelve board stack assemblies. The 128 MCP signals are the input via connectors shown. Each board stack instruments a pair of MCPs and is enclosed in an aluminum cage surrounded by chiller plates.

FIG. 20. State machine diagram of the neutrino trigger.

on each board stack provides all control signals necessary to operate and readout the ASICs and other auxiliary devices. The FPGA interfaces to the back-end data acquisition system for register control and data transmission via fiberoptic cable. To coordinate timing between the 12 modules of the mTC, each board stack accepts a central distributed clock via RJ45 connector. Another RJ45 connector is wired to the FPGA JTAG interfaces, allowing remote programming of the FPGA firmware *in situ*.

B. Clock distribution and triggering

All 12 board stacks communicate with a custom Printed Circuit Board (PCB), designated Clock and JTAG In PCI (CAJIPCI), over differential pairs of CAT7a cable. The CAJIPCI provides a low jitter ($\sigma_t < 2$ ps) clock to the front-end modules. Front-end board stacks provide a module-level trigger to the CAJIPCI over another differential pair on the same cable, and the CAJIPCI responds with a system level trigger over a third pair. The final differential pair can be used to perform flow control and limit trigger rates to the front-end electronics.

Three separate trigger levels are defined for the experiment. The lowest level trigger is a level 0 (L0) trigger, defined as a single channel trigger bit from an IRS ASIC. Thresholds for these triggers can be set via adjustment of an on-chip DAC. The 128 L0 triggers on a board stack are monitored by the FPGA. When the number of coincident triggers falls between two user programmable thresholds, a level 1 (L1) trigger is issued and sent to the CAJIPCI. The CAJIPCI, in turn, monitors L1 triggers from the 12 front-end board stacks and issues a level 2 (L2) trigger to the front-end modules under user-defined conditions.

A basic L2 trigger can be calculated based on the number of coincident L1 triggers. This basic trigger is appropriate for signals that fall mainly in a narrow time window (e.g., neutrons, gammas, and cosmic ray muons). A neutrino L2 trigger must monitor for both the prompt and delayed signals, so it includes an initial "arming" period when it detects a prompt
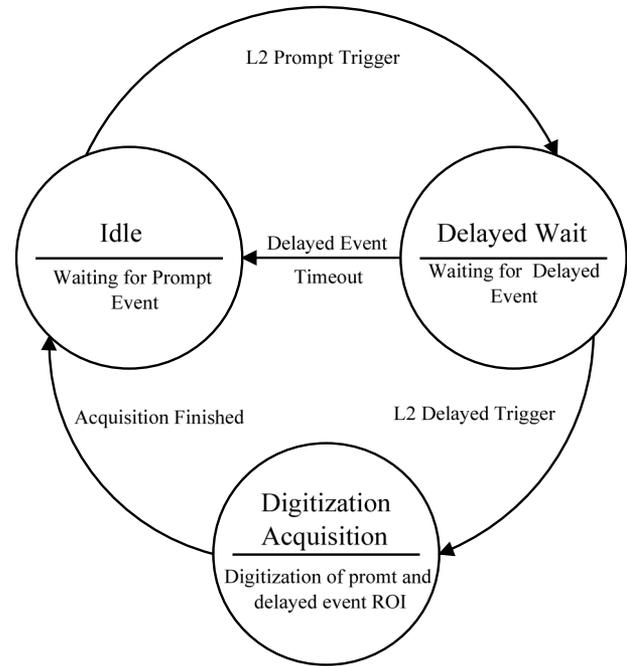
signal, and a second stage to issue a trigger upon receipt of a delayed signal. This logic is shown in Fig. 20. The timeout for the delayed trigger is typically set to ~12 μs, the length of the IRS storage array. Longer times between prompt and delayed signals are possible based on the analog memory management scheme used for the IRS, and this may be explored in future upgrades.

C. Data acquisition and software

Upon receipt of an L2 trigger, data from the front-end modules are digitized, readout by the front-end FPGA, and sent over fiberoptic cables using a gigabit Ethernet (1 GbE) interface. These data are received by commercial PCIe Ethernet cards, running on a rack-mount server, which can be operated directly or via a network connection. Operations from powering up to collecting data may be performed remotely over this network, with an aim for full remote control.

Data acquisition is implemented primarily in C++, with a Python API to the system that can be used for configuration as well as monitoring and slow control of the detector. Several algorithms and programs have been developed to perform automated startup, initialization, and real-time fine tuning of the electronics. Before physics data acquisition begins, threshold scans and pedestals are collected for each channel and stored in files for each board stack. Timing parameters for each chip are taken, adjusted, and stored to ensure optimum calibration. Once this procedure is completed, the files can be used for repeated data runs. For more detail on the required electronics calibrations, see Section V.

Once the detector is initialized and data are taken, the data can be observed and analyzed with tools developed by the collaboration. Several event viewers have been developed and analytical methods for analysis and event reconstruction

are being employed using packages such as MATLAB, C++, and ROOT.

D. Support systems

Power is supplied by HV and low voltage (LV) units from modular power supplies. These are mounted in the rack underneath the mTC's aluminum enclosure and are operated remotely. The front-end power consumption is approximately 330 W, so cooling of the electronics is crucial. Commercially available hard-drive chiller plates are used and mounted on the electronics card cages with deionized water as a coolant at a total flow rate of ~2 GPM. When the detector is operating, the temperature of the ASICs is monitored during operation to ensure the safety and stability of the electronics and typically is stable in the range ~30–35 °C, depending on ASIC position.

V. CALIBRATION

A number of calibrations are required to operate the detector and remove systematic biases, including electronics effects (both amplitude and timing), MCP efficiency and gain, and calibration with physics processes. We describe each in more detail here.

A. Electronics calibrations

The architecture of the IRS ASIC utilizes individual capacitors and comparators for each of the 32 768 storage cells of an input channel. Variations in the fabrication process create sample-to-sample differences in threshold voltages for the comparators, resulting in a fixed-pattern voltage structure that must be removed from digitized waveforms. These are known as "pedestals" and are evaluated by collecting events with no signal input. This may be done, for example, with the MCP high voltage turned off, or with the high voltage on but using software triggers that are uncorrelated with any signal inputs. Pedestals are typically collected at the beginning of a run period. Over 50×10^6 pedestals are required to run all channels of the detector at their full sampling depth. An example waveform following pedestal subtraction can be seen in Fig. 18.

Further feature extraction is performed on pedestal subtracted data, including estimates of pulse height and pulse times. Pulse timing is estimated using an offline, software-based, constant-fraction discrimination method, with time defined by the crossing of the signal over a set percentage of the pulse height, typically around 50%. Linear interpolation is used to determine this time with much higher granularity than the 370 ps spacing of the individual samples. To achieve the best possible timing resolution, further calibrations are necessary to remove ASIC fabrication effects. The delay line used to generate the fast sampling signals within the ASIC is a current-starved inverter chain. As with the storage cell comparators, process variations cause threshold variations in these inverters, leading to non-uniform timing distributions from sample-to-sample. This manifests as a fixed-pattern timing structure that is unique to each ASIC, which we refer to as the "fine timing

calibration." A total of 128 timing offsets must be calculated for each ASIC, one for each stage of the delay line. Typical spreads in timing values are 10%-15% of the nominal sampling delay. For example, in our standard operating mode with the IRS running at 2.7 GSa/s, the mean timing delay is 370 ps, with a spread of roughly $\sigma_t \approx 13$–55 ps.

To perform this timing calibration, we inject MCP-like pulses into the electronics at known delays relative to the sampling clock. By stepping the delay of these pulses in fine increments (as low as 15 ps), we can calculate a pulse time in units of sampling cells and cross reference it against the known delay, allowing us to map out the fine time structure within the ASIC. A total of 24 576 of these timing values (128 sample delays × 16 ASICs × 12 board stacks) are stored for the full detector.

Following these fine timing calibrations, we must then align all channels of the mTC to a common time reference. Although all channels sample synchronously based on the distributed master clock, various skews are introduced throughout the system from the ASIC structure, PCB routing delays, cable lengths, etc. To characterize these delays, we use a laser system, described below, to inject signal at fixed times into each face of the detector. By aligning timing of pulses for all channels relative to one another, we can measure these skews for each of the 1536 channels and remove them for any subsequent analysis.

B. Laser sources

A precision timing laser system (Advanced Laser Systems EIG1000D) can inject signal into any of the 6 faces of the mTC through a "needle" fiber connector installed in the space between the MCPs. Variable neutral density filters can be inserted between the laser diode output and the input to the fiberoptic connections that inject into the mTC. This allows studies at adjustable light levels, from single PE and up. Stepper motors can be used to move optical elements and select the injection point of the laser or adjust the laser attenuation. The laser controller is triggered by the timing distribution board, allowing optical pulses to be injected at adjustable times relative to the master sampling clock. These features allow the laser to be used as an automated *in situ* calibration or validation source.

C. MCP gains

All 24 MCPs operate on independent high voltage channels, allowing selection of gain independently for each tube. A specification from the manufacturer is provided for each MCP unit with HV settings at 10^5 gain. We have further measured gain curves by observing the common-last-dynode of each MCP for laser and cosmic ray muon signals. Specific HV settings vary by operating mode, as the expected number of PE detected covers a very broad range from tens of thousands of PE (e.g., for cosmic ray muons fully traversing the detector) to under 100 PE (for the delayed neutron capture from IBD).

Since gain can vary considerably across the pixels of an MCP, we must further calibrate the gains of each individual channel. This is typically done by measuring single PE
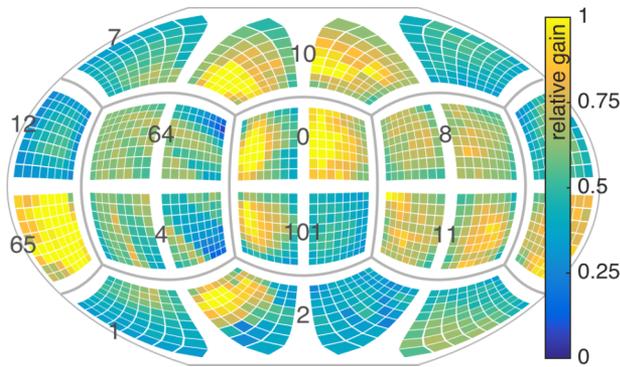
FIG. 21. Display of a preliminary relative gain map obtained from laser data. This includes electronics effects, such as variation in amplifier gain.

signals across the detector, either injected using the laser or by observing MCP dark pulses, and normalizing their mean amplitudes to one another. This further provides a conversion factor from digitized counts to an estimated number of PE for each recorded pulse. Quantum efficiency can be similarly calibrated on a pixel-by-pixel basis, using the calibration laser running in a mode where we collect primarily single PE pulses.

An example of a gain map obtained from laser data is shown in Fig. 21.

D. Cosmic muons

Cosmic ray muons provide one method for validating the calibrated performance of the detector, as they have a known energy deposition and time-stable rate. The fairly stable flux at sea level (~1 Hz through the detector) has a mean energy of about 2 GeV and is peaked near the vertical but falls off towards the horizon gently with a $\cos^2$ of zenith angle. In typical running conditions, one would expect some variation due to the local overburden. The minimum ionizing energy loss rate for polyvinyltoluene[15] (the plastic in the mTC's scintillator) is 1.956 MeV cm^2/g with density 1.02 g/cm^3, so the net (mean) energy loss rate in the cube should be about $dE/dx = 2.3$ MeV/cm.

To acquire muon data, a low gain is set on the MCPs to avoid saturation, and the trigger levels are changed accordingly. The previously determined electronics and gain calibrations are used to analyze the resulting data, and muon tracks can be fitted through the detector, as shown in Fig. 22. An

FIG. 23. Preliminary distributions of reconstructed cosmic ray muon parameters using data collected with the mTC, showing reconstructed incident angle (top left), energy deposition (top right), muon track length within the mTC (bottom left), and muon energy deposition per unit track length (bottom right). These distributions reflect a 500 event data sample. Results are expected to improve as calibrations continue.

example of reconstructed muon parameters for a preliminary data set is shown in Fig. 23.

VI. RECONSTRUCTION

The fast timing of the mTC's electronics, coupled with the excellent spatial resolution of the MCP channels, allows for high-quality reconstruction of subatomic events. Reconstruction is generally subdivided into two categories: unconstrained and constrained. *Unconstrained reconstruction* techniques like simple back-projection make possible the recovery in space and time of any arbitrarily distributed pattern of energy, while *constrained reconstruction* techniques — the simplest being a single point-source fit — allow for the exploitation of *a priori* knowledge about the event and provide more accurate reconstructions as long as correct assumptions are applied.
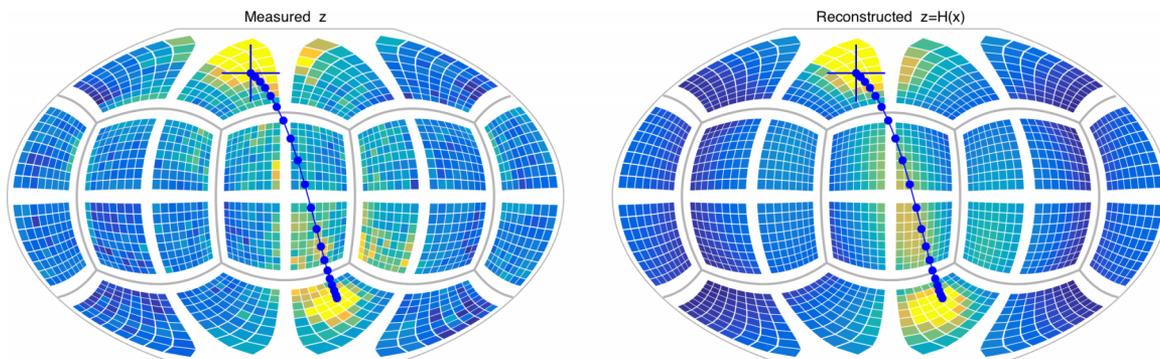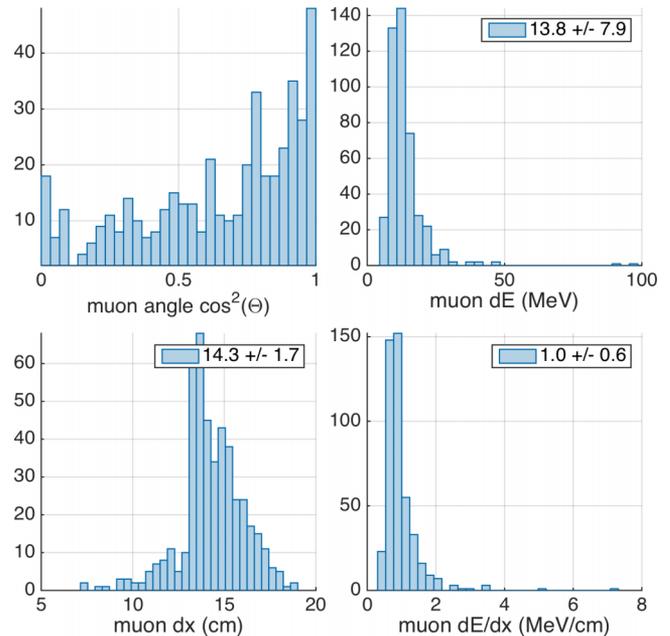

FIG. 22. Event display for a muon measured in mTC (left) and the expected light distribution for the best fit reconstructed path of the muon (right).

In general, the likelihood of observing a single PE z from a single point-source θ is

$$p(z|\theta) = \Lambda_t P_\Omega P_\gamma P_T Q, \qquad (16)$$

where Λ_t is the temporal likelihood, P_Ω is the solid angle probability, P_γ is the un-attenuated energy probability, P_T is the transmission (or non-reflecting) probability, and Q is the PMT quantum efficiency. Equation (16) then forms the basis of our likelihood function, defining the likelihood of point-source θ given measurements z,

$$p(\theta|z) = \prod_j p(z_j|\theta) p(\theta), \qquad (17)$$

where the likelihood $p(z_j|\theta)$ of measurement j with prior $p(\theta)$ is simply an evaluation of the measurement space created by θ at z_j. Equation (17) extends to multiple point sources as well,

$$p(\theta|z) = \prod_j \sum_i w_i p(z_j|\theta_i) p(\theta_i). \qquad (18)$$

For point source i, the likelihood $p(z_j|\theta_i)$ of measurement j given source i with weight w_i and prior $p(\theta_i)$ is simply an evaluation of the measurement space created by θ_i at z_j. This measurement space is defined by a point source position P_θ at time t_θ and is a function of several detector and scintillator characteristics including the following:

- scintillation spectrum, yield and decay constant(s);
- Cherenkov spectrum;
- quenching factors for heavy particles;
- scintillator attenuation length;
- re-emission efficiency of attenuated photons;
- refraction indices of the scintillator and PMT glass;
- PMT QE; and
- time and energy calibrations.

Equation (18) forms the basis for a variety of parameter estimators in the mTC. Any number of complex shapes (i.e., muon tracks, neutron scatters, and a complete antineutrino event) can be built up by using a collection of these simple point sources.

A. Candidate cuts

Measured events in the mTC pass through several candidate cuts before they are considered as possible $\bar{\nu}_e$ candidates. These cuts and their related candidate efficiencies are shown in Fig. 24. The cuts are implemented to both improve the quality of the $\bar{\nu}_e$ events and also to reduce the likelihood of backgrounds entering into the $\bar{\nu}_e$ candidate dataset.

The 5 mm edge cuts reject events with either prompt or delayed vertices <5 mm from the wall. Since the mTC is a single volume detector, the MCPs are directly adjacent to the scintillation volume. Points which are fit too close to the wall tend to suffer from low reconstruction quality, and this cut serves to reject these potentially unreliable fits. Another reason for the edge cut is to reduce the likelihood of a positron from leaving the detector volume, which could result in severe underestimation of the $\bar{\nu}_e$ energy. This edge cut reduces the detector fiducial volume by 20%, from 2.2 to 1.7 l.
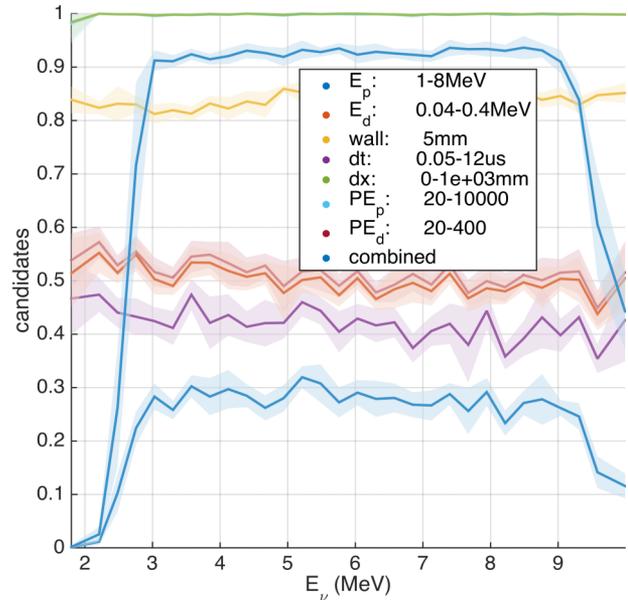

FIG. 24. Simulated mTC $\bar{\nu}_e$ candidate efficiency vs. $\bar{\nu}_e$ energy.

There are time cuts on the prompt-delayed dt as well; these are a 50 ns floor and a 12 μs ceiling (hardware imposed). The floor is designed to prevent late prompt PE's from entering the delayed signal dataset.

We impose energy and PE cuts as well. For the prompt signal, we accept energies of 1-8 MeV and PE counts between 20 and 10 000 PEs. The delayed signal has much stricter energy cuts, as it has a more consistent energy output; we accept delayed candidates with between 20-400 PEs and 40-400 keV.

$\bar{\nu}_e$ candidates must meet all these requirements in order to be accepted into the $\bar{\nu}_e$ candidate pool. In the mTC, we find about 30% $\bar{\nu}_e$ candidate efficiency at 3-4 MeV. The dominant source of efficiency loss is neutrons leaving the detector volume, which happens 45% of the time, and from neutrons leaving the 12 μs time window, which happens 30% of the time. These two causes alone reduce the mTC $\bar{\nu}_e$ candidate efficiency to <40%; the other cuts only have minor effect.

B. Performance

While the mTC has not yet detected any real world antineutrinos, its performance has been modeled through many GEANT and MATLAB MC simulations. Figure 25 shows the expected $\bar{\nu}_e$ energy resolution of the mTC across the 2-10 MeV reactor $\bar{\nu}_e$ energy spectrum, which peaks at 3-4 MeV. Our mean energy resolution is about 11% 1σ, including outliers in the long tail, or as low as 5% if outliers are ignored. Most outliers are due to higher energy positrons leaving the detector, resulting in significant under-estimation of their true energy. The prevalence of these occurrences decreases as the wall cuts are expanded.

Figure 26 shows the corresponding energy resolution as a function of $\bar{\nu}_e$ energy rather than weighted by the reactor spectrum as in Fig. 25. A nice coincidence is seen here: the best energy resolution is enjoyed at the peak of the reactor spectrum, with the resolution suffering at lower energies due to
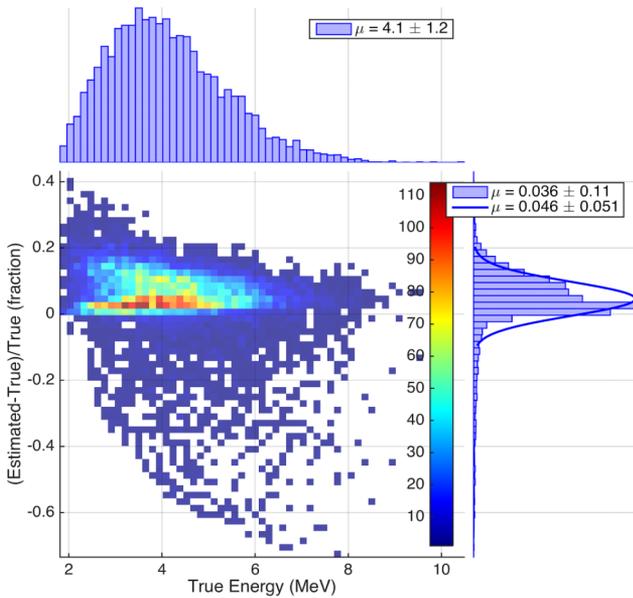
FIG. 25. Monte Carlo (MC) simulation results showing $\bar{\nu}_e$ energy resolution in mTC (y axis) over the reactor antineutrino spectrum (x axis).
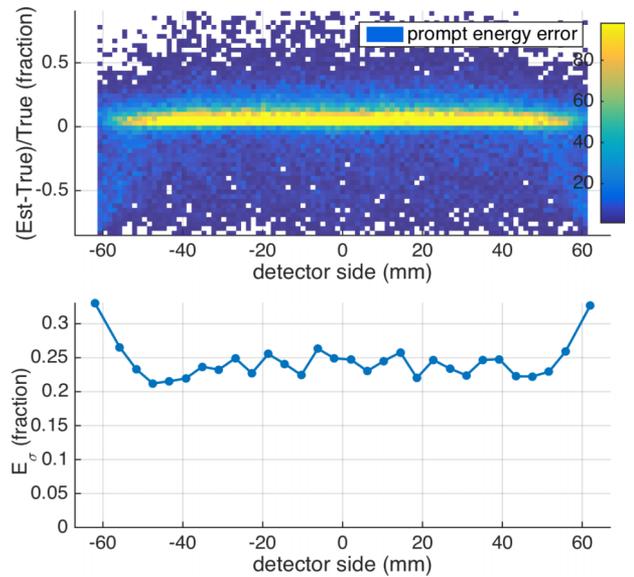

FIG. 27. Simulated mTC $\bar{\nu}_e$ energy resolution vs. vertex location, averaged over $\bar{\nu}_e$ energy with a flat input spectrum. The detector center is at 0 mm, and the detector wall is at 67 mm.

lack of light, and suffering at higher energies due to the high energy positron tracks leaving the detector more frequently. Figure 27 shows the same energy resolution values vs. vertex within the detector, indicating that resolution suffers near the detector walls.

Figure 28 shows the prompt and delayed vertex resolution as a function of $\bar{\nu}_e$ energy. In this context, "vertex" means the $\bar{\nu}_e$ interaction point for the prompt signal (the start of the e^+ track) and the capture location of the neutron for the delayed signal. The prompt vertex fits tend to bias towards the center of the e^+ track rather than its start, and both the prompt and delayed vertex location fits are smeared by the spatially dispersed energy depositions of the prompt (2×511 keV) and delayed (1×470 keV) gammas.

Figure 28 also shows that the prompt vertex resolution suffers at low $\bar{\nu}_e$ energies due to lack of light, and at higher $\bar{\nu}_e$ energies due to longer e^+ tracks (as the center of the track distances itself from its start point). Figure 29 shows Monte Carlo $\bar{\nu}_e$ angle reconstructions in the mTC and puts the mTC angular resolution in context by comparison with the CHOOZ detectors[16] and hypothetical 138 kT TREND detector.[35]

Figure 29 shows the angle error distributions of mTC, CHOOZ, and TREND over the −1 to 1 $\cos(\theta)$ range, where $\cos(\theta) = 1$ corresponds to zero error and $\cos(\theta) = -1$ corresponds to 180° of angle error. The color maps on the unit sphere represent several thousand reconstructions of Monte Carlo events and serve as a more intuitive measure of how well the mTC reconstructs direction. Though the mTC hypothetically

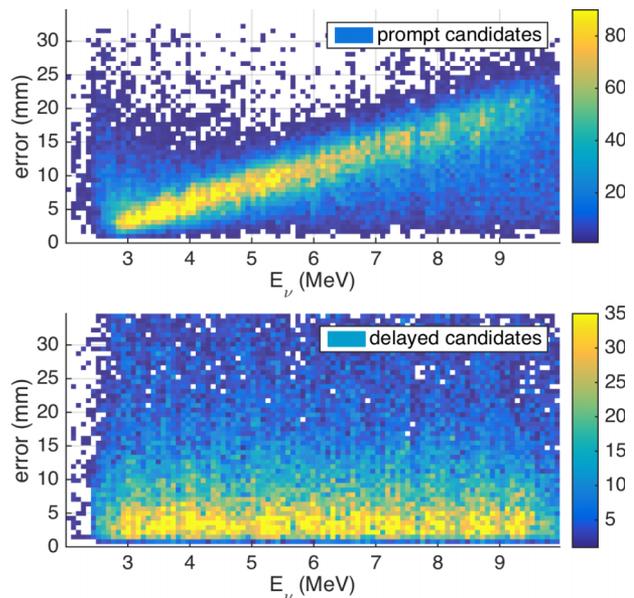FIG. 26. Simulated mTC $\bar{\nu}_e$ energy resolution vs. $\bar{\nu}_e$ energy.

FIG. 28. Simulated mTC $\bar{\nu}_e$ prompt and delayed vertex resolution vs. $\bar{\nu}_e$ energy.
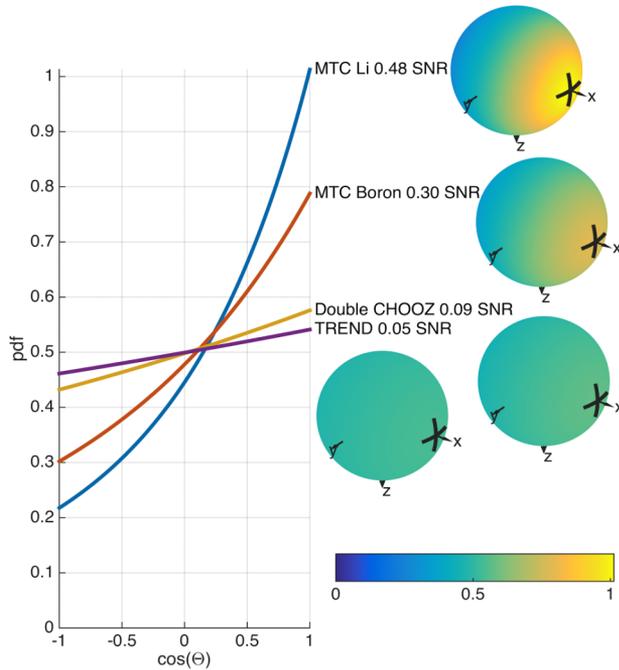
FIG. 29. Angular $\bar{\nu}_e$ resolution comparison between mTC, CHOOZ, and TREND. CHOOZ and double CHOOZ detectors both provide nearly identical angular resolutions as a result of identical near and far detector designs.

exceeds the $\bar{\nu}_e$ angular resolution of the CHOOZ detector, they are both in reality extremely poor at directional determination from a single $\bar{\nu}_e$ and require great statistics to reduce the uncertainty on any angle fit.

The angular resolution metric we employ is the vector Signal to Noise Ratio (vector SNR). The vector SNR is the magnitude of the vector mean divided by the standard deviation in any of the 3 dimensions (which should all share similar uncertainties) for a given population of vectors. In our application, these vectors are the reconstruction vectors connecting the delayed signal vertices to the prompt signal vertices. Such a group of vectors should, with some uncertainty, point back towards the $\bar{\nu}_e$ source.

We use this metric rather than the more common angle 1σ because the uncertainty is so great as to wrap completely around the sphere, rendering simpler 1-dimensional methods meaningless. An alternative metric for directional statistics is the von Mises–Fisher distribution, which provides a "concentration parameter" that increases as the angular distribution decreases.

In the mTC, our mean reconstruction vector (from delayed vertex to prompt vertex) is 10 mm long, and the 1σ standard deviation of these vectors is 32 mm, giving us a vector SNR of 10 mm/32 mm = 0.3. In the CHOOZ $\bar{\nu}_e$ detector, the mean reconstruction vector is 17 mm long with a 190 mm 1σ uncertainty about each axis, producing an SNR of 17 mm/190 mm = 0.09. The simulated TREND SNR is 0.05. Figure 29 shows the angular distributions for these 3 detectors plotted over the -1 to 1 $\cos(\theta)$ range, as well as wrapped around a unit sphere on a common colormap. Also, a hypothetical mTC-detector is shown with 1.5% ^{6}Li-loaded plastic scintillator.

As per Fig. 29, the chances of reconstructing a $\bar{\nu}_e$ as originating from the correct hemisphere (i.e., forward or backward)

are 62% in the mTC, 54% in CHOOZ, and 52% in TREND. These values are obtained by simply integrating the 0-1 $\cos(\theta)$ probabilities in Figure 29.

The analytical equation defining the vector SNR distributions shown in Fig. 29 is Equation (44) in Jocher *et al.*[35] This equation defines a proper analytical Probability Density Function (PDF) (one that normalizes to unity) over the -1 to 1 $\cos(\theta)$ range and was derived specifically for the purpose of describing $\bar{\nu}_e$ directional resolution. We could not find evidence of its use previously in the field of directional statistics.

VII. CONCLUSIONS

The compact size of the miniTimeCube gives it the potential for many novel measurements. Preliminary results indicate that the mTC should attain a 1:1 $\bar{\nu}_e$ signal-to-noise ratio inside the shielding cave at the NIST reactor site. The mTC is designed to exploit fast timing for event reconstruction. While the fiducial volume may be smaller than desirable in certain aspects, the sizing coupled with the fast $\sim$100 ps electronics and high spatial resolution enables high levels of imaging and reconstruction not attainable by larger detectors. Our modeled antineutrino vertex and angular resolutions (10 mm vector and 32 mm of noise) indicate that the mTC should enjoy $\sim$3 times better directional resolution than existing large $\bar{\nu}_e$ detectors.

mTC energy resolution, at 11% (or 5% without outliers), is on par with other large antineutrino detectors, despite the challenges faced by uncertain gamma energy deposition from one IBD event to the next. The efficiency of the detector, at 30%, may be improved significantly by the addition of higher levels of neutron capture doping material, which would solve many of the current problems with neutron retention in time and space.

One could scale the mTC concept up to a larger detector — as is planned with NuLat — or build a networked array of such small detectors. These could be used to perform in-depth studies ranging from neutrino oscillation with novel detector arrangements at very short baselines or to explore nuclear security applications. Upgrades to the mTC are ongoing, with continuing calibration, electronics improvements, and reactor tests planned. Ultimately we believe that the mTC provides exciting opportunities for fast timing exploitation, and we look forward to publishing future results as they become available.

[1]P. Vogel, L. Wen, and C. Zhang, Nat. Commun. **6**, 6935 (2015); e-print arXiv:1503.01059 [hep-ex].

[2]J. G. Learned, KamLAND internal memo, 2009.

[3]M. B. Chadwick *et al.*, "ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data," Nuclear Data Sheets **112**(12), 2887–2996 (2011).

[4]J. G. Learned, e-print arXiv:0902.4009 [hep-ex] (2009).

[5]C. Lane *et al.*, e-print arXiv:1501.06935 [physics.ins-det] (2015).

[6]ELJEN, Scintillator EJ-254 data sheet, 2015.

[7]PHOTONIS, PLANACON XP85012 data sheet, 2013.

[8]C. L. Cowan *et al.*, Science **124**, 103 (1956).

[9]F. Reines and C. L. Cowan, Nature **178**, 446 (1956).

[10]G. Knoll, *Radiation Detection and Measurement* (Wiley, 2000).

[11]I. Pawelczak *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **751**, 62 (2014).

[12]A. Wright, P. Mosteiro, B. Loer, and F. Calaprice, Nucl. Instrum. Methods Phys. Res., Sect. A **644**, 18 (2011); e-print arXiv:1010.3609 [nucl-ex].

[13]T. Crane and M. Baker, in *Passive Nondestructive Assay of Nuclear Materials*, edited by D. Reilly, N. Ensslin, and H. Smith (Los Alamos National Laboratory, Los Alamos, 1991).

[14]ENDF, Nuclear Data Libraries, 2011.

[15]K. Olive *et al.* (Particle Data Group), Chin. Phys. C **38**, 090001 (2014).

[16]Y. Abe *et al.* (Double Chooz Collaboration), J. High Energy Phys. **1410**, 086 (2014); e-print arXiv:1406.7763 [hep-ex].

[17]R. de Meijer *et al.*, in *Neutrino Geophysics: Proceedings of Neutrino Sciences 2005*, edited by S. Dye (Springer, 2005); e-print arXiv:physics/0607049.

[18]L. Cheng *et al.*, Technical Report BNL-NIST-0803, Rev. 1, 2004.

[19]A. Hanson and D. Diamond, Technical Report BNL-96386-2011-IR, 2011.

[20]P. Vogel and J. F. Beacom, Phys. Rev. D **60**, 053003 (1999); e-print arXiv:hep-ph/9903554 [hep-ph].

[21]V. Kopeikin, L. Mikaelyan, and V. Sinev, Phys. At. Nucl. **67**, 1892 (2004); e-print arXiv:hep-ph/0410100 [hep-ph].

[22]D. Dwyer and T. Langford, Phys. Rev. Lett. **114**, 012502 (2015); e-print arXiv:1407.1281 [nucl-ex].

[23]T. Mueller *et al.*, Phys. Rev. C **83**, 054615 (2011); e-print arXiv:1101.2663 [hep-ex].

[24]S. Agostinelli *et al.* (GEANT4), Nucl. Instrum. Methods Phys. Res., Sect. A **506**, 250 (2003).

[25]J. Ashenfelter *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **806**, 401 (2016); e-print arXiv:1506.03547.

[26]C. Hagmann, D. Lange, J. Verbeke, and D. Wright, Technical Report UCRL-TM-229453, 2012.

[27]M. S. Gordon *et al.*, IEEE Trans. Nucl. Sci. **51**, 3427 (2004).

[28]The Ice Radio Sampler (IRS) is a descendant of the Buffered LABRADOR series of ASICs[36] and was originally designed for performing radio searches for ultra-high energy neutrinos using antennas embedded in Antarctic ice.

[29]B. Dey *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **775**, 112 (2014); e-print arXiv:1410.0075 [physics.ins-det].

[30]K. Bechtol *et al.*, Astropart. Phys. **36**, 156 (2012); e-print arXiv:1105.1832 [astro-ph.IM].

[31]P. Allison *et al.*, Astropart. Phys. **35**, 457 (2012); e-print arXiv:1105.2854 [astro-ph.IM].

[32]E. Delagnes *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **567**, 21 (2006).

[33]S. Ritt, R. Dinapoli, and U. Hartmann, Nucl. Instrum. Methods Phys. Res., Sect. A **623**, 486 (2010).

[34]E. Oberla *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **735**, 452 (2014); e-print arXiv:1309.4397 [physics.ins-det].

[35]G. R. Jocher *et al.*, Phys. Rep. **527**, 131 (2013); e-print arXiv:1307.2832 [physics.ins-det].

[36]L. L. Ruckman, G. S. Varner, and A. Wong, Nucl. Instrum. Methods Phys. Res., Sect. A **591**, 534 (2008); e-print arXiv:0802.2278 [physics.ins-det].