Sensitivity to oscillation with a sterile fourth generation neutrino from ultralow threshold neutrino-nucleus coherent scattering

Bhaskar Dutta, Yu Gao, Andrew Kubik, Rupak Mahapatra, Nader Mirabolfathi, Louis E. Strigari, and Joel W. Walker

Sensitivity to oscillation with a sterile fourth generation neutrino from ultra-low threshold neutrino-nucleus coherent scattering

Bhaskar Dutta[1], Yu Gao[1], Andrew Kubik[1], Rupak Mahapatra[1],
Nader Mirabolfathi[1], Louis E. Strigari[1], and Joel W. Walker[2]

[1] *Mitchell Institute for Fundamental Physics and Astronomy,*
Department of Physics and Astronomy, Texas A&M University, College Station, TX 77843-4242, USA and
[2] *Department of Physics, Sam Houston State University, Huntsville, TX 77341, USA*

We discuss prospects for probing short-range sterile neutrino oscillation using neutrino-nucleus coherent scattering with ultra-low energy ($\sim 10\,\text{eV}$ - $100\,\text{eV}$) recoil threshold cryogenic Ge detectors. The analysis is performed in the context of a specific and contemporary reactor-based experimental proposal, developed in cooperation with the Nuclear Science Center at Texas A&M University, and references developing technology based upon economical and scalable detector arrays. The baseline of the experiment is substantially shorter than existing measurements, as near as about 2 meters from the reactor core, and is moreover variable, extending continuously up to a range of about 10 meters. This proximity and variety combine to provide extraordinary sensitivity to a wide spectrum of oscillation scales, while facilitating the tidy cancellation of leading systematic uncertainties in the reactor source and environment. With 100 eV sensitivity, for exposures on the order of 200 kg·y, we project an estimated sensitivity to first/fourth neutrino oscillation with a mass gap $\Delta m^2 \sim 1\,\text{eV}^2$ at an amplitude $\sin^2 2\theta \sim 10^{-1}$, or $\Delta m^2 \sim 0.2\,\text{eV}^2$ at unit amplitude. Larger exposures, around 5,000 kg·y, together with 10 eV sensitivity are capable of probing more than an additional order of magnitude in amplitude.

I. INTRODUCTION

Several recent short baseline neutrino experiments hint at the presence of additional neutrinos beyond the three active components in the Standard Model (SM). The radioactive source experiments of the GALLEX [1] and SAGE [2] Solar neutrino detectors have found indications of a deficit of electron neutrinos [3, 4]. Independently, very short baseline neutrino experiments with distances of < 100 m find evidence for a deficit of electron anti-neutrinos [5]. At least one additional sterile neutrino with a mass splitting of $\Delta m_{14}^2 \sim 1\,\text{eV}^2$ and an approximate 10% admixture between the electron flavor neutrino and the new sterile neutrino can potentially accommodate both of these results. More sophisticated proposals include the introduction of a hidden twin three-generation neutrino sector [6], which would hypothetically feature the testable short-baseline appearance of SM neutrinos from decays in flight. In addition, anomalies in data collected by the LSND and MiniBooNE [7–9] collaborations may likewise be explained by the invocation of sterile neutrinos.

Theoretical models have been developed to incorporate neutrinos at the ~ 1 eV mass scale. The strongest constraints on these models come from cosmology, in particular the constraints on the number of additional light degrees of freedom from Big Bang Nucleosynthesis (BBN) [10] and Planck measurements [11]. Though these constraints severely limit the theoretical models that have been considered, they can be evaded by introducing a coupling between a vector boson with $\sim$ MeV mass scale to the sterile neutrinos [12–14]. In addition to cosmology, Solar neutrinos can test for the presence of a 4th generation sterile neutrino [15].

Dedicated experiments are under development to confirm or rule out the existence of an additional sterile neutrino. These experiments are designed to detect electron anti-neutrinos on a baseline $\sim 1 - 20$ m, which is characteristic of the hypothetical sterile neutrino oscillation length. Ref. [16] discusses the prospects for using scintillators to detect electron anti-neutrinos via inverse beta decay in close proximity to a reactor core, while Refs. [17–19] discuss the prospects for using a radioactive source in close proximity to a scintillator. Ref. [20] first suggested the possibility of probing sterile neutrino oscillations through coherent elastic neutrino-nucleus scattering (CEνNS). Ref. [21] has recently updated this notion, in the specific context of low-energy pion and muon decay-at-rest neutrino sources.

In this paper we discuss the prospects for using the CEνNS channel and a new experimental design to search for sterile neutrinos. In spite of the large predicted cross section, CEνNS has yet to be detected, primarily because the state of the art of detector technology to this date has been unable to deliver sufficiently low threshold sensitivity to register deposition of the kinetic energy of the heavy recoiling nucleus. In particular, experimental programs that have previously discussed the prospects for detecting CEνNS from nuclear reactors [22] have typically referenced nuclear recoil thresholds at the keV scale or greater.

We project experimental sensitivities referencing the application of ultra-low threshold (as low as 10 eV) Ge and/or Si detectors [23] to the measurement of CEνNS. We consider first the more conservative recoil energy threshold of 100 eV for near-term phase-1 experiments. There are many detector technologies that are rapidly approaching the 100 eV energy threshold in the current generation of dark matter search experiments [24–26], and we project a target of this order to be feasible for first generation data taking, as further justified in Section III. For high-mass long-exposure runs, such as a ton-scale 5-year follow up experiment, we restrict analysis to the more speculative 10 eV energy threshold scenario, which will require substantial further research and development toward the goal of single electron resolution phonon-mediated ionization detectors [23]. Our first-generation experimental motivation is, however, direct and imminent, referencing the incremental evolution of in-hand technology based upon economical and modularly scalable detector arrays in conjunction with a research reactor site capable of providing large flux rates and enabling distances-from-core as near as one meter [1], or between two and three meters after allowing for the additional depth of shielding materials suggested by detailed simulation. Specifically, the proposed anti-neutrino source is a megawatt-class TRIGA-type pool reactor stocked with low-enriched ($\sim 20\%$) ^{235}U, which is administrated by the Nuclear Science Center at Texas A&M University (TAMU). This adjacency geometrically enhances the neutrino flux to within about an order of that typifying experiments at a 30 m baseline from a gigawatt-class power reactor source. Additionally, it provides an experimental length scale and median energy that are naturally calibrated for probing some of the most interesting regions of the Δm_{14}^2 parameter space. Moreover, the core of this reactor design is mobile, readily facilitating the collection of data at several propagation lengths, which improves sensitivity to the oscillatory spatial character of the signal and induces a cancellation of leading systematic error components. Additional background on the TAMU research nuclear reactor facility (including isotopic fuel ratios and the computation of the expected anti-neutrino flux), the fabrication and performance of ultra-low threshold Ge and Si nuclear recoil detectors, and the physics applications of neutrino-atom scattering are provided in Refs. [23, 27], as well as the references therein. Note, in particular, that most research reactors are not continuously operated, although there is no technical reason which prevents this, resources permitting; when referencing time, we refer to reactor live-time.

This paper is organized as follows. In Section II, we predict the CEνNS rate for our planned detector setup, including the possible presence of an additional sterile neutrino. In Section III, we characterize detector technology, experimental backgrounds, shielding requirements, and sources of systematic error. In Section IV, we establish a logarithmic observation baseline schedule with exposure compensation for the geometric flux dilution with distance. In Section V, we estimate statistical sensitivity of the proposed experiment to Δm_{14}^2 and $\sin^2 2\theta_{14}$ for various experimental baselines L. In Section VI, we establish a procedure in semi-closed form for quantifying the statistical preference of an oscillation hypothesis in the presence of data, and demonstrate the cancellation of leading systematic uncertainties. In Section VII, we describe a higher-order statistical method, which goes beyond certain approximations and simplifications made in the prior description. In Section VIII, we summarize and present our conclusions.

II. SCATTERING RATE PREDICTIONS INCLUDING STERILE NEUTRINOS

In this section we theoretically establish the CEνNS rate at the envisioned detector setup. We predict both the SM rate and the rate in the presence of a hypothetical additional sterile neutrino.

The coherent elastic nuclear scattering of neutrinos (with sufficiently low energy, typically $\sim$ MeV) is a long-standing prediction of the Standard Model that has been theoretically well-studied [28]. The differential cross-section for SM scattering of a neutrino with energy E_ν from a target particle of mass M and kinetic recoil E_R is, in terms of the applicable vector $q_V \equiv q_L + q_R$ and axial $q_A \equiv q_L - q_R$ charges,

$$\frac{d\sigma}{dE_R} = \frac{G_F^2 M}{2\pi}\left[(q_V + q_A)^2 + (q_V - q_A)^2\left(1 - \frac{E_R}{E_\nu}\right)^2 - (q_V^2 - q_A^2)\frac{ME_R}{E_\nu^2}\right]. \tag{1}$$

For anti-neutrino scattering there is a relative negative phase between (q_V, q_A) associated with the parity-flip, which is accommodated by the prescription $(q_V, q_A) \equiv (T_3 - 2Q\sin^2\theta_W, -T_3)$. For coherent nuclear scattering, these terms should be summed over the quark content of protons and neutrons, and either multiplied by the respective counts (Z, N) of each (in the vector case) or multiplied by the respective differential counts $(Z^+ - Z^-, N^+ - N^-)$ of up and

---

[1] Very near approach of the mobile reactor core (the active region of which is an approximately 0.35 m cube with several additional centimeters of housing material at each side boundary) to the thermal column identified for housing the planned experiment is presently obstructed by a ~ 0.7 m graphite block, the shielding benefits of which may be sub-optimal to the current application. Pending detachment of this block, it would become physically possible to locate the center of the reactor (closer than) one meter from the center of a Germanium or Silicon detector array.

down spins (in the axial case) [29]. This sum over nuclear constituents at the coupling level, prior to squaring in the amplitude, is the essence of the nuclear coherency boost.

In order to compute the cumulative expected Standard Model anti-neutrino scattering rate (cf. Ref. [29]), it is necessary to integrate in the region of the E_ν vs. E_R plane that is above $E_\nu > E_\nu^{\min} \equiv (E_R + \sqrt{2ME_R + E_R^2})/2$, which is the minimal neutrino energy (i.e. the inversion of the expression for the maximum recoil $E_R^{\max} \equiv 2E_\nu^2/(M + 2E_\nu)$ achievable in a collision with no glancing component) required to trigger a given recoil, and between the i^{th} binned boundary pair $E_R^{i\downarrow} < E_R < E_R^{i\uparrow}$ of the detector recoil. The integrand is a product of the previously described differential cross section and the normalized anti-neutrino energy spectral distribution $\phi(E_\nu) \equiv dN_\nu/dE_\nu \div N_\nu$, as well as the incident anti-neutrino flux Φ, the detector mass M, and the exposure time T. The expected number $N_{\text{Exp}}^{i,n}$ of CEνNS scattering events in a detector of composite mass M_{Det} from nuclei with mass M in a recoil energy bin E_R^i for an exposure T_n at a distance L_n is then

$$N_{\text{Exp}}^{i,n} = \Phi_0 \times T_n \times \frac{L_0^2}{L_n^2} \times \frac{M_{\text{Det}}}{M} \times \int_{E_\nu^{\min}(E_R^{i\downarrow})}^{\infty} dE_\nu \, \phi(E_\nu) \int_{E_R^{i\downarrow}}^{\min\{E_R^{i\uparrow}, E_R^{\max}(E_\nu)\}} dE_R \, \frac{d\sigma}{dE_R}(E_\nu, E_R) \,. \qquad (2)$$

Nuclear scattering will generally be dominated by the vector charge, and in the limit of vanishing axial charge the residual functional dependence $1 - ME_R/2E_\nu^2$ interpolates between a large cross-section at zero recoil and a vanishing cross-section at kinematic cut-off. The large mass-denominator in $E_R^{\max}$ highlights the necessity of ultra-low threshold detectors for observation of the heavily boosted CEνNS feature. Specifically, a typical MeV-scale incident neutrino will be downgraded by almost 5 orders of magnitude in the recoil (in the most kinematically favorable case) with a 100 GeV order nucleus, to around 20 eV. In more detail, we calculate that, in order to capture about half of the scattering from fission neutrinos with a mean energy of 1.5 MeV, a detector threshold around 50, 20 eV is required in Si, Ge. In this energy regime CEνNS scattering dominates over electron cloud scattering, by a factor of about 500 at 10 eV, and about 30 at 500 eV; the expected scattering rates become comparable above about 1 keV, where the available kinematic recoil of a nuclear mass from MeV neutrino scattering approaches kinematic freeze-out. For recoil thresholds of 10 or 100 eV, the minimal neutrino energy required to trigger a detection in Ge is around 0.6 or 1.9 MeV. This is potentially of note when considering the impact of additional low-energy neutrinos from secondary β-processes induced by reactor neutrons and gammas.

For concreteness, we have adopted the anti-neutrino energy spectral distribution $\phi(E_\nu)$ of Ref. [30] to model reactor neutrinos with energies E_ν above 2 MeV. This spectrum has been extrapolated from the correlated observation of electrons emitted by the fission products of ^{235}U irradiated with thermal neutrons. The flux is down by about a factor of 20 at 5 MeV, and around six magnitude orders at 10 MeV. Below 2 MeV (the threshold for inverse β-decay $\bar{\nu}_e + p \rightarrow e^+ + n$ is $E_\nu > 1.8$ MeV) there is no experimental data, and we employ the theoretically established curve of Ref. [31]. For reference, the intrinsic anti-neutrino production rate of the TAMU research reactor is approximately 1.9×10^{17} per second, yielding a flux at a mean distance-from-core L_0 corresponding to 1 m of $\Phi_0 \simeq 1.5 \times 10^{12} \, [\text{s} \cdot \text{cm}^2]^{-1}$ (dividing by $4\pi \times 100^2 \sim 10^5$).

All plotted and tabulated numerical calculations in this document are performed in the context of a Germanium target, averaging over natural isotopic abundances, where the leading contribution (at just over one third) is from ^{74}Ge, followed by ^{72}Ge and ^{70}Ge (at around one quarter and one fifth respectively), as well as ^{73}Ge and ^{76}Ge (at less than one tenth each). Typical values of $N_{\text{Exp}}^{i,n}$ for Ge in the proposed experimental configuration are provided in TABLE I, normalized to an exposure of one year with a one kilogram detector at one meter from core. Keeping in mind that the practical separation of detector and reactor centers may typically be somewhat larger, it is convenient here to provide baseline signal rates normalized to one meter for numerical simplicity of the rescaling law from unit length to other experimental baselines. This is of particular relevance for an oscillation experiment, wherein it is advantageous to probe several different length scales. The energy binning adopted is tuned here to achieve an approximately equitable distribution of scattering events. It is consistent with the expected detector recoil energy resolution of approximately 10%, down to a demonstrated (projected) absolute RMS resolution of around 7 (2) eV$_{ee}$ [23]. Moreover, it produces a manageable collection of curves that are adequately spaced so as to be structurally distinguishable and statistically well-populated, yet sufficiently finely grained so as to preserve features of the underlying continuum. Note, however, that the specific bin boundaries indicated in TABLE I are purely illustrative, insomuch as any practical binning scheme will ultimately be very sensitive to the future precision calibration of low-energy nuclear responses, as described further in Section III. In the case of a recoil sensitivity threshold substantially above 10 eV, the inapplicable bins are discarded.

We now consider the modification of these rates in the presence of a fourth sterile neutrino flavor, with a mass gap from the SM triplet on the order of $\Delta m_{14}^2 \simeq 1$ eV^2. Our targeted experimental baseline is on the order of a few to several meters (starting from the minimal length consistent with shielding requirements, e.g. 2-3 m, and extending to the vicinity of 10 m), such that SM oscillation to mu and/or tau flavors (which occur over much longer distances) is entirely decoupled. In this case, the probability $P_{(\alpha \rightarrow \beta)}$ for oscillation between the two decoupled neutrino flavors

TABLE I: The SM expected count of CEνNS scattering events N_{Exp}^i per nuclear kinetic recoil energy E_{R}^i bin i in 1 kg of Ge at a mean distance from core of 1 m for an integrated exposure of 1 [y]. The event rate drops precipitously above 1 keV. The presented binning is merely suggestive, and will ultimately be strongly dependent upon precise calibration of the nuclear response function at low energies.

E_{R}^i eV	10-20	20-35	35-55	55-80	80-110	110-155	155-220	220-350	350-1000
N_{Exp}^i	853	961	962	911	818	874	820	862	782

(α, β) is

$$P_{(\alpha \to \beta)} = \sin^2 \left[2\theta \right] \times \sin^2 \left[\frac{\Delta m^2 L}{4 E_\nu} \right]. \tag{3}$$

The wavelength associated with this propagation may be expressed as

$$\lambda = 4.97 \, [\text{m}] \times \left\{ \frac{E_\nu}{1 \, [\text{MeV}]} \right\} \times \left\{ \frac{1 \, \text{eV}^2}{\Delta m^2} \right\}. \tag{4}$$


FIG. 1: Fractional deviation from expected SM CEνNS event rates in Ge due to electron anti-neutrino oscillation with a sterile fourth flavor. Separate basis curves γ^i are displayed for nine binned windows in the nuclear kinetic recoil energy deposition. A tenth (bold, dashed) curve demonstrates the cumulative unbinned event deviation over all recoils above 10 eV. For a given energy band, values approaching 1 indicate more full depletion relative to the SM, whereas values nearer to zero indicate more marginal depletion. The vertical axis is inversely rescaled by the amplitude factor $\sin^2 2\theta_{14}$, and predictions for any targeted amplitude may be immediately read off by replacing the upper limit (1.0) with the applicable value. The horizontal axis indicates the product of the distance from core in meter units and the mass-square difference Δm_{14}^2 in units of eV^2. For example, at a gap of 1 eV^2, this axis is read literally in meters, whereas the outer scale bound at 100 would correspond instead to 10 m for $\Delta m_{14}^2 = 10$ eV^2. The presented binning is merely suggestive, and will ultimately be strongly dependent upon precise calibration of the nuclear response function at low energies.

It is useful to define a dimensionless quantity γ_i that represents depletion in the observed CEνNS scattering rate N_{Osc}^i relative to the SM expectation N_{Exp}^i due to oscillation with a sterile fourth generation neutrino (for some binning index i in the energy width and/or detector location).

$$\gamma_i(\Delta m_{14}^2 L) \equiv \frac{1 - (N_{\text{Osc}}^i / N_{\text{Exp}}^i)}{\sin^2 2\theta_{14}}. \tag{5}$$

We may then explicitly establish the convolved oscillation shape functionals $\gamma_i(\Delta m_{14}^2 L)$ defined in Eq. (5).

$$\gamma_i(\Delta m_{14}^2 L) = \left\langle \sin^2\left[\frac{\Delta m_{14}^2 L}{4E_\nu}\right]\right\rangle_{E_\nu} \equiv \iint dE_\nu \, d\sigma \, \phi \times \sin^2\left[\frac{\Delta m_{14}^2 L}{4E_\nu}\right] \div \iint dE_\nu \, d\sigma \, \phi \qquad (6)$$

Integration bounds and function dependencies have been suppressed in Eq. (6), but are identical to those expressed explicitly in Eq. (2). In particular, $d\sigma \equiv dE_R \frac{d\sigma}{dE_R}(E_\nu, E_R)$ represents a differential in the cross-section to be integrated over recoil energy, and $\phi \equiv \phi(E_\nu)$ represents the reactor anti-neutrino source spectrum. There is no residual dependence in the γ_i upon the exposure time, source flux normalization, or detector mass. Once computed (for example, by sampling at fixed increments in the product $\Delta m^2 L$, and interpolating to a continuous numerical function), The dimensionless $\gamma_i(\Delta m_{14}^2 L)$ wholly encapsulate all theoretical aspects of a putative oscillation signal's length and mass scale dependence, including implicit dependence upon the anti-neutrino source, and the physics of the recoil detection mechanism.

FIG. (1) plots γ_i, employing the TABLE I binning, after convolution of a suitable reactor source spectrum with the CEνNS detector response in Ge, as a continuous function of the product of the mass gap Δm_{14}^2 and the distance from core L. The coefficients L (horizontal axis) and $1/\sin^2 2\theta_{14}$ (vertical axis) are adopted in order to make the intrinsic scale invariance of the elementary event profiles manifest. As such, the depicted curves are not associated with any particular benchmark point in the model space, but are instead directly applicable to the full range of targeted parameters by simple renormalization of the axes with respect to the amplitude and phase of Eq. (3).

At very small values of $\Delta m_{14}^2 L$, i.e. at lengths L much smaller than the relevant Eq. (4) wavelength λ, the oscillation fraction will be inadequate for observation. Conversely, at excessively large values of $\Delta m_{14}^2 L$, where the propagation length represents several half-wavelengths, dispersion will wash out all discernible features, ($\gamma_i \Rightarrow 1/2$). The recoil energy binning is observed to partially mitigate the expected wavelength dispersion at both long and short distance scales, as is evident in comparison with the unbinned integration of all recoils. By resolving individual binned response curves, the window of sensitivity to oscillation in the mass gap m_{14}^2 parameter is widened by almost an order of magnitude at fixed L. By physically modulating the experimental baseline L (as facilitated by ready mobility of the reactor core at the proposed experimental site), this window may be extended even further.

III. CHARACTERIZATION OF DETECTORS, BACKGROUNDS AND EXPERIMENTAL UNCERTAINTIES

The march toward detector technology sensitive to nuclear recoil as soft as 100 eV is proceeding rapidly [24–26]. In particular, electron recoil thresholds of 56 eV$_{ee}$ [23], corresponding under the assumption of Lindhard scaling [32] to a nuclear recoil threshold as low as 250 eV, have already been actively demonstrated by CDMSlite at Soudan [24]. Continuous research and development effort is being exerted toward advance on this front, with fabrication and testing of improved design prototypes currently at various stages of progress. TABLE II summarizes vital characteristics of the first-generation detectors envisioned for deployment in the described experiment, including recoil threshold, net mass, and energy resolution. Specifically, an RMS resolution of 7 eV$_{ee}$ has now been locally demonstrated in a wafer of 1 cm thickness [23]), and scaling to 3.3 cm detectors with the same demonstrated bias voltage will ostensibly provide for a resolution of $7 \div 3.3 \sim 2$ eV$_{ee}$. This corresponds to a projected electron-equivalent recoil threshold of approximately 10 eV$_{ee}$, which scales under various extrapolations of the Lindhard factor to a threshold for nuclear recoils that is between 50 and 100 eV, allowing for a wide uncertainty pending the direct calibration of this low-energy regime.

Reaching the lowest 10 eV recoil sensitivity included in the present theoretical treatment will require further technological advances, although this goal would seem to be a plausible one with regards to basic underlying physics. Specifically, it will require the eventual realization of high-bias phonon-mediated ionization detectors with single-electron sensitivity [23]. The Lindhard factor at such low energies, or even at 100 eV, is yet to be established through direct nuclear recoil energy calibration. Current generation scattering experiments with a neutron beam targeted on commercial low-energy Ge detectors are capable of reaching 0.5 keV nuclear recoil thresholds, extending the work of Ref. [33]. Nuclear recoil calibration at very low energies, down to around 100 eV, is expected to be performed as part of the SuperCDMS SNOLAB program via the exposure of high-voltage detectors to either a neutron beam or activated source in an underground test facility. The ultimate goal is calibration down to the order of 10 eV nuclear recoils. In this regime, at the edge of single electron production, the practically achievable threshold will be sensitive to composition of the target material; for example, whereas a 10 eV threshold might be realistic in Ge, a number closer to 20 eV might be more appropriate to a material such as Si. At higher energies, on the order of 100 eV, the technologically limiting factors are less sensitive to such distinctions.

The placement of sensitive experimental apparatus in the near-field of a working nuclear reactor presents unique challenges associated with characterizing and modulating this radiologically intense setting. Primary backgrounds

TABLE II: We summarize expected characteristics of the first generation germanium detectors to be available for application to the described reactor sterile neutrino oscillation experiment. Demonstrated electron recoil resolution [23] is projected onto the target wafer thickness, and extrapolated under the assumption of Lindhard [32] scaling to a likely range of energy thresholds for the sensitivity to nuclear recoils. We allow for a conservative error factor of 2 in this projection, since the direct calibration of nuclear recoils in this energy regime remains to be established.

Dimensions ($\varnothing \times$ h)	100 mm $\times$ 33 mm
Mass	1.4 kg
Resolution σ $(E = 0)$	2 eV$_{\text{ee}}$
Electron $E_R^{\min}$	10 eV$_{\text{ee}}$
Lindhard Factor	$0.1 - 0.2$
Projected Nuclear $E_R^{\min}$	$50 - 100$ eV

inextricably linked with the targeted anti-neutrino flux include prompt gammas and neutrons sourced from the core. The energy spectrum and production rate of these backgrounds are predicted using a MCNP [34] core model developed at the Nuclear Science Center at Texas A&M University. Additionally, there are secondary background products derived from interactions with various casing and shielding materials. Sub-leading radiation levels from the decay of quasi-stable daughter products within the reactor, and from activated structural materials surrounding experimental cavity, persist even after the primary fission reaction has been powered down. Further, any surface-level experiment must also contend with muons and muon-induced neutrons from cosmic rays, as well as ambient gammas. The rate and spectral shape of detected recoils from such backgrounds must be either measured directly or estimated using simulation, and a shielding design must be introduced that is capable of reducing these backgrounds to a ceiling order-commensurate with (or below) that of the expected neutrino recoil rate.

Dual programs have been initiated for simulation in GEANT4 [35] of the experimental environment (including prototype shielding) and detector response to reactor neutron and gamma sources, and for the direct in-situ measurement of the same. The background shape and rate normalization now compare favorably between these two approaches. Certain details of the background estimation are provided subsequently in this section. The iterative process of measurement, simulation, and shielding design continues, and is described more comprehensively (and also more quantitatively) in a parallel publication with a primarily experimental focus [36]. In particular, the interested reader may there review in detail the nuclear simulation of neutron and gamma reactor flux rates (FIGS. 4), measurement of the gamma spectrum at various reactor power settings and core distances (FIGS. 7,8), comparison of the measured background shape and scaling against environmental simulation (FIGS. 9,10), and estimates of the shielding-controlled neutron and gamma flux rates in the signal region (FIG. 14).

In support of the present, primarily theoretical, analysis, we will begin from a target specification of the maximal background rate consistent with sensitivity goals, and then ask whether projected estimates of the corresponding shielding requirements are compatible with practical design limitations. Specifically, for observations in the nearest field, we will stipulate a unified (gamma and neutron) background rate in the detector of no more than 100 dru (differential rate unit, i.e. events/kg/day/keV) within the 10 to 1,000 eV region of interest, and will also secondarily consider a reduced background level of 10 dru. For observations performed at a more extended distance from the reactor, which benefit from additional natural depletion via reduction of the solid angle and increased line-of-sight integration depth through the pool (and also possibly a greater depth of applied shielding), but which likewise suffer from geometric reduction in the neutrino flux, we will target 10 or 1 dru, respectively. These values are suggested as controllable bounds relative to the presented projection for the CEνNS rate in Ge, which corresponds in turn to about 20 dru at one meter from core, about 5 dru at two meters, and about 1 dru at five meters. We have folded the background target into projections for the experimental reach, have projected the amount of data collection necessary in order to reach a measurement to 1% statistical resolution of these backgrounds in the experimental sideband above 1-2 keV, and have clarified the expectations for a shielding design capable of facilitating this goal.

Prompt gamma and neutron byproducts of fission in the reactor core are radiated with an integrated flux comparable to that of the signal neutrinos, i.e. on the order of 10^{12} [s·cm^2]$^{-1}$ at a distance of about one meter. These backgrounds diminish with distance-from-core (in a shield-dependent manner), but do not participate in the neutrino oscillation, and thus may potentially inhibit discrimination of the oscillation amplitude $\sin^2 2\theta_{14}$. Reactor gammas and neutrons are generated with an energy spectrum as given by the MCNP core model, produced at the face of the reactor core for a given reactor position. This has allowed for an estimation of the flux and deposited recoil energy at the detector location, when combined with a model geometry of the experimental hall (including an accurate description

of all materials) constructed in the GEANT4 framework. To test the simulation, in-situ measurements of the gamma background were taken with a HPGe (roughly 0.5 kg) detector and minimal lead shielding, located at the proposed experimental location. Energy deposition spectra were obtained with the core at various positions and power outputs, and compared to the prediction of the GEANT4 simulation. The simulation was found to accurately predict both the energy spectrum (on the order of 1% agreement) in the detector, as well as the scaling of the rate with distance from the core to within 10%. A preliminary shielding design was then added to the GEANT4 geometry model to asses the background expected in the full experimental setup. This shielding included both approximately 1.3 m of borated (5%) polyethylene as neutron shield and 30 cm of lead as gamma shield. Results indicate that backgrounds may be controlled to roughly the required order by adopting a shielding design corresponding to a mean core-to-detector separation of at least about 2 m. Such shielding was also found to be adequate for reduction of the ambient gamma background, which is much smaller than the reactor counterpart. A key lesson extracted from the early simulation data is that the interleaving of various shielding layers is essential to the suppression of backgrounds. The reasons for this are that each neutron capture will induce emission of a gamma, and likewise, high energy gamma rays may induce a secondary neutron cascade. Therefore, the sequential alternation of shielding materials respectively most effective against neutrons (e.g. poly) or gammas (e.g. lead) may be necessary in order to prevent the subsequent regeneration of a previously controlled background species.

Further optimization of the shielding design, including the described interleaving, is now underway. The described physical baseline of 2 m is apparently compatible with the proposed experimental objectives, and is presently adopted as a minimal practical separation of the detector and reactor centers. By increasing this separation, and back-filling the intervening space with additional shielding materials, it appears possible to further substantially reduce reactor backgrounds below the specified maximum. Mobility of the reactor core within its pool environment provides a free mechanism for variably enhancing early damping of the neutron component without any reconfiguration of the experimental column due to the considerable efficacy of simple water shielding against this background. Given that all backgrounds will tend initially to fall much more rapidly with distance in an appropriate shielding material than the geometric r^{-2} neutrino depletion, and then ultimately level out into a regime of deep suppression, one generically expects a metric of signal S to background B significance such as $S/\sqrt{S+B}$ to become maximized for some intermediate distance scale. As such, it remains possible that the experimental sensitivity to certain phenomena may be elevated by extending the minimal separation baseline, e.g. to three meters.

The interaction of a final state neutron with the detector will be dominantly in the form of a nuclear recoil, whereas incident gammas will interact prominently with the atomic electron cloud. At large recoils, independent phonon and ionization measurements facilitate an event-by-event discrimination of these two types of interaction. However, at lower recoil energies, below about 1 keV, convergence of the respective yield curves precludes such an identification. In this regime, which corresponds to the region of interest for the detection of CEνNS, it may instead be advantageous to employ a large bias voltage in order to amplify the secondary Luke-Neganov phonon (i.e. phonons shed by drifting charge carriers) signal far above the energy scale of the original recoil event. In so doing, the threshold of sensitivity to very soft recoils may be dramatically extended, with the fundamental energy scale reconstructed from inference of the number of electron-hole pairs created in the impact, as inferred in turn from the phonon calorimetry, which is proportional to this quanta times the applied voltage. Incidentally, this also implies that the available energy binning is not strictly arbitrary, but is rather dependent upon the stepwise function relating quantized production of electron-hole pairs to recoil energy. An energy resolution below 10 eV is possible [23], which is consistent with the spirit of the example binning in TABLE I. Monitoring of backgrounds will be performed with dedicated detectors, including both scintillation-based detectors placed at strategic locations within the proposed shielding, and a dedicated CDMS iZip-type detector mounted alongside the neutrino detectors. These monitoring detectors will provide normalization and discrimination for backgrounds above 1 keV recoil energy, which will provide a strong constraint on the background in the signal region when coupled with the simulation described previously.

The signal rate associated with CEνNS in Ge is suppressed by a factor of about 200 from 10 eV to 1 keV, and is practically non-existent above 2 keV. By contrast, the reactor backgrounds extend far above this scale, and may thus be precisely calibrated by observation of the experimental side-band region where CEνNS events are unpopulated. Simulation and preliminary observations suggest that the recoil spectrum in this regime is essentially flat, as consistent with expectations for low-energy Compton scattering. It will be necessary to even more carefully model the background shape in the signal region for purposes of the planned experiment, but it is possible already to make certain important observations. Statistical calibration of the background normalization to, say, 1% will require a number of events on the order of $N \simeq 10^4$ in order to satisfy $\sqrt{N}/N \simeq 0.01$. Taking, for example, a conservative side-band region extending from 2 keV up to about 10 keV (in reality, there will be usable normalization data far above this scale) this resolution would be achievable after approximately just a single running day with a single kilogram detector at a background rate of 100 dru, and in less than two weeks at 10 dru. Although this analysis is quite simplistic, it conveys an essential point: especially in the regime that backgrounds are large enough to compete with signal, then they are also quite easy to calibrate for the purpose of generating a data-driven background estimate in the signal region of interest.

In addition to the described reactor backgrounds, there are cosmogenic backgrounds, primarily in the form of muon-induced neutrons (e.g. via spallation in the concrete reactor pool housing), which do not decline with distance from core[2]; backgrounds of this type are particularly difficult for samples taken with the reactor at the far positions, as their uniform deposition rate at all detector baselines integrates to a much more substantial fraction of all events as the exposure time T_n is escalated to compensate for the geometric flux dilution. It is possible, however, to estimate this rate and to make an appropriate correction to each observation; the correction factor may be modeled computationally, or (better) extracted from data as a DC signal component that does not conform to a power law or oscillatory profile, or (best) measured directly during reactor off-cycles (although one must be careful to account for the residual activity of medium-lived daughter products (e.g. from the neutron capture-induced breeding of ^{238}U to ^{239}U, which beta-decays over 2-3 days to ^{239}Pu via ^{239}Np). In-situ measurements of the muon rate at the experimental site have now been collected, and show a reduction of roughly a factor of 4 due to the concrete overburden (about 13 mwe) of the reactor pool wall. Direct discrimination of the associated (unified) neutron event rate is furthermore achievable in the same manner previously described for reactor-sourced neutrons, via implementation of a dedicated detection module with supplementary ionization data capture. Ultimately, this potentially dangerous background is expected to be quite well controlled relative to an intrinsic rate of around a quarter hertz per kilogram. Beyond the described overburden, active veto technology is available with approximately 97% efficiency, and (most simply and most importantly) a further thousand-fold reduction is achieved through a kinematic cut above the CEνNS threshold $E_{\rm R} \lesssim 1$ keV.

Given the close proximity of source and detector described in this document, it is natural to ask when one should expect the point-like approximation of each physically extended object to remain valid. In particular, the spatial distribution of active flux sources (as well as the somewhat more narrowly confined distribution of measurement positions) will induce additional smearing in the basis-functions presented in FIG. (1). Such smearing may be expected to manifest as perturbative corrections in the ratio of the core size, i.e. 0.3 m, over the experimental baseline, which is typically close to (and may be substantially more than) an order of magnitude larger. Given the smallness of this parameter, we expect that the dominant effects along these lines arise instead from the previously modeled convolution of the oscillation response with the reactor energy spectrum and the recoil differential cross-section. Incidentally, such correction terms would break the scale-invariance manifest in FIG. (1). Although the baseline projections made in this document should be insensitive to these details at leading order, it will be important to more completely simulate such finite-size effects as the described project approaches the data collection phase.

It is worthwhile also to consider several additional sources of systematic experimental uncertainty, which may impose a ceiling on the expected statistical resolution. A very important uncertainty exists with regards to the overall anti-neutrino flux normalization (and to a lesser degree, shape). Errors are propagated from imperfect knowledge of the reactor thermal power, and from the extrapolation of this power into the associated anti-neutrino spectrum. Reactor operators are able to provide precision measurements of the thermal output power, as well as estimates (based on simulation with the code MCNP [37]) of the isotopic fuel composition and fission fractions f_i/F, where f_i is the absolute fission rate of species i and $F \equiv \sum f_i$. Uncertainty estimates on the order of 2-3% are typical, although it may be possible to reduce this to around a half of a percent (cf. Ref. [38]). However, it remains to be clarified precisely how the cited improvements may scale from the context of a commercial power reactor to that of a TRIGA research reactor. Calibration of direct recoil observations in the experimental sideband are also able to independently aid determination of the appropriate normalization factor.

The experimental apparatus is furthermore itself vulnerable at some level to both binary (false or failed trigger response) and continuous (smearing of the recoil energy resolution) errors, although these effects are expected to be rather small, with efficiencies rapidly approaching 100% within the fiducial volume and just above the recoil threshold. Millisecond recovery time in the detector modules eliminates substantial danger of event pileup. Finally, additional (minor) uncertainties may be propagated from errors in input parameters (e.g. th renormalization group evolved gauge couplings and Weinberg angle or the Z-boson mass), limitations in the analysis (e.g. the omission of electron scattering relative to CEνNS events or imperfect accounting of the dispersive spatial integration about the mean distance between reactor and detector centers), or even the inadvertent observation of alternate modes of new physics (e.g. Z' scattering or neutrino magnetic moment scattering).

The leading component of the described systematics is expected to proportionally scale the SM and oscillated (loss of signal) event rate at all energies and all distances, or to present an additive background event profile which may be estimated and calibrated in the experimental sideband. In the search for exotic interaction vertices, it is very useful

---

[2] We emphasize here that the experimental plan is to hold the detector arrays and shielding at a fixed location, as the reactor core is track-mounted and mobile. This is a significant advantage of the host reactor setting. Benefits of this architecture include increased mechanical stability of of the experimental apparatus, and increased geometric stability of each detector's complex relationship with its surroundings. Specifically, this carries over to an expectation for immutability of the cosmogenic backgrounds.

to combine scattering observations from multiple nuclei, such as Ge and Si, in order to cancel systematics and resolve signatures in the differential coupling to protons and neutrons. In the present case, the unique spatial variation of the signal response provides an intrinsic mechanism for the reduction of systematics with a single scattering target. In this environment, there is an advantage to employing a dominantly Ge based detector paradigm, as it delivers a threefold advantage in the CEνNS interaction rate per kilogram at low threshold. To the extent that systematics are primarily attributable to a universal scaling error, they may be dealt with in a particularly tidy fashion, as will be clarified in Section VI. Generalized methods applicable beyond this simplified case are described in Section VII.

In summary, the described experimental environment is host to a difficult conglomeration of particle backgrounds, each presenting their own unique profile and complications. However, preliminary study indicates that a careful combination of shielding, vetoes, simulation, and calibration may be sufficient to firstly curtail competing rates within the order of definite projections for the SM CEνNS rate, and secondly characterize the shape and scale of unified residual recoil backgrounds and systematics to a precision and accuracy facilitating their subtraction. It is important here to emphasize that the CEνNS signal is not an inherently rare one; although it is has thus far eluded experimental detection, technology sufficient to register the associated soft nuclear recoils, in conjunction with a source presenting a sufficiently large neutrino flux, will certainly deliver observations of coherent nuclear scattering with very robust statistics. Such a setting constitutes an ideal laboratory for probing deviations from the SM in the neutrino sector. Specifically, the context of mixing with a sterile neutrino sector provides a profound handle for resolving new physics embedded within even an imperfectly characterized background via fitting of an oscillation template correlated across length scales and across the spectrum of recoil energies, facilitating the cancellation of leading systematics with a non-oscillatory profile.

IV. EXPOSURE PER BASELINE OBSERVATIONAL SCHEDULE

A rough estimate of the exposure necessary to attain statistical significance may be constructed as follows. From Eq. (3), the maximal anti-neutrino disappearance fraction is $\sin^2 2\theta$, assuming extremization of the position-dependent term (every length L will realize this criterion for certain discretized incident neutrino energies E_ν). Neglecting systematic effects for the time being, the statistical significance may be gauged by the ratio of the extremal event deficit $N_{\mathrm{Exp}} \sin^2 2\theta$ over the fluctuation $\sqrt{N_{\mathrm{Exp}}}$ in the SM expected event rate N_{Exp}. Integrating across energy bins, the net annual SM CEνNS expectation for a $M = 1$ kg Ge detector with 10 eV (100 eV) kinetic recoil sensitivity at a distance of 1 m from a megawatt nuclear reactor source is approximately $N_{\mathrm{Exp}} = 8 \times 10^3$ (4×10^3) events. Setting the significance ratio to unity, and taking into account the dilution of flux with distance L^2 from the core, the minimal exposure time in a given sampling bin is around

$$T_{\mathrm{Min}} \simeq 1 \ [\mathrm{y}] \times \frac{1.3 \times 10^{-4}}{\sin^4 2\theta} \times \left\{ \frac{1 \ [\mathrm{kg}]}{M} \right\} \times \left\{ \frac{L}{1 \ [\mathrm{m}]} \right\}^2 . \tag{7}$$

For example, adopting the stipulated reactor power with a single kilogram detector, and taking an oscillation amplitude $\sin^2 2\theta \simeq 0.1$, the minimal integration time for onset of statistical resolution would be about five days at a distance of $L = 1$ m, or about 120 days at a distance of $L = 5$ m. The escalation to a 100 kg detector would offset a ten-fold reduction in the signal amplitude to $\sin^2 2\theta \simeq 0.01$ with identical exposure. In case of a 100 eV threshold, the event rate is reduced by about 50% and the sensitivity to fast oscillations is reduced by about 30%. However, forfeiture of the low-energy spectral data implies a loss of independent short-distance information, which is potentially more damaging than the simple loss of statistics.

In order to resolve the underlying oscillatory character of the signal, one is clearly motivated to sample multiple points in the oscillation profile. To some extent, this can be accomplished passively, by sampling multiple neutrino energies E_ν (or multiple kinetic recoil energies E_R^i) at a fixed distance. In this manner, cf. FIG. (1), regular trends in the event deficit might be resolved within a single vertical constant-L slice of the binned response curves. However, it is obviously preferable to complement this approach with sample data that is literally extended in space, such that various candidate signal wavelengths may be probed directly. If this is possible, then independent likelihood optimizations of the $\sin^2 2\theta_{14}$ and Δm_{14}^2 parameters within an oscillation template may be extracted from each binned energy range, and subsequently combined into a unified signal fit and error estimate.

Intuition suggests that the best statistical power in an oscillation template fit will emerge from multiple samples located commensurately with the signal half-wavelength. As highlighted in FIG. (1), data samples in the extreme near-field are very poorly suited for resolving oscillation features, which have not had a sufficient baseline over which to mature. Conversely, measurements in the very far-field will have dispersed beyond the level at which isolated features may be resolved. Equally tempered spacings at $\lambda/4$ or wider experience an elevated danger of coincidentally aligning with the signal troughs. These observations suggest a logarithmic sample locating schedule, with a local density that separates adjacent measurements by something like a quarter of their mean distance from core.

The prescription for the location L_n of the n^{th} sample point out of $(N + 1)$ is as follows, in units of the primary $(n = 0)$ observation scale $L_0 \equiv 1$, as a function of the scaled position of the final $(n = N)$ observation at $L_N \equiv L$.

$$L_n = L^{n/N} \tag{8}$$

Near detector samples have a very distinct statistical advantage in terms of the $1/L_n^2$ flux enhancement, which radically reduces requisite integration times relative to those of far samples. In order to accrue a matching event count at longer distances, an offsetting quadratic compensation $T_n \propto L_n^2$ in the exposure time is required. The constant of proportionality may be established by constraining the sum over exposures to match a specified net interval T, after which the time T_n per sample location L_n is prescribed as follows.

$$T_n = T \times L^{2n/N} \times \left(\frac{L^{2/N} - 1}{L^{2+2/N} - 1} \right) \tag{9}$$

The increased temporal cost of far detector samples is attenuated by elongation of the sample spacing. Like the fretting intervals on a guitar, each "octave" (doubling of the experimental baseline) will contain an identical number of sampling locations in this prescription, although it will necessitate a four-fold increase in the cumulative exposure time to achieve similar statistical significance. It should be emphasized that the described accommodation provides for the delivery of equal statistical weight from all sample locations, but it simultaneously escalates the impact of backgrounds, primarily cosmogenic, that do not abate with distance in the far samples.

V.   ESTIMATION OF EXPECTED SENSITIVITY

FIG. 2: Figures depict contours in the χ^2 statistical significance, specifically $S/\sqrt{B}$, where the signal S corresponds to an event deficit attributable to oscillation with a sterile fourth generation neutrino relative to the background expectation B for the SM electron antineutrino CEνNS event rate in Ge. Scale invariant projections are made in terms of the mass gap Δm_{14}^2 eV^2 and the amplitude $\sin^2 2\theta_{14}$, as a product or quotient with the experimental baseline L m, respectively, and times the square root of exposure in the latter case. At unit distance, exposure time and mass these factors drop out, and the axes may be read traditionally. The detector recoil energy is unbinned in this analysis. The left- and right-hand panel integrates all recoils above the more conservative and more optimistic thresholds of 100 and 10 eV, respectively.

The oscillated event expectation N_{Osc}^1 in the i^{th} E_{R}^i recoil energy bin is given in terms of the baseline SM expectation N_{Exp}^i, the oscillation amplitude $\sin^2 2\theta_{14}$, and the convolved deviation shape functional $\gamma_i(\Delta m_{14}^2 L)$ of Eq. (5).

$$N_{\text{Osc}}^i = N_{\text{Exp}}^i \times \left\{ 1 - \sin^2(2\theta_{14}) \, \gamma_i(\Delta m_{14}^2 L) \right\} \tag{10}$$

In the absence of data, it is still quite possible to roughly estimate the sensitivity of a counting experiment to deviations from the null result in a simplified manner. Referencing Eq. (10), we construct a χ^2 statistic comparing the deviation-squared of the oscillated signal N_{Osc}^i from the SM CEνNS expectation (which is construed now as background) to the

associated statistical uncertainty $\sigma_i \sim \sqrt{N_{\text{Exp}}^i}$ in this background rate. We sum over C observation channels, where the binning index i momentarily performs double duty, labeling both the targeted range of recoil energies and the detector location.

$$\chi^2 \equiv \sum_{i=1}^{C} \frac{(N_{\text{Osc}}^i - N_{\text{Exp}}^i)^2}{N_{\text{Exp}}^i} = \sin^4 2\theta_{14} \times \sum_{i=1}^{C} \gamma_i^2 \, N_{\text{Exp}}^i \tag{11}$$

We may (and will) further consider inclusion of projections for the unreduced non-CEνNS experimental backgrounds (as described in Section III) into a unified event profile baseline, above which the amplitude of sterile neutrino induced oscillation should be statistically visible. We justify neglect of systematic uncertainties on the shape and normalization of the unified background event profile in two ways. Firstly, as described in the prior section, the amount of data collection required to suppress uncertainty in the sideband calibration of non-CEνNS backgrounds has been demonstrated to be well within reach. Secondly, as will be elaborated in the subsequent sections, a more complete analysis using log-likelihood techniques with an oscillatory template hypothesis may be formulated to induce cancellation of leading order uncertainties that do not present with an oscillating event profile. Note further (for example) that a 1 year run will give on the order of 10^5 recoils at one meter and the flux normalization would carry a 0.3% (0.5% for $E_R > 100$ eV) uncertainty, suggesting (in this case) that systematic flux uncertainties in reactor neutrino and neutron backgrounds are subdominant for $\sin^2 2\theta \gtrsim 0.01$.

In the limit where many stochastically dispersed binning channels C are sampled with an approximately uniform distribution of expected counts $N_{\text{Exp}}^i \simeq N_{\text{Tot}}/C$, the value of Eq. (11) will converge to $\chi^2 \to 3/8 \, N_{\text{Tot}} \sin^4 2\theta$, where the numerical coefficient represents a fourth moment $\langle \sin^4 \rangle = 3/8$ of the sinusoid embedded within γ_i. The result is independent of the number of channels C, and is identical to the scenario where samples are unbinned. This indicates that statistical significance of the deviation declines in this scenario with the isolation of samples into multiple bins, because the fixed χ^2 value is then distributed over more degrees of freedom C. The result is readily understood, and is attributable to the fact that the sign of γ_i is always positive, i.e. the sterile neutrino always effects a downward fluctuation in the event rate.

The χ^2 significance of the oscillation-induced anti-neutrino deficit relative to the statistical background at a single experimental baseline L, and with no binning in the nuclear kinetic recoil, is projected in FIGS. (2) as a function of Δm_{14}^2 and $\sin^2 2\theta_{14}$. Specifically, the exhibited contours correspond to integer values in the range of 1 to 5 for the test statistic $S/\sqrt{B}$, interpreting S as the event depletion attributable to sterile neutrino oscillation effects, and B as the unified SM CEνNS event expectation. Results are presented for both conservative and optimistic nuclear recoil thresholds of 100 and 10 eV. In order to manifest the generality and elegance of underlying scaling symmetries with respect to distance L and the exposure time T and mass M, backgrounds which do not respect the same scaling law are neglected in these figures. As expected from Eq. (3) and FIG. (1), observability is greatly diminished in the vertical axis whenever $(\Delta m_{14}^2 \, \text{eV}^2 \times L \, [\text{m}] \ll 1)$, as there is insufficient phase evolution. Likewise, as suggested by Eq. (11), observability in the horizontal axis is hampered by reduction of the oscillation amplitude $\sin^2 2\theta_{14}$, reduction of the exposure, or by elongation of the separation from core (via geometric reduction in the neutrino flux as $N_{\text{Exp}}^i \propto 1/L^2$).

FIGS. 3 show the specific sensitivity projected for a 100 kg payload with a 2 yr exposure, using the phase-1 recoil sensitivity threshold of 10 eV. Figures are presented with both a 2 and 5 m separation between the detector and the reactor core. Additionally, unified neutron, gamma, and cosmic (etc.) background event rates of 100 and 10 dru are folded into the analysis in the near and far positions, respectively. For purposes of rough comparison (approaches to limit setting and handling of systematics are not identical), global fit contours at 95% confidence for short-baseline (blue dashed) and ν_e disappearance (red solid) from Ref. [39], as well as projected limits from SOX [19] and Solar + Kamland [40], are overlaid. As indicated, there is good sensitivity overlap with the primary region of interest, although the described phase-1 scenario is limited with respect to probing the low mass-gap parameter space.

FIGS. 4 repeat the prior analysis with escalation to a 1 ton payload and a 5 year exposure. Simultaneously, we adopt for this phase-2 analysis, an improved (speculative) recoil sensitivity threshold of 10 eV. In keeping with this more optimistic and forward-looking scenario, we also assume shielding improvements such that background rate is limited to 10 and 1 dru at 2 and 5 meters, respectively. The projected discovery contours for this second analysis essentially covers the allowed space of Δm_{14}^2 and $\sin^2 2\theta_{14}$ values associated with global fits to reactor and gallium experiments [39], and projected sensitivity to the $\bar{\nu}_e$-sterile mixing is competitive with that of the SOX experiment [19].

If the reactor source and the expected SM background rate N_{Exp}^i are very well understood, then the uniform depletion of events at long distance scales in the $(\gamma_i \Rightarrow 1/2)$ dispersion limit is still quite visible. However, a constant deficit of this type is not intrinsically separable from systematic errors in the event rate normalization. Again, the inherent advantage of the oscillation signal is its spatial structure, which allows for the handy cancellation of systematics in the underlying event rate. On the other hand, this also implies that the detector will then only be sensitive to oscillation wavelengths λ on the order of the distance to core L. The methodology appropriate to an analysis of this second type is developed subsequently.

VI.   MEASUREMENT OF PARAMETERS AND QUANTIFICATION OF STATISTICAL BOUNDS

When data is available, either of the experimental or Monte Carlo variety, an analysis quite different than that described in Section V may be conducted, wherein the observed counts N_{Obs}^i in the i^{th} distance and/or energy bin are compared to the expectation value N_{Osc}^i, cf. Eq. (10), under an oscillation hypothesis template. Allowing for an undetermined renormalization $\alpha \simeq 1$ of the overall event rate expectation, the probability of observing a deviation of N_{Obs}^i events relative to the rate αN_{Osc}^i is normally distributed with respect to the bin uncertainty width σ_i.

$$\mathcal{P}_i \equiv \mathcal{P}(N_{\text{Obs}}^i) = \exp\left\{ -\frac{\left[N_{\text{Obs}}^i - \alpha N_{\text{Osc}}^i \right]^2}{\sigma_i^2} \right\} \tag{12}$$

Given definite values for the binned observation N_{Obs}^i, expectation N_{Exp}^i and uncertainty σ_i, we may invert the prior interpretation such that the likelihood $\mathcal{L}_i \equiv \mathcal{L}(\alpha, \sin^2 2\theta, \Delta m^2)$ of the actual underlying global scale, amplitude, and phase parameters aligning with some specified triplet of floating values is numerically proportional to (or, conveniently, equivalent to) the Eq. (12) probability density $\mathcal{L}_i \equiv \mathcal{P}_i$. The joint likelihood $\mathcal{L}_C = \prod_{i=1}^{C} \mathcal{L}_i$, which is a product of likelihoods for each of the C binned observation channels, remains functionally dependent upon just three global degrees of freedom. The negative-log $\widetilde{\mathcal{L}_C} \equiv -\ln \mathcal{L}_C$ of this quantity varies monotonically with the original likelihood, and is minimized at the same point in parameter space at which the former is maximized; it furthermore converts the product into a sum over χ^2-type terms.

$$\widetilde{\mathcal{L}_C} = \sum_{i=1}^{C} \frac{\left[N_{\text{Obs}}^i - \alpha N_{\text{Osc}}^i \right]^2}{\sigma_i^2} \tag{13}$$

The joint log-likelihood $\mathcal{L}_C$ may be minimized numerically in order to fix each of the parameters α, $\sin^2 2\theta$ and Δm^2, while establishing error bounds associated with deterioration of the fit as their values flow away from the optimization. However, it is instructive and potentially useful to first semi-analytically reduce this expression with regards to the scale renormalization α. Setting $\partial \widetilde{\mathcal{L}_C} / \partial \alpha = 0$ via Eq. (13), and referencing Eq. (10) with the compact notation
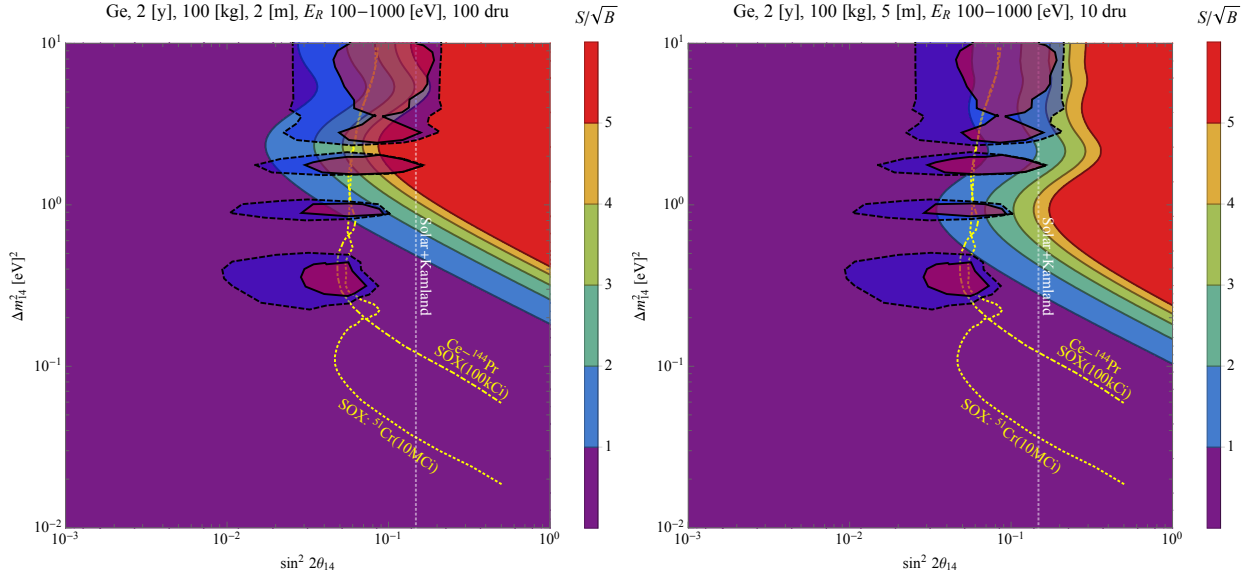

FIG. 3: Phase-1 prospective limit of $\nu_e - \nu_s$ mixing parameters with 100 kg Ge detector mass and 2 year effective exposure at a sample distance of 2 m (left) or 5 m (right) from the reactor. A minimum recoil threshold of 100 eV is adopted. Contours are shown in the χ^2 statistical significance, specifically $S/\sqrt{B}$, where the signal S corresponds to the oscillation event deficit and the background B sums the SM electron anti-neutrino CEνNS event rate together with a postulate for the unified reactor plus cosmic, etc. event rate. This additional background component is modeled with a strength of 100 or 10 dru at the near and far positions, respectively. Global 95% short-baseline (blue dashed) and ν_e disappearance (red solid) fits [39] and projected SOX [19] and Solar + Kamland [40] limits are plotted for rough comparison.

$\widehat{s} \equiv \sin^2 2\theta_{14}$, we get the following.

$$\sum_{i=1}^{C} \frac{\left[N^i_{\mathrm{Obs}} - \alpha\, N^i_{\mathrm{Exp}} \times \{1 - \widehat{s}\,\gamma_i\}\right] \times N^i_{\mathrm{Exp}} \times \{1 - \widehat{s}\,\gamma_i\}}{\sigma_i^2} = 0 \tag{14}$$

This conditional does not, in general, lead to a compact closed-form expression for α, and certainly not one that is independent of the detailed structure of the shape functionals γ_i. However, certain well-motivated approximations will induce a substantial simplification. Firstly, we will now restrict the interpretation of Eqs. (13,14) to a single binned energy range at a time, e.g. one of the TABLE I aggregations, restricting the summation index n to trace only over the range of distances L_n from core. Employing the Eq. (9) time allocation convention, the expected event count N^i_{Exp} in the i^{th} energy bin is constant with respect to sampling done across positions; this is an approximation, with associated systematic error contributions, e.g. insomuch as it relies upon the uniform subtraction of previously elaborated cosmogenic and related backgrounds that do not decline with distance-from-core. Given this, it is reasonable to likewise stipulate an energy-binned error width σ_i that does not fluctuate with the position n. Finally, we will marginalize over the value of the $\sin^2$ spatial oscillation shape $\gamma_i \to \langle \gamma_{i,n} \rangle = 1/2$. This choice implies that we must later be somewhat careful to elect a uniform distribution of spatial samples, with respect the expected fluctuations of $\gamma_i(\Delta m^2 L)$ about the mean. However, that is not difficult, as the distances scales L_n most sensitive to the oscillatory signal component at a given mass gap Δm^2 are precisely known. With these modifications, Eq. (14) reduces to the following condition with respect to the mean observed count per observation $\langle N^i_{\mathrm{Obs}} \rangle \equiv \sum_{n=0}^{N} N^{i,n}_{\mathrm{Obs}} / N$ (or the obvious modification thereof when n is restricted to a subset of $[\,0, N\,]$ that ensures $\langle \gamma_{i,n} \rangle \simeq 1/2$).

$$\alpha = \frac{\langle N^i_{\mathrm{Obs}} \rangle}{N^i_{\mathrm{Exp}} \times \{1 - \widehat{s}/2\}} \tag{15}$$

The interpretation of this result is straightforward; the midline of the expected oscillation depicted in FIG. (1) has been renormalized to coincide with the mean of observations over various experimental baselines L_n. This scaling induces the cancellation of leading systematic errors associated, for example, with the reactor flux normalization. Now, we will assign the uncertainty width $\sigma_i \to \sqrt{\langle N^i_{\mathrm{Osc}} \rangle}$ attending each observation in keeping with the dominant residual statistical fluctuation. It is simultaneously convenient to define a pair of measures of the oscillatory deviation


FIG. 4: Phase-2 prospective limit of $\nu_e - \nu_s$ mixing parameters with 1000 kg Ge detector mass and 5 year effective exposure at a sample distance of 2 m (left) or 5 m (right) from the reactor. A minimum recoil threshold of 10 eV is adopted. The additional background component is modeled with a strength of 10 or 1 dru at the near and far positions, respectively. Contours and overlays are as described in FIGS. 3.

from the mean, for the observed and expected signal, respectively.

$$\Delta_{\text{Obs}}^{i,n} \equiv \frac{N_{\text{Obs}}^{i,n}}{\langle N_{\text{Obs}}^i \rangle} - 1 \tag{16}$$

$$\Delta_{\text{Osc}}^{i,n} \equiv \hat{s} \times \left\{ \frac{1 - 2\gamma_{i,n}}{2 - \hat{s}} \right\} \tag{17}$$

In terms of these factors, the joint log-likelihood referenced in Eq. (13) for all observations in a given energy bin may be succinctly rephrased as follows, where the subscript on $\widetilde{\mathcal{L}_2^i}$ references the pair of remaining degrees of freedom in the template, namely $\sin^2 2\theta_{14}$ and Δm_{14}^2.

$$\widetilde{\mathcal{L}_2^i} = \langle N_{\text{Obs}}^i \rangle \times \sum_{n=0}^{N} \left\{ \Delta_{\text{Obs}}^{i,n} - \Delta_{\text{Osc}}^{i,n} \right\}^2 \tag{18}$$

The goodness of a fit to data may be quantified by the comparison of this factor to the zeroth order fit $\widetilde{\mathcal{L}_0^i}$, where parameterization freedom has been redundantly eliminated by simultaneous application of the SM limits $\hat{s} \to 0$ and $\gamma_i \to 1/2$, which imply $\Delta_{\text{Osc}}^{i,n} \to 0$.

$$\widetilde{\mathcal{L}_0^i} = \sum_{n=0}^{N} \frac{\left\{ N_{\text{Obs}}^{i,n} - \langle N_{\text{Obs}}^i \rangle \right\}^2}{\langle N_{\text{Obs}}^i \rangle} = (N+1) \times \frac{\sigma_i^2}{\langle N_{\text{Obs}}^i \rangle} \geq (N+1) \tag{19}$$

The expected variance σ_i^2 should be comparable to $\langle N_{\text{Obs}}^i \rangle$ if it arises from a purely statistical mechanism, but may be (substantially) larger if the signal has a substantial oscillatory component that is not well modeled by the flat zeroth order template. The result is expected to grow in proportion to the number of elected bins. There is expected to be an advantage for $\widetilde{\mathcal{L}_2^i}$ over $\widetilde{\mathcal{L}_0^i}$ with respect to the inclusion of additional sample points if and only if the signal is within the strongly oscillating central region of the $\Delta m^2 L$ space, cf. FIG. (1). Judicious preselection of the indices n to be retained in each region of the likelihood optimization should be performed in a manner consistent with the $\langle \gamma_{i,n} \rangle \simeq 1/2$ condition.

Wilk's theorem states that twice the difference of negative-log-likelihoods for nestable model templates is approximately χ_D^2 distributed, with degrees of freedom D equal to the difference in number of optimized parameters. Specifically, the criterion for a significant improvement, at a type-I error level p, for the 2-parameter template fit over the constrained fit is as follows, where the cumulative distribution function (CDF) represents the fraction of parameter space bounded within a multi-dimensional (typically Gaussian) integration out to some "radius" χ.

$$\widetilde{\mathcal{L}_2^i} \leq \widetilde{\mathcal{L}_0^i} - \text{CDF}^{-1}(\chi_2^2, 1 - p)/2 \tag{20}$$

The inverse CDF is simply the χ_D^2 boundary value in D dimensions for which a fraction p of possible outcomes would be considered more extreme. For example, with $D = 2$, the inverse CDF χ^2-values for confidences levels $(1 - p)$ corresponding to $\{68, 95, 99.7\}\%$, i.e. $\{1, 2, 3\} \times \sigma_i$, are $\{2.3, 6.2, 11.8\}$, respectively. This is a one-tailed test, insomuch as the negative-log-likelihood of the parameterized fit must necessarily meet or exceed that of the constrained fit. Eq. (20) may alternatively be inverted to solve for the p-value, which may be converted into an equivalent significance multiple $\mathfrak{N} \times \sigma_i$ of the one-dimensional Gaussian standard deviation.

$$p = 1 - \text{CDF}\big(\chi_2^2, 2 \times \big[\widetilde{\mathcal{L}_0^i} - \widetilde{\mathcal{L}_2^i}\big]\big) \tag{21}$$

$$\mathfrak{N} = \sqrt{\text{CDF}^{-1}(\chi_1^2, 1 - p)} \tag{22}$$

If the oscillation template is demonstrated to be significantly superior to the null template in a given energy bin i, then uncertainty bounds on the optimized parameters $\sin^2 2\theta_i$ and Δm_i^2 may be established via the computation of likelihood profiles. In turn, the one-dimensional likelihood $\widetilde{\mathcal{L}_1^{i,j}}(q)$ of the j^{th} parameter fit is sampled for values q^j adjacent to the optimization q_0^j, while all conjugate parameters are re-optimized. Then, the cutoff at a level p corresponds to the value of q^i for which $\widetilde{\mathcal{L}_1^i}(q^j) - \widetilde{\mathcal{L}_1^i}(q_0^j)$ is equal to $\text{CDF}^{-1}(\chi_1^2, 1 - p)/2$. Results for each parameter may then be consolidated across independent energy bins in order to establish combined limits.

It is expected that the sensitivity delivered by this type of analysis (or similarly for the generalized approach sketched in the final section), where one takes full advantage of the additional handles provided by the oscillatory

event profile (including embedded correlations across distance and recoil energy bins), could substantially exceed the somewhat naïve sensitivity estimates computed in Section V by comparing the unbinned signal amplitude to the scale of fluctuations in the unified CEνNS observation rate (including non-CEνNS backgrounds) at a single experimental baseline. It is of great interest to the authors to undertake a full trial analysis of simulated data with Monte Carlo event population, folding in realistic backgrounds and experimental systematics, applying the statistical techniques described in the present work to the extraction of signal and the setting of limits. Only by such means is it possible to extract information encoded in the variable experimental baseline and spectral decomposition of the recoil energy spectrum while deeply suppressing systematic uncertainties in the underlying non-oscillatory event profile. We briefly lay out the formalism in Section VII. This exercise, which would require optimization of the realistic measurement plan and exceeds the scope of the current document, is reserved for future work.

VII. BACKGROUNDS AND A GENERALIZED FIT

The recoil measurement will face several backgrounds that need to be taken into account. In this section we discuss a general formalism of the likelihood fit, and a proof-of-principle treatment of the major background categories listed in Section III. To build a general template for the recoil event rate, let us formulate the total recoil event number at each detector location as

$$N_i = T_i \int_{E_{R,i}^{min}}^{E_{R,i}^{max}} dE_{\mathrm{R}} \left(\frac{L_0}{L_i}\right)^2 \left(\alpha f_\nu(E_{\mathrm{R}})\frac{dN_\nu}{dE_{\mathrm{R}}} + \beta f_n(E_{\mathrm{R}})\frac{dN_n}{dE_{\mathrm{R}}}\right) + \gamma f_{\mu n}(E_{\mathrm{R}})\frac{dN_{\mu n}}{dE_{\mathrm{R}}}, \tag{23}$$

where i is the index for different detector locations at a distance L_i from the reactor core, T_i is the time exposure, L_0 is a reference distance. E_{R} is the recoil energy. dN/dE_r represents the flux from reactor neutrino recoils ($\frac{dN_\nu}{dE_{\mathrm{R}}}$), the reactor gamma-ray induced (secondary neutron) events ($\frac{dN_\gamma}{dE_{\mathrm{R}}}$), and the cosmic muon induced neutrons ($\frac{dN_{\mu n}}{dE_{\mathrm{R}}}$) which have penetrated through the cement overburden. The f function denotes the energy-dependent detection efficiencies. We ignore the 'direct' reactor gamma-ray recoils, which can be efficiently eliminated by shielding.

$\{\alpha, \beta, \gamma\}$ are a set of nuisance parameters that allow each component to float within their normalization uncertainties $\{\Delta\alpha, \Delta\beta, \Delta\gamma\}$. By including active-sterile oscillations, the reactor neutrino events will depend on sterile neutrino mixing parameters $\{\Delta m^2, \sin^2 2\theta\}$,

$$N_i = N_i(\{\Delta m^2, \sin^2 2\theta\}; \{\alpha, \beta, \gamma\}). \tag{24}$$

The relative sizes between the oscillation signal and the normalization uncertainties play an important role in the fitting strategy. For a large sterile mixing angle, e.g. greater than the uncertainty on the reactor flux $\sin^2 2\theta > \Delta\alpha$, ν recoils yield a deviation (compared to the non-oscillation flux) no less than half of the $\sin^2 2\theta$ amplitude over the entire energy range, and it is desirable to combine the data at all recoil energy to enhance likelihood significance.

However, when $\sin^2 2\theta$ is comparable or less than the reactor flux uncertainty, the flattened half-amplitude-decrease from sterile neutrino oscillation becomes indistinguishable from a overall fluctuation in the normalization of the reactor neutrino flux. In this case it can be more desirable to focus on the range of E_{R} where the oscillation has not flattened from dispersion (see the central range of FIG 1).

With the estimated 100 dru (10 dru) background rates, a high-statistics calibration in the side-band can determine the leading background normalizations to be less than 1 percent accuracy in days (weeks), by running a calibration fit for the nuisance parameters,

$$\{\alpha_0, \beta_0, \gamma_0\} \text{ minimizes } \sum_i \frac{[N_i(\{0,0\}; \{\alpha, \beta, \gamma\}) - N_i^{obs}]^2}{N_i^{obs} + sys^2}, \tag{25}$$

where null sterile mixing parameters are assumed for calculating exclusion limits. Here the systematics will include any uncertainty other than the fluctuations represented by $\{\alpha, \beta, \gamma\}$. Next we can perform a fit to the sterile mixing parameters,

$$\chi^2 = \sum_i \frac{[N_i(\{\Delta m^2, \sin^2 2\theta\}; \{\alpha, \beta, \gamma\}) - N_i^{obs}]^2}{N_i^{obs} + sys^2} + \frac{(\alpha - \alpha_0)^2}{(\Delta\alpha)^2} + \frac{(\beta - \beta_0)^2}{(\Delta\beta)^2} + \frac{(\gamma - \gamma_0)^2}{(\Delta\gamma)^2} \tag{26}$$

At each $\{\Delta m^2, \sin^2 2\theta\}$ point, $\{\alpha, \beta, \gamma\}$ are marginalized over to obtain the best fit χ^2 value. The additional nuisance parameter terms take care of the effect of globally shifting the flux normalizations, and leave the number of degrees of freedom unchanged. In the statistics-dominated limit, the nuisance parameters will be tied closely to their central values, and the above formula reduces to Eq. 11.

The variation in energy and spatial distribution of backgrounds require the calibration process to be carried out at each reactor-detector distance of measurement. Optimization between observation time and detector location needs to be done in accordance with the total experimental time span. Well-placed cuts can also improve the signal to background ratio, for instance, requiring $E_R < 200\text{eV}$ can efficiently control cosmic-muon induced recoils, whose flat energy distribution weighs more on higher E_R in comparison to that of signal recoil events.

VIII.   CONCLUSIONS AND SUMMARY

In this paper we have discussed a variable multi-meter-scale baseline measurement of active-sterile neutrino oscillations in collaboration with the Nuclear Science Center at Texas A&M University. Our motivation is to test for the presence of a 4th generation sterile neutrino, which has been hinted at by several experiments [1–5]. A 4th generation sterile neutrino at the mass scale $\sim$ eV can be accommodated cosmologically, and is consistent with Solar neutrino data.

We have shown that the adjustable reactor-detector distance gives our measurement scheme a great advantage to sample different areas of the mixing parameter space, and that limits are competitive with projected sensitivities from SOX [19] and from detection of Solar neutrinos [40] with sufficient exposure. A distance as near as 2-3 m emphasizes sensitivity to the mixing angle, and farther positions provide improved reach for a small 4th generation neutrino mass-square difference. Measurements at 2 and 5 meters, with a total 100 Kg Ge detector mass over an effective exposure of 2 years probes a large fraction of the Δm^2 parameter space that is allowed by short-baseline experiments. With a conservative recoil energy threshold of 100 eV we project an estimated sensitivity to first/fourth neutrino oscillation with a mass gap $\Delta m^2 \sim 1\,\text{eV}^2$ at an amplitude $\sin^2 2\theta \sim 10^{-1}$, or $\Delta m^2 \sim 0.2\,\text{eV}^2$ at unit amplitude.

Depending on the performance of the phase-1 experiment, a phase-2 experiment with a 1000 kg payload may be proposed, to operate over a 5 year period to provide the maximum sensitivity achievable from the reactor experiment at TAMU. Such a next-generation experiment would be able to fully utilize and benefit from rapid technological progress in the reduction of minimal recoil sensitivities (proceeding in parallel for applications to low mass dark matter search experiments), which may conceivably be extended down to the 10 eV order. In this case, it is possible to probe more than an additional order of magnitude in amplitude. Such a second phase experiment might additionally be situated at an alternate site with higher neutrino flux for the ultimate sensitivity to sterile neutrinos.

We have presented results accounting for the inclusion of reactor and cosmogenic backgrounds and statistical uncertainties on the total experimental event rate. We have provided an estimate for systematics, which include the anti-neutrino flux normalization, and uncertainties in the other reactor and non-reactor backgrounds. We have shown that spatial variation of the signal response provides an intrinsic mechanism for the reduction of systematics, and have discussed methods for parameter estimation in the presence of both statistical and systematic uncertainties.

Neutrino mass provides the best evidence for physics beyond the Standard Model. Independent of the aforementioned experiments that may hint at the presence of an additional sterile neutrino, this provides clear motivation for us to understand the behavior of neutrinos over all accessible interaction channels and energy scales. The experiment that we describe provides a unique channel to probe neutrino interactions, and a promising method to search for physics beyond the Standard Model.

[1] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten, Phys. Lett. **B685**, 47 (2010), 1001.2731.
[2] J. N. Abdurashitov et al. (SAGE), Phys. Rev. **C80**, 015807 (2009), 0901.2200.
[3] C. Giunti and M. Laveder, Mod. Phys. Lett. **A22**, 2499 (2007), hep-ph/0610352.
[4] C. Giunti and M. Laveder, Phys. Rev. **C83**, 065504 (2011), 1006.3244.

[5] G. Mention, M. Fechner, T. Lasserre, T. A. Mueller, D. Lhuillier, M. Cribier, and A. Letourneau, Phys. Rev. **D83**, 073006 (2011), 1101.2755.

[6] Y. Bai, R. Lu, S. Lu, J. Salvado, and B. A. Stefanek, Phys. Rev. **D93**, 073004 (2016), 1512.05357.

[7] A. Aguilar-Arevalo et al. (LSND), Phys. Rev. **D64**, 112007 (2001), hep-ex/0104049.

[8] A. A. Aguilar-Arevalo et al. (MiniBooNE), Phys. Rev. Lett. **105**, 181801 (2010), 1007.1150.

[9] A. A. Aguilar-Arevalo et al. (MiniBooNE), Phys. Rev. Lett. **110**, 161801 (2013), 1207.4809.

[10] G. Steigman, Adv. High Energy Phys. **2012**, 268321 (2012), 1208.0032.

[11] P. A. R. Ade et al. (Planck), Astron. Astrophys. **571**, A16 (2014), 1303.5076.

[12] S. Hannestad, R. S. Hansen, and T. Tram, Phys. Rev. Lett. **112**, 031802 (2014), 1310.5926.

[13] B. Dasgupta and J. Kopp, Phys. Rev. Lett. **112**, 031803 (2014), 1310.6337.

[14] K. C. Y. Ng and J. F. Beacom, Phys. Rev. **D90**, 065035 (2014), [Erratum: Phys. Rev.D90,no.8,089904(2014)], 1404.2288.

[15] A. Palazzo, Phys. Rev. **D83**, 113013 (2011), 1105.1705.

[16] J. Ashenfelter et al. (PROSPECT), in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013) Minneapolis, MN, USA, July 29-August 6, 2013* (2013), 1309.7647, URL `http://inspirehep.net/record/1256022/files/arXiv:1309.7647.pdf`.

[17] D. A. Dwyer, K. M. Heeger, B. R. Littlejohn, and P. Vogel, Phys. Rev. **D87**, 093002 (2013), 1109.6036.

[18] Y. Gao and D. Marfatia, Phys. Lett. **B723**, 164 (2013), 1302.5725.

[19] G. Bellini et al. (Borexino), JHEP **08**, 038 (2013), 1304.7721.

[20] A. Drukier and L. Stodolsky, Phys. Rev. **D30**, 2295 (1984), [,395(1984)].

[21] A. J. Anderson, J. M. Conrad, E. Figueroa-Feliciano, C. Ignarra, G. Karagiorgi, K. Scholberg, M. H. Shaevitz, and J. Spitz, Phys. Rev. **D86**, 013004 (2012), 1201.3805.

[22] H. T. Wong, J. Li, and Z. Y. Zhou (TEXONO), in *1st Yamada Symposium on Neutrinos and Dark Matter in Physics (YS-1 and NDM03) Nara, Japan, June 9-14, 2003* (2003), hep-ex/0307001.

[23] N. Mirabolfathi, H. R. Harris, R. Mahapatra, K. Sundqvist, A. Jastram, B. Serfass, D. Faiez, and B. Sadoulet (2015), 1510.00999.

[24] R. Agnese et al. (SuperCDMS), Phys. Rev. Lett. **116**, 071301 (2016), 1509.02448.

[25] J. R. T. de Mello Neto et al. (DAMIC), in *Proceedings, 34th International Cosmic Ray Conference (ICRC 2015)* (2015), 1510.02126, URL `http://inspirehep.net/record/1396680/files/arXiv:1510.02126.pdf`.

[26] C. E. Aalseth et al. (CoGeNT), Phys. Rev. **D88**, 012002 (2013), 1208.5737.

[27] B. Dutta, R. Mahapatra, L. E. Strigari, and J. W. Walker, Phys. Rev. **D93**, 013015 (2016), 1508.07981.

[28] D. Z. Freedman, Phys. Rev. **D9**, 1389 (1974).

[29] J. Barranco, O. G. Miranda, and T. I. Rashba, JHEP **12**, 021 (2005), hep-ph/0508299.

[30] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch, Phys. Lett. **B160**, 325 (1985).

[31] V. I. Kopeikin, Phys. Atom. Nucl. **75**, 143 (2012), [Yad. Fiz.75N2,165(2012)].

[32] J. Lindhard, V. Nielsen, M. Scharff, and P. Thomsen, Mat. Fys. Medd. Dan. Vid. Selsk **33**, 10 (1963).

[33] A. K. Jastram, Ph.D. thesis, Texas A&M University (2015), URL `http://hdl.handle.net/1969.1/156438`.

[34] D. Pelowitz et al., *MCNPX Users Manual Version 2.7.0* (2011), LA-CP-11-00438.

[35] S. Agostinelli et al. (GEANT4), Nucl. Instrum. Meth. **A506**, 250 (2003).

[36] G. Agnolet et al. (MINER) (2016), 1609.02066.

[37] T. Goorley et al., Nucl. Tech. **180**, 298 (2012).

[38] J. Cao, Nucl. Phys. Proc. Suppl. **229-232**, 205 (2012), 1101.2266.

[39] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz, JHEP **05**, 050 (2013), 1303.3011.

[40] J. Billard, L. Strigari, and E. Figueroa-Feliciano, Phys. Rev. **D91**, 095023 (2015), 1409.0050.