Feasibility of a Nuclear Recoil Observatory with Directional Sensitivity to WIMPs and Solar Neutrinos

B. Simpson[1]

Abstract

Now that conventional WIMP dark matter searches are approaching the neutrino floor, there has been a resurgence of interest in the possibility of introducing recoil direction sensitivity into the field. Such directional sensitivity would offer the powerful prospect of reaching below this floor, introducing both the possibility of identifying a clear signature for dark matter particles in the galaxy below this level but also of exploiting observation of coherent neutrino scattering from the Sun and other sources with directional sensitivity. We survey the experimental status of all technologies proposed to date, and perform a cost-benefit analysis to identify the optimal choice in different WIMP and neutrino scenarios. Based on our findings, we propose a large-scale directional nuclear recoil observatory with directional WIMP sensitivity below the neutrino floor and capability to explore Solar neutrino coherent scattering with direction sensitivity

Keywords: keyword1, keyword2

Contents

1. Introduction

[Section organizer: all]

The aim of this paper is to lay out the science case and goals for a large galactic recoil observatory, to show that the goals are achievable in principle, to compare the capability of different technologies to reach those goals. Here is a citation [**?**].

2. Science Case for a large Nuclear Recoil Observatory

[Section organizer: Katie Mack]

2.1. WIMP Scattering

2.1.1. WIMP scattering review

[KM]

- Detection overview

  The primary method of direct detection is nuclear recoil, in which the aim is the detection of the momentum transfer from a dark matter particle to a target nucleus. Experiments have been carried out using a wide range of targets, with recoil detection via charge, light, or heat (phonon) signals. Direct detection experiments have produced limits on the properties of WIMPs in the parameter space of mass and WIMP-nucleon cross section, with the tightest limits challenging favored supersymmetric WIMP models and reaching thermal production cross sections. However, several experiments have also reported detections that may be consistent with dark matter interpretations, while being inconsistent with existing limits. One notable example is the DAMA/LIBRA collaboration, which has reported a signal in annual modulation over 14 years and at a signal significance of 9.3σ. As DAMA/LIBRA is unique in using a NaI crystal target, efforts are being made to reproduce the experiment in the Southern Hemisphere to rule out target-specific effects and to eliminate seasonal variations as an explanation for the effect.

  Meanwhile, directional detection presents a new opportunity for discovery in this space. With directional capability, detectors have a strongly enhanced ability to remove backgrounds, through the reliance on the expectation that the WIMP wind should originate primarily from roughly the direction of the constellation Cygnus, due to the direction of the motion of the Sun through the Galactic WIMP halo. Directional capability will make potential WIMP detection more reliable and robust through (1) confirmation of the connection between the events and the Galactic halo, and (2) elimination of backgrounds associated with solar neutrinos at low interaction cross sections (which come from the direction of the Sun) and with backgrounds from the detector's immediate surroundings (which will not correlate with the direction of Cygnus).

  Ongoing directional detection experiments such as DRIFT-II have provided a proof of concept and upper limits. Our proposed detection method will provide the opportunity to strongly improve these limits and potentially to detect the Galactic dark matter, and to study the structure of the halo. In addition, it will have the prospect to study unique phenomenology through sensitivity to coherent neutrino scattering, which presents a "wall" in detection space for non-directional experiments [**?**].

  We present below a brief summary of current limits on WIMP properties from non-directional and directional detectors, as well as prospects for unique discovery with directional detectors.

- Current limits

  Current limits on WIMP scattering from direct detection experiments are generally expressed in the parameter space of the WIMP-nucleon scattering cross section and the WIMP mass, where the detection threshold depends on the mass of the target nuclei and the energy threshold of the detector's sensitivity to nuclear recoils. In recent years, several experiments have produced signals consistent with WIMP recoil events, but the majority of detection efforts have produced lower limits, and there are presently no candidate detections that are consistent with the results of all experiments.

The event rate for nuclear recoils is given by:

$$\frac{\mathrm{d}R}{\mathrm{d}E}(E,t) = \frac{\rho_0}{m_\chi m_A}\int v f(\mathbf{v},t)\frac{\mathrm{d}\sigma}{\mathrm{d}E}(E,v)\mathrm{d}^3 v \tag{1}$$

where ρ_0 is the local dark matter mass density, m_χ is the dark matter particle mass, m_A is the nucleus mass, v is the dark matter velocity in the detector rest frame, $f(\mathbf{v},t)$ is the velocity distribution, and the derivative is the differential cross section, which can be written as:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E} = \frac{m_A}{2\mu^2 v^2}\left(\sigma_0^{SI}F_{SI}^2(E) + \sigma_0^{SD}F_{SD}^2(E)\right). \tag{2}$$

Here, the first term includes the spin-independent cross section and form factor and the second includes the spin-dependent cross section and form factor. The factor μ_A is the WIMP-nucleus reduced mass. Constraints on the spin-independent and spin-dependent cross sections vary, as they employ different assumptions about the interactions between the WIMP and the nucleon. In the spin-dependent case, the interaction probability is not amplified for heavier target nuclei, and the constraints tend to be weaker.

The strongest constraints available apply to the spin-independent cross section. Figure 1, from [**?**], shows a selection of constraints from direct detection experiments, along with the allowed detection regions due to results from the DAMA/LIBRA experiment and CDMS-Si.

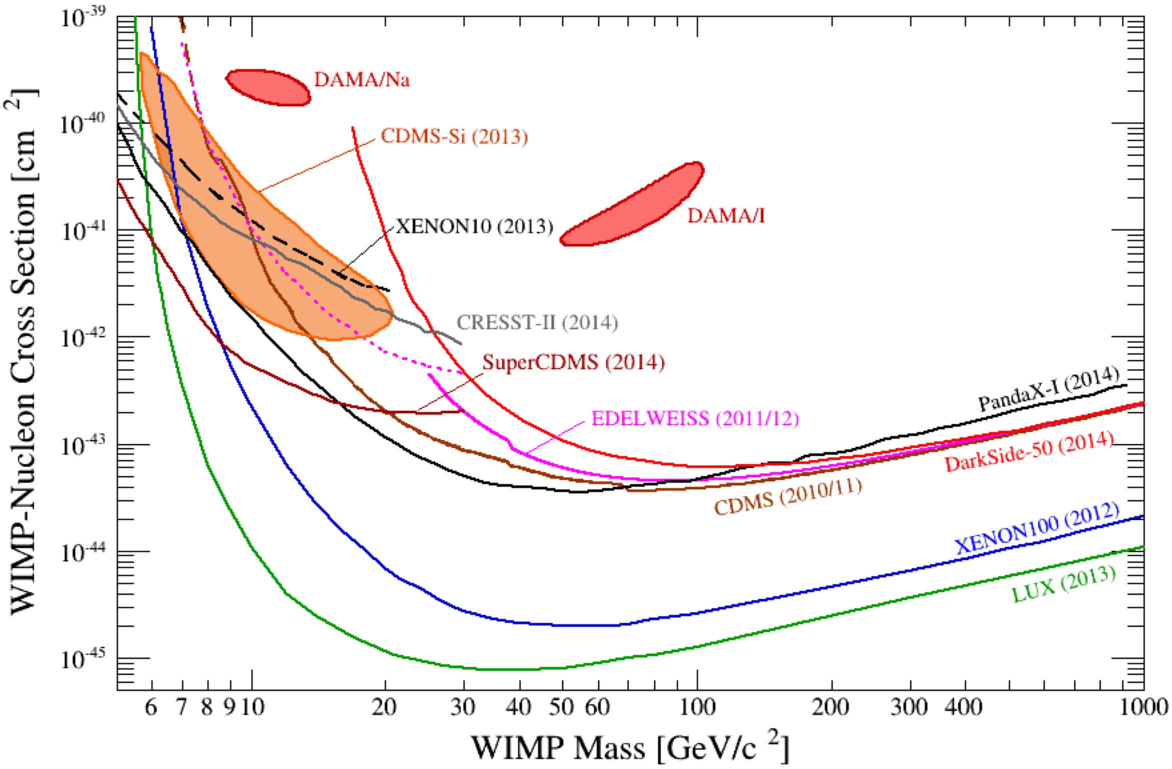

Figure 1: Constraints on the spin-independent cross section, from [**?**]. Constraints and detection regions for different experiments are labelled in the plot. The detection regions shown correspond to the DAMA/LIBRA experiment and CDMS-Si.

2.1.2. Galactic signal detection below the neutrino floor

It was anticipated in early work on direct dark matter detection that large detectors would eventually become sensitive to coherent scattering between neutrinos and nuclei [**?**]. For the keV nuclear recoil energy scales observed in direct detection experiments Solar, diffuse supernovae and atmospheric neutrinos all constitute a significant background for detector exposures beyond the ton-year scale [**? ? ?**]. Because neutrinos are impossible to shield they represent the ultimate background for the direct detection of WIMPs. Moreover, because the nuclear recoil energy spectra induced by coherent neutrino-nucleus scattering mimics the spectra for WIMPs of certain masses, the discovery of these characteristic masses is limited due to the sizable systematic uncertainty on the expected neutrino flux. The range of cross sections that are reached by an experiment that has sufficient sensitivity to be also subject to a dominating neutrino background is known as the "neutrino floor" [**?**]. The shape of the neutrino floor is dependent on the flux of each neutrino background component as well as, importantly, the uncertainty on this flux. The most notable and threatening feature in the neutrino floor is the shoulder just below WIMP masses of ~10 GeV due to the large flux and low energies of Solar neutrinos. The most important of these are the neutrinos originating from ^{8}B decay. In a xenon experiment the nuclear recoil signal due to a 6 GeV WIMP with a SI cross section around 5×10^{-45} cm^2 is well matched by ^{8}B neutrinos. Towards slightly larger masses (10 - 30 GeV) the neutrino floor is set by the diffuse supernova neutrino background (DSNB) due to the cumulative emission of neutrinos from a cosmological history of supernovae. The expected flux of the DSNB is extremely low ($\sim$ 80 cm^2 s^{-1} [**?**]) so the neutrino floor at these intermediate masses falls by several orders of magnitude in cross section. Towards masses beyond 100 GeV the neutrino floor is induced by the low energy tail of atmospheric neutrinos from cosmic ray interactions in the upper atmosphere. Atmospheric neutrinos are the only significant background contributing neutrino energies above 100 MeV. The low energy tail of atmospheric neutrinos is difficult to both measure and theoretically predict [**?**] so currently has uncertainties of around 20% [**?**].

A central challenge for the next generation of dark matter experiment is how to continue the search for dark matter WIMPs to cross sections below the neutrino floor. However, it is important to emphasise that despite the nomenclature, the neutrino floor is not a hard limit to direct detection. This is because the neutrino background is not strictly irreducible, even in conventional experiments. While the nuclear recoil energies of coherent neutrino-nucleus scattering and WIMP-nucleus scattering are very similar, the spectra do not exhibit perfect matching, even for masses best mimicked by neutrinos. As initially shown by Ruppin *et al* [**?**], the neutrino background can be subtracted with recoil energy information alone for very high statistics due to the slight differences in the tails of the recoil energy distributions. Unfortunately this requires prohibitively large experimental exposures, usually in excess of 1000 ton-years. It has also been shown that for some of the additional operators posited in the non-relativistic effective field theory formalism, the recoil spectra are sufficiently distinct from neutrinos to allow their discrimination with fewer events than in the standard SI or SD cases [**? ?**]. However the overlap between the WIMP signal and neutrino background spectra is worsened - independent of particle physics - once astrophysical uncertainties are taken into account [**?**].

Given that the next generation of ton-scale experiment is expected to become sensitive to coherent neutrino-nucleus scattering, it is pertinent to search for alternative and more powerful methods of subtracting the background. The most basic approach to alleviate the background is to exploit the complementarity between target nuclei of differing masses and nuclear content. For the SI neutrino floor it has been shown that this approach only leads to a marginal improvement in alleviating the neutrino background, however in the case of SD interactions the differences in nuclear spin contents make complementarity a more viable strategy [**?**]. It was also shown by Davis [**?**] that the use of timing information also allows the low mass neutrino floor to be overcome with slightly lower statistics. This approach exploits both the annual modulation of the dark matter signal due to the relative Galactic motion of the Earth and the Sun, as well as the annual modulation in the Solar neutrino flux due to the eccentricity of the Earth's orbit.

Directionality presents by far the most attractive prospect for circumventing the neutrino floor because the unique angular signatures of both dark matter and Solar neutrinos allows optimum discrimination between signal and background. This was first shown in Ref. [**?**] in the context of conventional low pressure gas TPCs and in Ref. [**?**] for experiments using a range of readout strategies. The effect of directional information has also been explored in ideas using nuclear emulsions [**?**] and spin-polarised helium-3 [**?**]. The general consensus is that in an idealised directional experiment there is effectively *no* neutrino floor. The crucial factor that enables this is that over the course of the year the Sun does not pass through the constellation of Cygnus. The angular distance between Cygnus and the Sun undergoes a sinusoidal modulation which peaks in September at around 120° and is a minimum during March at around $\sim 60°$. Because Solar neutrino recoils can only point with angles less than 90° from the Solar direction, this implies that over long periods during a year there are large WIMP signal regions across the sky where it is guaranteed that the number of Solar neutrino events is zero (ignoring the effects of angular resolution). On the other hand, the advantage of directional detectors for dealing with the neutrino floor at higher masses is not as significant due to the greater angular dispersion exhibited by the remaining backgrounds [?]. The directionality of non-Solar neutrinos is much less well understood. Whilst the DSNB is certainly expected to be isotropic, one would not naively expect the same to be true for atmospheric neutrinos. Indeed FLUKA simulations of low energy neutrinos have shown an enhancement in the flux towards the horizon [? ?]. However, for the atmospheric neutrino floor this phenomenon turns out to be unimportant in part because the directionality is fixed in the reference frame of the detector, as opposed to the Galactic signal which transits across the sky over a sidereal day. Additionally the coherent scattering process acts to wash out much of the horizontal directional preference meaning the recoil sky due to atmospheric neutrinos is also very close to isotropic in appearance. So directionality is much less powerful at circumventing the neutrino floor beyond 100 GeV, however an ideal directional detector can still out-perform an equivalent conventional experiment by a factor of a few.

Reference [?] also demonstrates the effect of a range of experimental limitations suffered by directional detectors, namely lower dimensional readout strategies, head-tail sense recognition and angular resolution. It was found that at low masses the projection of recoil tracks onto 1- or 2-dimensional planes, whilst slightly harming the progression of discovery limits past the neutrino floor (compared with full 3-d information), was not the most limiting factor. Head-tail recognition on the other hand has been shown persistently to be a major limitation for directional detection in general. This is true also for subtracting neutrino events, especially when combined with lower dimensional readout strategies. For a 1-d experiment with no sense recognition, the discovery limits made under the neutrino background are only marginally better than with no directional information at all. Furthermore, comparison between various strategies found that having a fully 3-d experiment without head-tail is slightly better than a 2-d experiment with head-tail. This is only true for Solar neutrinos however, for the high mass neutrino floor it was found that the lack of sense recognition is disastrous, independent of the dimensionality of the recoil direction measurements. This is because of the significant increase in overlap between WIMP and neutrino event rates after the folding of the forward and backward directions. Finally for the question of angular resolution, it has been argued that a resolution better than 30° should suffice to achieve an order of magnitude improvement over a non-directional experiment when the Solar neutrino background is considered. The restriction imposed by angular resolution is less important when considering higher enegry neutrinos, but in this case the benefit offered by directionality is less pronounced overall.

2.1.3. WIMP astrophysics

[KM]

A wide range of observations across Galactic to cosmological scales present strong evidence for the existence of dark matter as an unseen component of the Universe and a dominant contribution to the mass budget of galaxies, clusters, and the cosmic web. From measurements of the gravitational potential within the Milky Way, we can infer the distribution of dark matter locally and begin to reconstruct the full dark matter halo. While estimates of the local density of dark matter (within a few kiloparsecs of the Sun) have converged around a value of $\rho \approx 0.008 M_{\odot}\mathrm{pc}^{-3}$ (see, e.g., [? ?]), there is still a great deal of uncertainty surrounding the velocity distribution of the dark matter [? ?], which impacts the direct detection rate via the distribution function $f(\mathbf{v}, t)$, as well as the motion of the Sun through the halo [?]. These uncertainties impact the reliability of signal modelling, but may also present an opportunity for discovery via directional dark matter detection where other methods have limited power. For example, it has been shown that directional detection can place constraints on the velocity structure of the halo [? ? ?], which can give insight into its structure formation history, as well as dark matter's fundamental properties. Similarly, simulation work has suggested that direct detection has the potential to illuminate dark matter streams [?], debris flows [?], and other hidden features of the dark matter halo. While very massive streams can be indirectly seen via studies of disrupted stellar systems, small streams and low-mass dark matter halos are likely to have undetectable levels of influence on luminous matter; an opportunity to better understand the level of clumpiness of the Galaxy's dark matter halo is a key potential advantage of directional detection. Any insight we may gain into the small-scale structure of dark matter halos has the potential to produce hints of non-vanilla dark matter models (such as warm dark matter, self-interacting dark matter, etc) and to illuminate structure formation processes. It also has the potential to be fully unique probe, as there are currently very few observational handles on the small-scale structure and mass function of dark matter.

The main limiting factors in the ability of directional detectors to constrain WIMP astrophysics are the precision obtained on the incoming direction of particles and the energy distribution of recoils [?]; we discuss the experimental prospects for direction and energy precision in Sections 3 and 4.

2.1.4. *Particle models and directionality*

[KM, COH]

- Advantages of directionality for distinguishing models [[Discuss: warm dark matter, self-interacting dark matter; see papers by Peter etc.]]
- Inelastic dark matter models can give rise to enhanced signal discrimination power in directional detectors [**?**]. This is because the recoils are more focused in the forward direction because slower WIMPs can't scatter with enough energy to induce an excited state. Directional detectors can also disentangle elastic and inelastic events in models which allow for both [**?**].
- The non-relativistic effective field theory framework proposes a set of additional non-standard operators that include all Hermitian, Galilean and rotation-invariant interactions constructed out of the low energy degrees of freedom involved in the DM-nucleus interaction [**?**]. Certain examples, those which are dependent on the transverse velocity of the interaction, give rise to unique ring-like angular signatures [**? ?**]. This means that directional detectors are potentially more powerful than conventional experiments in distinguishing between these particular operators.
- If dark matter exists in the form of 'darkonium' bound states composed of two or more particles (as is predicted in some configurations of asymmetric dark matter models) it has been shown that there may be angular signatures observable in directional experiments that may constrain their properties [**?**].

2.2. *Solar Neutrino Coherent Scattering*

[PB] The coherent scattering of neutrinos off nuclei was predicted over 40 years ago with the realization of the neutral weak currents [**?**]. This standard model process remains unobserved due to daunting detection requirements: ~keV nuclear recoil thresholds, kilogram to ton-scale target masses, and low backgrounds. Due to the small weak charge of the proton, the coherence results in an enhanced neutrino-nucleon cross-section that is approximately proportional to the square of the number of neutrons in the nucleus. A few years after the coherent neutrino scattering prediction, and, ironically, before the conception of the first dark matter direct detection experiments, the possibility of using this enhanced process to develop a "neutrino observatory" was put forward [**?**]. A cornucopia of physics searches were envisioned using neutrinos from stopped-pion beams, reactor neutrinos, supernova, solar neutrinos and even neutrinos of a geological origin.

Shortly thereafter, the first generation of dark matter experiments began to search for the scattering of WIMPs of their detectors, where the signature was a low-energy nuclear recoil. These experiments have dramatically improved their sensitivities over the last three decades by simultaneously increasing the target masses, as well as reducing background nuclear recoils. Today the irony lies with the fact that the unshieldable recoils that result from coherent neutrino scattering will soon be a source of background for the next generation of dark matter direct detection experiments [**?**][**?**][**?**][**?**]. Without the ability to separate the neutrino recoils, the progress in WIMP detection sensitivity will be halted. On the other hand, an experiment that can successfully separate and identify these neutrino events can not only proceed past the so-called "neutrino floor", but can also realize the long-awaited vision of a "neutrino observatory". A detector with directional sensitivity has the potential to do just that.

2.2.1. *Solar neutrino scattering review*

In the coherent neutrino scattering process, coherence is only satisfied when the initial and final states of the nucleus are identical, limiting this enhancement to neutral current scattering. The coherence condition, where the neutrino scatters off all nucleons in a nucleus in phase, is also only maintained when the wavelength of the momentum transfer is larger than that size of the target nucleus. Full coherence for all scatters is only guaranteed for low energy neutrinos – less than 10's MeV, depending on the target size. The standard model total cross section for the process can be approximate (neglecting neglecting axial vector terms that arise from unpaired nucleons):

$$\sigma = \frac{G_F^2}{4\pi}\left[Z(4\sin^2\theta_W - 1) + N^2\right]E_\nu^2|F(q)|^2 \tag{3}$$

Where G_F is the Fermi constant, Z is the number of protons, N is the number of neutrons, θ_W is the Weinberg angle, q is the momentum transfer, E_ν is the energy of the nucleus and θ is the scattering in the lab frame. It is evident that the cross section also increases with the square of the energy of the neutrinos; however, while the form-factor condition–which comes in as $|F(q)|^2$–is easily satisfied for Solar neutrinos, the total cross section begins to suffer from decoherence with supernova neutrinos, and neutrinos from stopped pion beams. A detector with an energy threshold of zero can expect to see several hundred to a few thousand recoils from solar neutrinos per ton-year of exposure, depending on the target mass [**?**].

The differential cross-section with recoil energy can be approximated as:

$$\frac{d\sigma}{dE_{rec}} = \frac{G_F^2}{8\pi}\left[Z(4\sin^2\theta_W - 1) + N^2\right] M\left(2 - \frac{E_{rec}M}{E_\nu^2}\right) \quad (4)$$

Where E_{rec} is the recoil energy of the target nucleus, and M is the mass of the target nucleus.

A more realistic scenario for estimating count rates can be made assuming a ^{19}F target, for example, and a 5 (10) keV threshold for observing nuclear recoils. This results in an expectation of ∼90 (15) background recoils per ton-year, from solar neutrinos alone [**?**].

2.2.2. Advantages of directional detection

It is possible to alleviate the constraints that these solar neutrino recoils place on any dark matter search by taking advantage of the expected directional response recoils due to both the putative WIMPs, as well as those from solar neutrinos. The coherent neutrino scattering differential cross section with respect to the recoil angle can be written as:

$$\frac{d\sigma}{d(\cos\theta)} = \frac{G_F^2}{8\pi}\left[Z(4\sin^2\theta_W - 1) + N^2\right] E_\nu^2(1 + \cos\theta) \quad (5)$$

The resulting recoils are thus biased to the forward direction, away from the location of the Sun. As the solar position changes diurnally with respect to the expected direction of the WIMP wind, an analysis of the recoil direction of events in the detector should reduce the impact of this background. A similar separation could be imagined for terrestrial, atmospheric and diffuse galactic supernova neutrino backgrounds–each with their own characteristic directionality and energy scale.

2.2.3. Science with source and detector

2.3. Other Physics

2.3.1. Non-solar neutrinos

[KS] Neutrinos with energies less than a few tens of MeV [anything else besides supernova and solar? Low-energy atmospheric... a section on stopped-pion nus?]

- Supernova neutrinos: A core-collapse supernova will emit an enormous fluence of neutrinos over a few tens of seconds time scale. The neutrinos in the burst will have a few to a few tens of MeV of energy, and will include all flavors of neutrinos and antineutrinos with roughly equal luminosity [**?**].

  Dark-matter detectors with very low recoil energy thresholds are sensitive to a supernova neutrino burst via coherent elastic neutrino-nucleus scattering. The order of magnitude is a few events per ton of detector material for a supernova at ∼10 kpc (near the most likely distance to the supernova [**?**]), and statistics will scale linearly with detector mass and as the inverse square of distance to the supernova. Such a detection would be valuable due to its sensivity to the entire flux, given that most other detectors online are sensitive primarily to the $\bar{\nu}_e$ (in water, scintillator detectors) and ν_e (in argon, lead detectors) components of the flux [**?**]. Furthermore, some neutrino spectral information can be reconstructed from the measured nuclear recoil spectrum.

  The advantages of directionality for the detection of supernova burst neutrinos via CEvNS are several: first, obviously, directional information about the source will be of value to observers in electromagnetic wavelengths and in gravitational waves who want to make prompt observations of the supernova event in real time. Currently, only detectors able to make directional measurements of elastic scattering on electrons have good pointing ability (and Super-K is the only current instance). Even if there is no obviously bright supernova event (as may be the case for a failed supernova), directional information will be able to narrow down the possible progenitors. Finally, the direction information can be used on an event-by-event basis to reconstruct an more precise neutrino energy.

2.3.2. Axions and exotic models

[PB, KM, COH]

[COH] some possible ideas:

- Solar neutrino physics e.g. Refs. [**?** **?**].
- Detecting/pointing Galactic SN as in Ref. [**?**].
- Beyond the SM neutrino interactions [**?** **?**].
- Diffuse supernova neutrino background as a test of cosmology [**?** **?**].
- Atmospheric neutrinos, low energies still hard to measure [**?**].
- Detecting dark matter and Solar axions, if axioelectric effect can be demonstrated in gaseous target [**?** **?**]

3. Existing Directional Detection Technologies

[Section organizer: James Battat]

Contributors to this section:

- James Battat jbattat@wellesley.edu
- Elisabetta Baracchini baracch@gmail.com (esp. with "emerging tech." such as columnar recombination, nanotubes, anisotropic scintillators, DNA, etc.)

Directional detection can be achived by a direct reconstruction of the nuclear recoil geometry (e.g. by building a tracking detector), or by an indirect proxy for the recoil direction (*e.g.* a detector whose response depends on the relative alignment of the recoil and the detector axes). A detailed and critical assessment of directional readout technologies is provided in Ref. [**?**].

3.1. Detectors that reconstruct the recoil track

The currently active directional experiments all aim to reconstruct the geometry of the recoil track. Of these, most make use of a low-pressure gas Time Projection Chamber (TPC), in which the track geometry is measured in 1D or 2D or 3D. In addition to gas-based TPCs, track reconstruction at the FIXME sub-millimeter scale has been demonstrated in solid emulsions. More exotic and at this point unvalidated technologies such as a customized matrix of DNA strands have been proposed as well.

3.1.1. Gas-based TPCs

james will do this

- Negative ion drift vs. Electron drift
- amplification device may be integral to readout (micromegas, MWPC) or separate (GEM)
- MWPC
- MPGD (micromegas, mupic, pixel chip)
- Optical

3.1.2. Nuclear Emulsions

James will populate this

[?]

see also EB's excerpt.

3.1.3. DNA strand detector

A highly novel recoil tracking detector makes use of customized DNA or RNA strands mounted in a matrix onto a nanometer-thick gold foil [?]. A WIMP would interact with and kick out a gold atom from the foil, and the recoiling gold atom would sever several DNA strands. Using well-established biological techniques such as polymerase chain reaction (PCR) and sequencing, it would be possible to identify the (x, y, z) coordinate of each severing event, thereby reconstructing the nuclear recoil axis (though not the vector direction). Originally proposed in 2012, there are no published experimental demonstrations of this technology.

3.1.4. Planar targets (graphene)

Nuclear recoils in a 3D (bulk) target suffer multiple interactions with the surrounding medium that scramble the recoil direction. In principle, the recoil direction can be more directly measured if the target is planar. Furthermore, planar targes can be fabricated from semiconductor materials in which the exictation energy is on the order of $\sim$ 1 eV, allowing even MeV-scale WIMPs to initiate electronic excitations. A recent proposal [?] suggests that 2D graphene could serve as a directional detector of sub-GeV WIMPs. This is a particularly interesting idea, especially given that no other directional technology can probe this WIMP mass scale. Although there has not been an experimental demonstration of this technology, it may be possible to do so within the PTOLEMY experiment (a relic neutrino search) [?].

3.2. Detectors that indirectly determine the recoil direction

3.2.1. Anisotropic scintillators

Solid scintillators (*e.g.* NaI and CsI) are commonly used in particle detection, and specifically in dark matter detection. Because of their large target mass and high-*A* content, they are particularly interesting for spin-independent WIMP searches. Some scintillators, such as $ZnWO_4$ and stilbene have been shown to exhibit a response that depends on the recoil ion direction relative to the crystal axes. In principle, this scintillation anisotropy can be used to infer the nuclear recoil track direction without direct reconstruction of the track geometry. Several groups have explored the possibility of using anisotropic scintillators for a directional dark matter search [? ? ? ? ? ?], though the magnitude of the anisotropy is too small for a sensitive directional WIMP search.

It is important to notice that none of this have yet proven anisotropic scintillation for low energy nuclear recoils. Therefore, all the quoted energy resolution, threshold and general performances are for general detection of alpha, beta and gamma radiation and not necessarily valid for nuclear recoils.

3.2.2. Columnar recombination

When heavy tracks ionize a medium, a column of electrons and ions gets created along the track direction. If no electric field is present, these particles will recombine producing a scintillation light signal. Since recombination probability depends on the proximity of electrons and ionized atoms, if an external electric field is applied, the amount of light produced will depend on the relative orientation of the field with respect to the ionizing track. A large angle, in fact, will lead electrons transversely away from the ions, generating a small recombination scintillating signal (R), while a small angle will bring electrons and ions closer together and produce a relative enhancement of the R signal with respect to the ionization signal (I). A precise measurement of the R/I ratio (charge/light) could therefore be used to ?sense? the directionality of the track without actually seeing it [?]. Since the direction is inferred from this ratio that is produced prior to the drift, all the limitations imposed by the degrading effects of diffusion, avalanche gain and reconstruction noise would be effectively largely reduced, possibly allowing the construction of large monolithic Xenon gas TPC at the ton-scale. With the Xe density at 10 bar being 0.05 gr/cm^3, a 1-ton detector could be realized with only 20 m^3.

Evidence for columnar recombination in alpha tracks was observed in dense Xenon [?], so the question still to be answered is if this can be seen for the much shorter nuclear recoils. Recent simulations [?] confirm how, with the proper cooling of the ionized electrons, the recombination probability should show directional sensitivity for track longer than about 2 μm in gaseous Xe at 10 bars, implying about 30 keV energy threshold. The main issue is to keep electrons thermalized near the ions in order to recombine efficiently, but unfortunately pure Xe do not satisfy this requirement due to the lack of inelastic scattering below 7 eV. This is the reason why the only only published work on the subject employed Trimethylamine (TMA) as dopant, because of its large inelastic cross-section, its UV-quenching properties and the possibility of exploiting a Penning effect. The transformation of the Xe^{+} image into the TMA^{+} molecular image and the columnar recombination happening on TMA^{+} ions would then provide light around 300 nm, a much more PMT-friendly light than the Xe emission spectrum. Unfortunately, despite enhancement of recombination with TMA was observed, no sign of scintillation light from recombination was detected and TMA was found to highly absorb Xe light without re-emitting it [**?**]. The use of alternative dopants, possibly generating negative ions drift, has recently been suggested but not yet tested.

While columnar recombination is intrisically sensitive to the axial track direction but not to its sense, the combination of two detector with the drift fields anti-aligned could be able to show head-tail sensitivity in a statistical way. The final performances of an experiment based on this technique, in terms of energy threshold and resolution, directionality performances and efficiencies, will highly depend on the readouts chosen to detect the light and charge produced in the process, and is therefore beyond the possibility of evaluation as for today.

3.2.3. Carbon nanotubes

Single wall aligned carbon nanotubes (CNTs) have been recently proposed as a DM target due to their an anisotropic response to neutral particles [**?**]. When a C ion gets scattered off the CNTs walls, in fact, if the right initial conditions are met, it sees the tube as empty and can travel with nearly no loss of energy (i.e. channeling). Numerical simulations have confirmed that different orientation of the CNTs axis with respect to the Cygnus constellation would give sensibly different channeling probabilities and therefore produce significantly different C ions current at the end of the nanotube.

The proposed detector concept by [**?**] is a brush of CNTs array closed at one end and opened at the other, inserted in a (low-pressure) TPC to detect the outgoing C ions down to low $\sim$ 10 keV. An R&D effort is currently on-going in Italy to test the channeling hypothesis for neutral particle scattering and the TPC detector approach. If these were proven successful, then an experiment based on this technique would profit from the higher density of CNTs (seems possible to reach about 10 kg on 100 thin stacked CNTs panel of 1×1 m^2 each) and possess about the same performances of a gaseous TPC-based approach, depending on the chosen readout. Other possible detector configuration (with solid target to detect the outcoming C ions, for example) could also be considered and would show significantly different performances.

3.3. Summary table

Requires a bit of thought...

- Energy resolution demonstrated.
- Axial reconstruction demonstrated? Down to what energy? With what angular resolution?
- Sense-recognition demonstrated? Down to what energy?
- Full-volume fiducialization demonstrated?
- Flexibility for different targets (mostly for gas-based TPCs – e.g. different gases, negative ion vs. electron drift)
- Technological readiness (including largest volume in operation, prospects for scaling up, some mention of cost per something (e.g. volume, or area, or ....)
- Background discrimination? This would potentially be a rather hard item to cover... gamma/recoil separation studies by Loomba et al.

  NEWAGE: 1e-6 at 50keVee at 100 Torr (CF4)

  DRIFT: 2e-7 above 30keVr at 41 Torr (CS2:CF4:O2)

MIMAC: ???

Pixel chips: ???

Dinesh: ???

- background level studies of material / components - perhaps already covered by Neil?

4. Quantitative Comparison of Directional WIMP and Solar Neutrino Sensitivity

[Section organizer: Sven Vahsen]
Contributors to this section:

- Sven Vahsen, sevahsen@hawaii.edu
- Kentaro Miuchi, miuchi@phys.sci.kobe-u.ac.jp
- Cosmin Deaconu, cozzyd@kicp.uchicago.edu
- Anne Green, anne.green@nottingham.ac.uk
- Ciaran O'Hare, ppxco@nottingham.ac.uk

This section compares the sensitivity of different simplified technologies to key science goals. This is done by developing a figure of merit that also considers cost, but assumes zero background is achievable. "Simplified technologies" means the comparison will be of idealized technologies; for instance wires versus strips versus 2D pixels (optical) versus 3D pixels, not a comparison of existing experiments. There will probably be a slight focus on TPC readouts in this section, as TPCs currently are the most studied in terms of performance, and furthest along in terms of technological readiness. This choice of focus means that the comparison becomes more realistic. That being said, comparing TPCs with the other approaches is also important, and should be included. We are still thinking about the how to do it in detail.

4.1. Simulation of nuclear recoils

This section describes the part of the simulation that takes a recoil momentum vector as input, and creates an ionization distribution in the detector. By default we use SRIM to create the ionization distribution, and create a uniform spatial distribution.Sven has code that builds PDFs based on SRIM, but his code ignores straggling. Cosmin has similar code that also includes straggling, and is willing to contribute. Kentaro also has code.

The detector target (low-pressure gas mixture, noble liquid, or emulsions) is part of the the SRIM simulation, so it should also be discussed here. I imagine that for each target chosen, we show plots of:

- SRIM range vs energy
- SRIM quenching factor vs energy
- SRIM stragling vs energy

A starting point for gas mixtures to simulate could be: CF_4, SF_6, CF_4+SF_2, SF_6+He

4.2. Simulation of detectors and readouts

This section describes simulation of the charge propagation and readout. Charge propagation can be described by Gaussian smearing for diffusion and ion attachment (?). Readout simulations will be simplified. Both Sven and Kentaro have existing code for this. We will consult the relevant experts for each readout type, as listed in parentheses below, and fill in the specs.

- TPC w/ MWPCs wires [D. Snowden-Ifft],
- TPC w/ strip readout [James]

- TPC w/ optical readout [James/Cosmin]
- TPC w/ mu-pic readout [Kentaro],
- TPC w/ GEMs and pixels [Sven, Elizabetta]. Proposed specs to use: 50 x 50 micron pixels. 50 micron equivalent time-binning for TOA. 100 electron noise. 2000 electron threshold. Dynamic range: ?
- emulsions
- TPC w/ noble liquid target

4.3. Algorithms to extract directional signals

Because the above detectors measure very different quantities, different algorithms are needed to reconstruct a directional signal. (For instance, wires don't reconstruct tracks.) How can we ensure that the comparison is done *fairly*, i.e. that we are optimally exploiting the directionality of each detector? Let proponents of detector propose method? It would be nice to keep this part as simple as possible. We can probably adopt what has been done in the papers of Green, O'hare, Billard et al. This could be a good place to get Billard and O'hare to help.

4.4. Directional power of detectors versus recoil energy

The final detector comparison is sensitive to astrophysics, gas optimization, detector performance, and cost. To decouple these effects, we here start out by quantifying directional performance versus recoil energy. This is done by estimating how many recoils each detector needs to observe, to discriminate a delta function (all recoils go in the same 3D direction) at 5-sigma from a flat recoil distribution, versus recoil energy. This goal here is to provide an intuitive results that clearly shows the recoil energy range where each technology is effective, and how directional it is.

4.5. Directional WIMP and Solar Neutrino Sensitivity

Compared to the previous section, in this section, and the next, we now also fold in the recoil distributions for a realistic physics scenarios, target interaction probability, and cost. We explain how we compare the sensitivity of directional technologies, including how we optimize for nuisance parameters such as gas pressure. The procedure for one physics goal is explained in detail, culminating in a final publicity plot where directional detector technologies are compared against each other and against non-directional ones. Such a publicity plot is one key goal for the paper. A candidate example plot could be ”# sigma that a galactic-coordinate dipole pointing back to CYGNUS, and a flat recoil distribution can be separated, per million dollars, for a 100 GeV WIMP” [with a specified cross-section]. Non-directional detectors probably score zero on this performance metric. If cost proves too hard to pin down, then we can instead show #sigma per cubic meter for each technology, chose a TPC with wires as the default, and tabulate the required cost for other technologies to become competitive.

Ciaran has agreed to provide the recoil distributions, probably as 3-vectors + time. These will then be interfaced with Sven's/Cosmin's/Kentaro's code.

(We also have to discuss here or elsewhere any form factor or angular distribution assumptions we make for WIMP and neutrino scattering.)

4.6. Figure of Merit for Specific Science Goals

The same procedure as in the previous section is now repeated for a number of physics goals. Again, we'll ask Ciaran to generate the 3-vectors for nuclear recoils from WIMPs and neutrons. This time, only the final result (the publicity plot) is given for each physics goal.

Candidate list of physics scenarios (will be revised based on physics case chapter):

- discover DAMA/LIBRA WIMP
- discover 100 GeV WIMP above neutrino floor
- discover 1TeV WIMP above neutrino floor
- discover WIMPS below neutrino floor (1, 10, 100, 1000 GeV)
- discover WIMP streams

4.7. Conclusion on Technology Choices

Follow this with discussion of optimal technology choices: Is there a general winner that emerges? Or one winner for high, and one for low energy recoil scenarios? Is the conclusion biased by the zero background assumption? How would it change if discrimination power is included? (Can we think of an easy way to do that?)

5. Zero Background Feasibility

[Section organizer: Neil Spooner]

Generally direct search dark matter experiments strive to control backgrounds sufficiently so as to achieve an expected rate of less than 1 background event recorded in the anticipated exposure time and target mass, essentially that we have zero background within the fiducial volume. An assumption that this is achievable for all the directional technologies compared here was an important caveat made in the previous section of this work. This next section now addresses the realities of this assumption. Specifically we seek to answer the question firstly whether or not directional low pressure gas TPCs can in principle achieve such low backgrounds at the experiment scale required to reach the proposed scientific goals, but further, whether particular TPC readout technologies, with their individual associated intrinsic backgrounds and discrimination capabilities, are more or less able to reach these goals. The latter aspect depends in part on assumptions about the radio-purity of internal detector materials likely involved, most notable for instance because this affects the total internal neutron background. So an alternative tack, adopted here also, is to estimate and compare the specification on material radio-purity required for success, for instance the 238U content in each case, then to comment on the achievability of these requirements.

As shown in [ref x] the additional particle identification properties of directional detectors mean that in principle they may in actuality be able to tolerate a non-zero level of isotropic nuclear recoil background, yet still be able to identify the signal of interest here for dark matter, a non-isotropic distribution of recoil directions. However, the level of tolerance will depend strongly on the capabilities of the technology and anyway will clearly reduce sensitivity overall. A maximum signal to background ratio of order x10 might be a reasonable upper limit in certain circumstances [ref]. Nevertheless, a good starting point for comparison purposes, adopted for this work, is to assume an aim of zero background.

The following sections present results and conclusions on these issues based on new GEANT4 detector Monte Carlos and other simulations of the various key background contributions. Although not necessarily mandatory, experience from many dark matter experiments demonstrates that full fiducialisation of the active detector volume is likely necessary to achieve the background goals. This aspect is addressed in Sec 5.1. The fundamental issue of neutron backgrounds, that result in nuclear recoil events likely indistinguishable from WIMP induced events, is addressed in Sec 5.2, considering separately contributions from cosmic ray muon neutrons, and rock and detector neutrons. The subsequent parts, Sec 5.3, 5.4 and 5.5, cover respectively simulations of gamma, radon related backgrounds and other possible surface backgrounds. For the majority of the technologies some generic conclusions can be drawn based on the commonality of the basic infrastructure needed for any TPC dark matter experiment, such as a deep site, passive shielding and containment vessel. The majority of any variance from this comes from details of the internal TPC structures, notably the readout planes. These aspects are together summarized in Sec 5.6.

5.1. Fiducialization

5.2. Neutron Backgrounds

Neutrons are a major concern for all direct search experiments because they can produce nuclear recoils just like WIMPs. However, there are various issues that make the requirements for mitigating against neutron backgrounds in a low pressure gas TPC significantly different from those cases involving conventional solid or liquid based detector technologies. Firstly, the potentially low sensitivity to light charged particles, muons, muon-induced secondary particles and electrons, means that these may not be recorded. Secondly, the low density of the target means neutrons are less likely to undergo double or multiple scatters. Both these factors potentially reduce options for vetoing neutron induced nuclear recoils, depending on the readout technology chosen. The former does depend critically on the degree of position segmentation of the readout and the energy threshold achievable in those individual readout channels, essentially the sensitivity to dE/dx. The issue of vetoing by recording multiple neutron scatters then depends on the contiguous size of the detector array. For instance, at 200 torr SF6 the mean free path of a typical background neutron is 60m. This would be the sort of scale required to have any benefit from detection of multiple neutron scatters.

Factors such as these, the uniqueness of the low pressure TPC technique and potentially powerful particle identification, mean that estimating neutron backgrounds by extrapolation from existing background simulations such as have been performed for massive xenon or bolometric detectors[ ], is not appropriate. The work presented here is thus based on a set of dedicated TPC Monte Carlos. Some relevant initial work on neutron backgrounds was previously undertaken by some of the authors here but focused on smaller TPC target masses of order 1-10kg [ ]. The new work presented here makes use of the latest updated GEANT4 and SOURCES packages and specifically targets the more complex situation of neutron background mitigation in the much larger experiments required to reach the goals of CYGNUS. As noted the procedure adopted is to start by examining aspects that are independent of the internal readout technology. This includes firstly the laboratory location, determined by the depth, rock composition and cavern geometry. Secondly, the outer passive shielding and any active veto system, and finally the containment vessel, modeling both its geometry and composition. The remit here is to investigate muon-induced neutrons resulting from cosmic-rays penetrating from the Earths surface and also neutrons produced by spontaneous fission and alpha-n reactions in the rock and shielding/vessel materials. The procedure thus requires simulation of the geometry, particle production, tracking and detection, the goal being to find the rate of neutron-induced nuclear recoils anticipated in different situations. From this can be determined requirements for such issues as the amount of passive shielding, the efficiency and form of any external veto and the form and purity of the vessel materials, such as required to achieve the goal of zero background. The issue of neutrons from internal detector components, that depends on details of the readout technology, is addressed last.

5.2.1. Laboratory and TPC Geometry

In order to explore some range of possible scenarios for potential experiments we adopt here two broad geometries and underground site characteristics used for the GEANT4 simulations. Details of these scenarios are given in Table 1. The first, Option 1, assumes a laboratory akin to that of the Italian Gran Sasso facility in depth and rock composition. The second one is designed to be broadly compatible with the Boulby Underground Laboratory in the UK or the WIPP site in the US, located in salt rock. The two geometries are illustrated in Fig. 1.

In both cases the rock was simulated to a depth of 3 m outwards from the cavern walls, with the appropriate rock composition. it was assumed that the cavern volume around the detector contained 1 atmosphere of an 80:20 nitrogen to oxygen mixture and that there were no other materials or objects present in the caverns. Regarding the detector vessel dimensions, the choice obviously depends on the pressure and gas composition adopted for the experiment of which there are many possibilities. For the purposes of making broad comparisons here, bearing in mind the science goals of CYGNUS, it was decided to assume use of SF6 gas at 50 Torr with volume sufficient to produce around 0.5 tons of target nuclei, in this case of fluorine. For a facility in salt rock there are usually restrictions on the height and width feasible but not the length. So for this option an elongated vessel of 5 x 5 x 40m was chosen. Other sites in hard rock, such as Gran Sasso, do not necessarily have such a height restriction. In this case a generic size of 10 x 10 x 10m was selected. Based on these dimensions engineering studies were made to determine a minimum total mass of vessel material required in each case. For instance for Option 1 a mass of 200 tons was found necessary. Real vacuum vessels of such size will require strengthening supports both inside and outside. However, for simplicity in simulations the mass of these was taken into account by applying an appropriate average increase in thickness to the vessel walls.

As stated, the background from internal TPC components will be affected by details of the readout design, covered later. Nevertheless, some generic assumptions can be made about other TPC structures required inside the detector which are likely common to any design. Most notable here is the central cathode and field cage. However, we note that the total area required for the former will also depend on the gas mixture adopted, since this influences the diffusion and hence determines the maximum drift distance that can be tolerated. For the comparisons here we assume a compromise drift distance of 50 cm, which yields a total cathode area of 2000 m2 in both options. The design of this is assumed here to comprise ultrathin cathode sheets supported on acrylic frames, of design similar to that demonstrated previously [ref]. The field cage itself can also be made of light acrylic components, with copper strips to act as the field rings. In these components the acrylic provides by far the dominant mass, conservatively estimated to be xx tons. These components are again approximated as sheets of appropriate thickness distributed in the TPC volume (see Fig. 1). An additional important potential source of internal background is the resistor chain required to feed voltages to the field cages since these are generally composted of ceramic materials. These are modeled as a mass of xxkg of ceramic composition containing x ppb U and y ppb Th, as typically measured [ref], distributed in rectangular strips along the sides of the field cages.

5.2.2. Rock Neutrons and Passive Shielding

The first simulations were performed to determine what thickness of passive neutron shielding is required around the CYGNUS detector to ensure an induced recoil rate from this source that is below 1 per year for each of the two detector options. It is recognized that an active veto shield is also likely needed to assist with rejection of muon related neutrons (see sec 5.2.3) and that this in practice could be fully, or partially, integrated with the passive shield. To allow for this the passive shield was modeled in two forms, namely a generic hydrocarbon, able to represent either passive material or a plastic or liquid scintillator, and water, also useable in passive form as part of a veto using Cherenkov radiation. In both cases it is assumed here that any containment structure or internal components, such as photomultipliers, are of sufficiently low background and low mass to be ignored. We note also that account needs to be taken of the energy threshold chosen, as determined in part by the science priorities. To allow for this we consider cases for 1 keV and 10 keV threshold. For this work the SOURCES code was used to generate neutrons from the decay chains of U and Th in the rock. The Watt spectrum was used to generate events from spontaneous fission whilst for alpha-n reaction events were obtained by using the relevant isotope lifetimes, energy spectra of alphas, reaction cross sections, alpha stopping powers etc. [continue description as in the Carson paper]

The results from these simulations are outlined in Table 2 and plotted in Fig. x. It can be seen from these that at least 1.5m of passive neutron shielding is required in all scenarios, though for the hard rock case of Option 1 this rises to around 1.7 m [these are guesses] due to the higher U content of that rock. For option 1 this amounts to a total mass of passive shielding of xx tons.

5.2.3. Vessel and TPC Neutrons

As seen in Sec 5.2.2 the external rock neutron flux can be controlled by passive shielding, as expected. Neutrons from internal detector radioactivity is know to be a harder challenge since control of this likely relies on selecting radio-pure materials, use of tricky internal shielding or innovative analysis techniques. The vacuum vessel, being the component with by far the largest mass, has the potential to dominate this aspect, followed by the outer passive shield, TPC field cage and resistors. Rather than assume values for the U and Th content of the vessel materials the approach taken here is to determine from the simulations what amount of U, Th contamination can be tolerated in each of these major components without compromising the criteria of ¡1 event recorded per year, at each energy threshold. Given the potential difficulty of obtain steel with low U, Th content simulations were also performed assuming an acrylic vessel. [describe any details of the simulations, see Carson etc]

The form of the neutron energy spectrum for each of the components, normalized to x ppb U and y ppb Th, is shown in Fig. x. Table x summarizes the results in terms of specifications on the maximum U, Th levels that can be tolerated in each scenario.

Some initial conclusions can be extracted from these results. Firstly, well selected steel has been measured to have U and Th content as low as typically x and y ppb respectively. This a factor x10 and x10 higher than would be tolerable according to the results obtained. To our knowledge no serious efforts to develop or pre-select steel for ultra-low background have been made so it is conceivable that steel with this level of contamination can be obtained. The alternative of acrylic looks more favourable since there has been extensive work on developing low U, Th material. For instance, levels as low as x and y PPB U and The have been reported [ref]. However, there are clearly significant mechanical challenges with this option. An alternative could be to mount non-structural acrylic shielding within a steel vessel, to shield off steel related neutrons. To explore this, further simulations were run using the Option 1 steel vessel design but with internal acrylic added. Results for this are shown in Table x. It can be seen that at of order 20cm (??) of internal acrylic would be needed to gain a factor x10 in neutron rate.

Regarding the internal TPC components, including resistor chain but excluding the readout planes, the specification on U and Th content appear achievable based on known levels as measured [probably not true?], for instance...

5.2.4. Muon-induced neutrons and active vetoing

The main question here I think is to determine what veto efficiency is needed and to make comments on how that would be achieved in practice. This could include an assumption on the sensitivity for vetoing muon neutrons by recording coincident EM in the TPC, which in turn depends on the readout.

The nuclear option is to go deeper underground.

Include some comment on double scatter vetoing both internal to the detector (hard) and external using the muon veto.

5.2.5. Neutron Conclusion

Summary of the design specs needed

5.3. Gamma Backgrounds

5.4. Radon and Radon Progeny Backgrounds

5.5. Surface and other Backgrounds

5.6. Comparison of Technologies for low background

6. Underground Sites and Engineering

[Section organizer: Neil Spooner]

This covers the requirements for and feasibility of achieving the necessary engineering and underground site infrastructure including the scientific argument for multiple sites? E.g., for a 1D detector, is there an advantage to distributing the same total target mass over multiple sites? Or do you get the same benefit from multiple orientations at the same site?

7. Conceptual Design Strategy

[Section organizer: all]

This summarizes the above technology discussions and briefly outlines possible scenarios and a straw man design for a Galactic Recoil Observatory

8. Conclusion

[Section organizer: all]

This section restates the science case in light of the technology discussion and provides comment on the likely feasibility, cost and design of a future large scale galactic recoil observatory.

References