ANNUAL REVIEWS

Annual Review of Nuclear and Particle Science

Experimental Considerations in Long-Baseline Neutrino Oscillation Measurements

Francesca Di Lodovico,[1] Ryan B. Patterson,[2] Masato Shiozawa,[3] and Elizabeth Worcester[4]

[1]Department of Physics, King's College London, London, United Kingdom

[2]Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, California, USA

[3]Kamioka Observatory, Institute for Cosmic Ray Research, The University of Tokyo, Kamioka-cho, Hida-city, Gifu, Japan

[4]Physics Department, Brookhaven National Laboratory, Upton, New York, USA; email: etw@bnl.gov

Annu. Rev. Nucl. Part. Sci. 2023. 73:69–93

First published as a Review in Advance on June 27, 2023

The *Annual Review of Nuclear and Particle Science* is online at nucl.annualreviews.org

https://doi.org/10.1146/annurev-nucl-102020-101615

Keywords

neutrinos, oscillation, long baseline

Abstract

Long-baseline neutrino oscillation experiments, which are among the largest neutrino experiments in the world, have extensive physics programs to make precision measurements of three-flavor oscillation parameters, search for physics beyond the Standard Model, and study neutrinos from astrophysical sources. In this article, experimental considerations, including oscillation phenomenology, detector and experiment design, and analysis strategies, are described, with a focus on the three-flavor oscillation measurements. Current and future experiments are discussed, and significant sources of systematic uncertainty, along with mitigation strategies, are emphasized as control of systematic uncertainty is critical for success in precise measurement of long-baseline oscillation parameters. This article is structured as a primer for those new to this area of experimental work.

Contents

1. INTRODUCTION 70
1.1. Oscillation Phenomenology 71
1.2. Oscillations at Long Baseline 73
1.3. Experimental Considerations 74
2. CURRENT AND FUTURE EXPERIMENTS 75
3. NEUTRINO BEAMS 77
4. NEUTRINO–NUCLEUS INTERACTIONS 79
5. FAR DETECTORS 81
5.1. Water Cherenkov Detectors 81
5.2. Tracking Calorimeter Detectors 82
6. NEAR DETECTORS 84
6.1. Functionally Identical Near Detectors 84
6.2. Multicomponent Detectors 84
6.3. Segmented Near Detectors 84
6.4. Argon Near Detectors 85
6.5. Movable Detectors 85
6.6. Additional Nuclear Targets 85
7. ANALYSIS STRATEGIES 85
7.1. Predicting Far Detector Spectra 86
7.2. Statistical Procedures 87
8. OTHER PHYSICS WITH LONG-BASELINE EXPERIMENTS 88
9. CONCLUSION 89

1. INTRODUCTION

Neutrino oscillation is the phenomenon whereby a neutrino created with a definite flavor (e.g., ν_e, ν_μ, ν_τ) may be detected with a different flavor sometime later. Neutrino oscillation is accommodated in the Standard Model of particle physics by extending the model to include nonzero neutrino masses along with some level of mixing between the neutrino flavor eigenstates and the neutrino mass eigenstates, typically denoted ν_1, ν_2, and ν_3.

Neutrino oscillations have been observed by a number of experiments using neutrinos produced by the Sun, by cosmic ray interactions in the atmosphere, by nuclear reactors, and by particle accelerators. The phenomenon was first established definitively around 2000 (see the sidebar titled Nobel Prize). Global fits to world neutrino oscillation data, such as in Reference 1, provide a comprehensive list of experimental inputs to our understanding of neutrino oscillation. Different neutrino sources offer different neutrino energies, flavors, and propagation times that, when coupled with various detector approaches, can offer very different views into the underlying structure of neutrino masses and flavor mixing.

This article focuses on one common and flexible experimental configuration—namely, that of long-baseline neutrino oscillation experiments. Long-baseline experiments provide access to a particularly wide range of neutrino physics and are highly relevant at present given the most pressing questions surrounding neutrino masses and mixings.

NOBEL PRIZE

The 2015 Nobel Prize in Physics was awarded in recognition of the first definitive observations of neutrino oscillations by the Super-Kamiokande and SNO Collaborations (2–4), a discovery that established that neutrinos indeed have mass. Super-Kamiokande observed a deficit of muon neutrinos produced in the atmosphere by cosmic ray interactions relative to expectation. In their 1998 result (2), they compared the number of events coming from above (thus directly from the atmosphere) and below (thus passing through the Earth, traveling a longer distance) and were able to observe an oscillation signature. The solar neutrino problem—a deficit in electron neutrinos observed relative to the number predicted by solar models—had been around since the Davis experiment at Homestake in the 1960s. Originally, many thought that the models were wrong. However, the SNO experiment (4) demonstrated that the total number of neutrinos from the Sun matched expectations and that the deficit was specifically from electron neutrinos, resolving the solar neutrino problem in favor of the oscillation hypothesis. Takaaki Kajita and Arthur B. McDonald were recognized by the Nobel Committee for their leadership of these groundbreaking experiments, which led the way for the highly active field of neutrino oscillation research that followed.

1.1. Oscillation Phenomenology

In the three-flavor neutrino paradigm, there are three neutrino mass eigenstates—ν_1, ν_2, and ν_3—with distinct, nonzero masses for at least two of the three, such that there are two mass splittings. The mass states are distinct from the flavor eigenstates, which are labeled ν_e, ν_μ, and ν_τ based on the charged lepton to which they couple in the charged-current (CC) weak interaction. The flavor eigenstates can be expressed as linear combinations of the mass eigenstates and vice versa, and the coefficients of these linear combinations form a mixing matrix called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix (5, 6):

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad 1.$$

This 3×3 matrix can be expressed in terms of three mixing angles [θ_{12} ($\sim 33^\circ$), θ_{13} ($\sim 9^\circ$), and θ_{23} ($\sim 49^\circ$)] and a potentially *CP*-violating complex phase (δ_{CP}), which is currently unknown, though current experiments are beginning to set limits on its value. The octant of θ_{23}, which is the question of whether the value of θ_{23} is exactly 45°, less than 45° (lower octant), or greater than 45° (upper octant), is unknown and is of interest because it could point to a previously unknown symmetry.

The mass states are defined such that Δm^2_{21} is the smaller mass splitting (order 10^{-4} eV^2) and $m^2_2 > m^2_1$. The remaining mass splitting, Δm^2_{32}, is of order 10^{-3}, and its sign is unknown. The condition $\Delta m^2_{32} > 0$ is referred to as normal ordering, while the condition $\Delta m^2_{32} < 0$ is called inverted ordering.[1] **Figure 1** shows the fractional flavor content and possible orderings of the neutrino mass states.

Neutrinos are created via weak interactions in definite flavor states. Oscillations occur as neutrinos travel through space because these initial flavor states are not eigenstates of the Hamiltonian that governs their propagation. The state thus evolves into an admixture of flavor states. The oscillation probability—that is, the probability of observing a particular flavor change—depends on the

[1]Some authors use the term hierarchy rather than ordering, though this usage seems to be losing favor over time. All the same, one will still find many references to mass hierarchy, normal hierarchy, and inverted hierarchy in the literature.

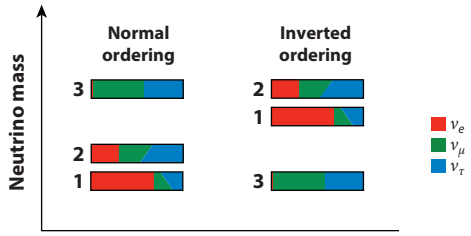

Figure 1

Fractional flavor content $|U_{\alpha i}|^2$ ($\alpha = e, \mu, \tau$) of the three mass eigenstates ν_i based on the current best-fit values of the mixing angles. δ_{CP} is varied from 0 (*bottom of each colored band*) to 180° (*top of each colored band*) for normal and inverted mass ordering on the left and right, respectively. The different colors correspond to the ν_e (*red*), ν_μ (*green*), and ν_τ (*blue*) fractions. Figure inspired by Reference 7.

elements of the PMNS mixing matrix (or equivalently the mixing angles and the complex phase δ_{CP}), on the differences between the squared masses of the mass eigenstates, and on the experimental ratio L/E, where L is the distance traveled (the baseline) and E is the neutrino energy. In a simplified scenario with only two neutrino types, or in practical scenarios where a two-neutrino approximation can be applied, the oscillation probability takes on a straightforward form:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right), \quad \text{for } \alpha \neq \beta. \qquad 2.$$

If $E/L \gg |\Delta m^2|$, then the oscillation probability is very small and no effect will be observed. Similarly, if $E/L \ll |\Delta m^2|$, many oscillations will occur between production and observation such that the oscillation term averages to 1/2. This allows experimental control for accessing different mixing parameters. **Table 1** shows the characteristic values for L, E, and the corresponding values of Δm^2 to which they are most sensitive. As a result of these characteristics, θ_{12} and Δm^2_{21} are historically referred to as the solar mixing parameters, and θ_{23} and Δm^2_{32} are called the atmospheric mixing parameters, because of the samples in which they were first measured. Note that in the case of solar neutrino oscillations, the oscillation probability is significantly modified by matter effects as the neutrinos travel through the dense and variable solar medium (8), and thus the expressions in this article are not directly applicable. Recent measurements of θ_{13} are dominated by short-baseline reactor experiments.

Table 1 Characteristic values of L and E for experiments performed using various neutrino sources and the corresponding ranges of $|\Delta m^2|$ to which they can be most sensitive to flavor oscillations in vacuum

Source	L (m)	E (MeV)	$\lvert\Delta m^2\rvert$ (eV^2)
Solar	10^{10}	1	10^{-10}
Atmospheric	10^4–10^7	10^2–10^5	10^{-1}–10^{-4}
Reactor			
Short baseline	10^2–10^3	1	10^{-2}–10^{-3}
Long baseline	10^4–10^5	1	10^{-4}–10^{-5}
Accelerator			
Short baseline	10^2	10^3–10^4	>0.1
Long baseline	10^5–10^6	10^3–10^4	10^{-2}–10^{-3}

Data from Reference 9.

Long-baseline accelerator experiments are largely insensitive to the solar parameters but are able to measure θ_{23}, Δm^2_{32}, θ_{13}, and δ_{CP} in a single experiment if the baseline is long enough, with sensitivity to θ_{13}, δ_{CP}, and the θ_{23} octant coming from the ν_e appearance channel. Experimental results suggesting the possibility of anomalous oscillations at values of L/E that cannot be explained by the three flavors described here have given rise to searches for sterile neutrinos (10): additional neutrino states that interact with matter only via gravity but that mix with the known neutrino states. Such searches are beyond the scope of this article, but long-baseline experiments do have sensitivity to these and other new phenomena.

1.2. Oscillations at Long Baseline

The primary signal channels for long-baseline oscillation experiments, ν_μ CC interactions and ν_e CC interactions, are identified using the flavor of the outgoing charged lepton. Because accelerator-based neutrino beams are dominated by muon neutrinos, observations of ν_μ and ν_e interactions are sometimes referred to as disappearance mode and appearance mode, respectively. The third potential channel, ν_τ interactions from $\nu_\mu \to \nu_\tau$ oscillation, introduces a number of unique experimental challenges due to the relatively high energy threshold for the ν_τ CC process (stemming from the τ mass) and due to the difficulty in identifying the short-lived τ particle in large neutrino detectors. Thus, the $\nu_\mu \to \nu_\tau$ oscillation channel is not discussed in this review and also does not feature centrally in many long-baseline experiments. One notable exception is the OPERA (Oscillation Project with Emulsion-Tracking Apparatus) experiment (11), which was designed from the start to directly observe $\nu_\mu \to \nu_\tau$ oscillations.

The $\nu_\mu \to \nu_e$ appearance probability for neutrinos traveling through a reasonably uniform medium (e.g., Earth's crust) is

$$
\begin{aligned}
P(\nu_\mu \to \nu_e) \simeq\ & \sin^2\theta_{23} \sin^2 2\theta_{13} \frac{\sin^2(\Delta_{31} - aL)}{(\Delta_{31} - aL)^2} \Delta^2_{31} \\
& + \sin 2\theta_{23} \sin 2\theta_{13} \sin 2\theta_{12} \frac{\sin(\Delta_{31} - aL)}{(\Delta_{31} - aL)} \Delta_{31} \frac{\sin(aL)}{(aL)} \Delta_{21} \cos(\Delta_{31} + \delta_{CP}) \\
& + \cos^2\theta_{23} \sin^2 2\theta_{12} \frac{\sin^2(aL)}{(aL)^2} \Delta^2_{21},
\end{aligned}
\tag{3.}
$$

where $\Delta_{ij} = \Delta m^2_{ij} L/4E$, $a = G_F N_e/\sqrt{2}$, G_F is the Fermi constant, N_e is the number density of electrons in the medium, L is the baseline, and E is the neutrino energy. Both δ_{CP} and a switch signs for the equivalent antineutrino oscillation probability, such that both *CP* violation and matter effects give rise to a neutrino–antineutrino asymmetry. The size of the *CP*-violating effect can be parameterized by the Jarlskog invariant (12):

$$
J^{\text{PMNS}}_{CP} \equiv \frac{1}{8} \sin 2\theta_{12} \sin 2\theta_{13} \sin 2\theta_{23} \cos\theta_{13} \sin\delta_{CP}. \tag{4.}
$$

The matter effect occurs because neutrinos passing through matter experience coherent forward scattering that modifies their propagation. Because ν_e particles can have both CC and neutral-current (NC) interactions with the electrons in matter without a flavor change, while ν_μ and ν_τ particles can have only NC flavor-preserving interactions with electrons, the effective Hamiltonian is modified in a way that enhances the appearance probability for neutrinos and suppresses it for antineutrinos in normal ordering; the opposite is true for inverted ordering. This matter–antimatter asymmetry is degenerate with that produced by *CP* violation. The size of the matter effect increases with energy, and thus the associated asymmetry increases in practice with

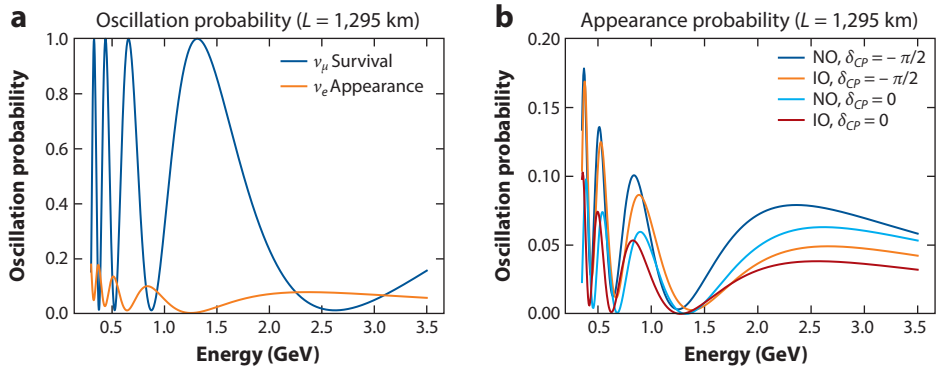

Figure 2

Visualizations of $P(\nu_\mu \rightarrow \nu_\mu)$ and $P(\nu_\mu \rightarrow \nu_e)$ as functions of neutrino energy for a baseline of 1,295 km. (*a*) The two oscillation channels for a fixed set of oscillation parameters. (*b*) The ν_e appearance probability under four different parameter scenarios, which vary between the two possibilities for neutrino mass ordering (normal or inverted) and two choices for δ_{CP} ($-\pi/2$ or 0). The resulting variations in the oscillation probability provide long-baseline experiments their sensitivity to these parameters. Abbreviations: IO, inverted ordering; NO, normal ordering.

baseline; for baselines greater than about 1,000 km, the asymmetry from matter effects is greater than that from maximal *CP* violation, allowing experimental disentanglement of the two effects.

Neutrino oscillation probabilities as a function of energy are illustrated in **Figure 2** assuming a DUNE-like[2] experimental baseline (1,295 km) and setting the values of oscillation parameters near the current global best fits. The ν_μ survival ($\nu_\mu \rightarrow \nu_\mu$) probability varies from zero to one while the ν_e appearance ($\nu_\mu \rightarrow \nu_e$) probability at the first oscillation maximum (rightmost peak in each panel of **Figure 2**) is below 10%. The remainder of the probability is accounted for by $\nu_\mu \rightarrow \nu_\tau$ oscillations. Experiments typically match their baseline and neutrino energy range to be sensitive to the region around the first oscillation maximum. For typical ν_μ energies from conventional neutrino beams (of order 1 GeV), the optimal baseline is of order hundreds of kilometers.

NC interactions are also studied at long-baseline experiments. In NC interactions, the incoming neutrino scatters via exchange of a Z boson. The outgoing neutrino is undetected, and there is no charged lepton by which to tag the incident neutrino's flavor. For this reason, measurement of NC interactions is not useful for studying specific flavor transitions. However, NC measurements are sensitive to all three neutrino flavors, so the measured rates in both the near and far detectors should be independent of all oscillation parameters. NC analyses can therefore be useful to search for anomalous neutrino disappearance, as in References 13–15.

1.3. Experimental Considerations

Long-baseline neutrino oscillation experiments are designed to measure oscillation parameters governing $\nu_1\nu_3$ and $\nu_2\nu_3$ mixing by observing ν_μ disappearance and ν_e and ν_τ appearance over a baseline of hundreds of kilometers. These experiments make use of conventional accelerator-based neutrino sources whose beams consist primarily of muon neutrinos or muon antineutrinos. The energy of the beam is chosen to match an experiment's baseline such that the far detector (FD) is at or near the location of first maximum oscillation probability. In some experiments the

[2]DUNE is a future long-baseline experiment (see Section 2 for a description).

FD is located slightly off-axis from the beam such that the neutrino energy spectrum at the FD is relatively narrowly peaked near the optimum L/E, while other experiments choose to place the FD on-axis to observe a broad range of energies and thus more of the oscillation period in L/E. In addition to the large FD, these experiments include a near detector (ND) placed close to the neutrino source. The ND provides vital control over systematic uncertainties by measuring neutrino interactions at a location where oscillations are still negligible.

To accumulate sufficient statistics at the FD, given the small cross section for neutrino interactions and the reduction in beam flux with increasing distance from the source, high-power beams and massive detectors are required. More details about neutrino beams and fluxes are given in Section 3. The FD for a long-baseline experiment must be many kilotons in mass at a minimum. Until an experiment becomes limited by systematic errors, the sensitivity of the experiment scales as the square root of exposure (the product of beam power, FD mass, and time).

The main signal modes are ν_μ CC and ν_e CC interactions, so the FD must be capable of distinguishing ν_μ from ν_e interactions by identifying the flavor of the outgoing lepton and must be capable of measuring the energy of final-state particles to estimate, in turn, the incoming neutrino's energy. NC interactions with a neutral pion in the final state can be misidentified as ν_e CC interactions if the two gamma particles from π^0 decay are mistaken for an electron, so detectors must be able to resolve the two gammas and/or distinguish gammas from electrons. Resolving details of the neutrino interaction final state can improve the neutrino energy estimate and, as described in Section 4, can also reduce the impact of uncertainties in the neutrino interaction model. More discussion of FDs is provided in Section 5.

NDs provide measurements of neutrino interactions before oscillations, which can be used to form expectations for the rates and spectra of different neutrino flavors at the FD after oscillations and, critically, to constrain systematic uncertainties related to neutrino flux, neutrino–nucleus interaction cross sections, and detector effects, depending on the design and composition of the ND. Because the neutrino flux is large close to the source, NDs collect very large samples of neutrino interactions and must be designed to make successful measurements in a higher-rate environment than FDs.

Two distinct ND strategies have been used in experiments historically. In one approach, the ND is designed to resemble the FD in form and function to the greatest extent possible, allowing relatively direct mitigation of systematic uncertainties as the relevant cross sections, detector effects, and analysis techniques are well matched between the detectors. In the other approach, ND measurements are used to constrain more overtly the various sources of systematic uncertainties, usually encapsulated in a detailed experimental model, by employing multiple detector technologies with higher performance requirements than the FD. The ambitious physics goals of next-generation experiments require ND designs that incorporate aspects of both approaches as well as new ND techniques not previously used. More details about ND design and analyses are provided in Section 6.

2. CURRENT AND FUTURE EXPERIMENTS

There are currently two operating accelerator-based, long-baseline neutrino oscillation experiments: T2K (Tokai to Kamioka) (16) in Japan and NOvA (NuMI Off-axis ν_e Appearance) (17) in the United States. These experiments have been taking data since 2009 and 2014, respectively, and both plan to continue collecting data into the mid-2020s. Both measure ν_μ disappearance and ν_e appearance in off-axis, narrow-band neutrino beams at baselines of 295 and 810 km, respectively. As discussed in Section 5, the NOvA ND and FD are functionally identical segmented liquid scintillator detectors, and T2K has a suite of detectors at its near site and uses the Super-K

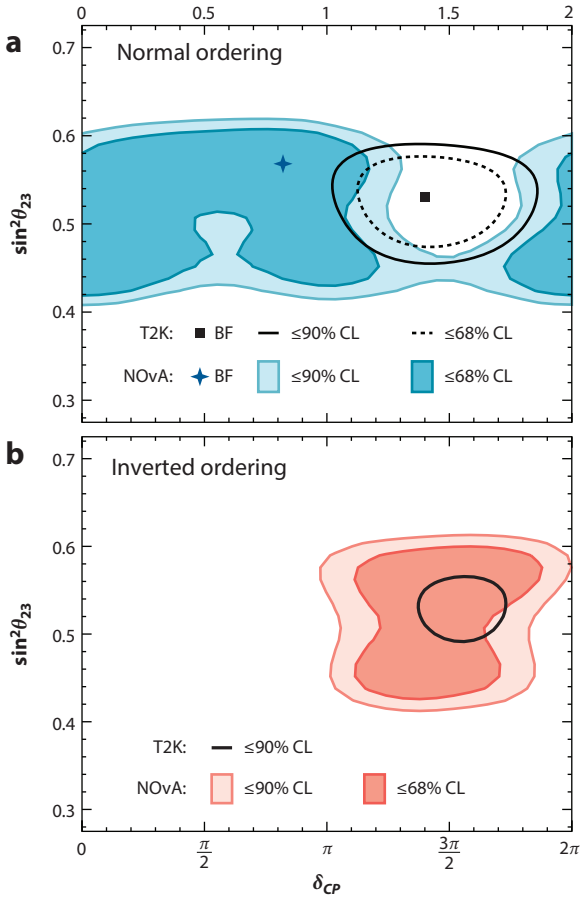

Figure 3

The best-fit (BF) values of $\sin^2\theta_{23}$ and δ_{CP} to NOvA and T2K data for the (*a*) normal neutrino mass ordering and (*b*) inverted neutrino mass ordering. Figure adapted from Reference 17 (CC BY 4.0).

(Super-Kamiokande) water Cherenkov detector (WCD) as its FD. NOvA measures neutrinos from the NuMI beam at Fermilab, while T2K measures neutrinos from the J-PARC beam. Both experiments measure ν_e and $\bar{\nu}_e$ appearance and have published measurements of Δm^2_{32}, $\sin^2 2\theta_{23}$, and δ_{CP}. **Figure 3** shows a comparison of the best-fit values for $\sin^2\theta_{23}$ and δ_{CP} for the two experiments, using external data to constrain θ_{13}. Both experiments slightly favor the normal ordering. T2K data have a large asymmetry in $\nu_\mu \rightarrow \nu_e$ versus $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, thus favoring *CP*-violating values of δ_{CP}, while NOvA does not have this asymmetry, resulting in a slight tension between the experiments (17). The NOvA and T2K Collaborations are pursuing joint fits to their oscillation data to take direct advantage of their design complementarity (e.g., very different balances of matter-induced and potential *CP*-violation-induced oscillation asymmetries) while also understanding and incorporating joint systematic uncertainty treatments where required. The experiments plan to release a first joint fit to their combined data in the coming year.

Both Fermilab and J-PARC/Kamioka plan to host next-generation neutrino experiments with the goal of definitively addressing the open questions in neutrino oscillations. The current generation of long-baseline oscillation experiments is statistically limited; next-generation experiments

Table 2 Summary of experimental details for current and future long-baseline neutrino oscillation experiments

	Location	Beam	Baseline	Near detector	Far detector
T2K	Japan (Tokai to Kamioka)	J-PARC 500 kW (upgrade to 1.3 MW)	295 km	Suite of detectors, on- and off-axis	Water Cherenkov, 22.5 kt fiducial, off-axis
NOvA	United States (Fermilab to Ash River, Minnesota)	NuMI 850 kW	810 km	Segmented liquid scintillator, off-axis	Segmented liquid scintillator, 14 kt active, off-axis
Hyper-K	Japan (Tokai to Kamioka)	J-PARC 1.3 MW	295 km	Suite of detectors, on- and off-axis	Water Cherenkov, 187 kt fiducial, off-axis
DUNE	United States (Fermilab to Lead, South Dakota)	LBNF 1.2 MW (upgrade to 2.4 MW)	1,285 km	Liquid argon time projection chamber plus suite of detectors, on-axis, movable off-axis	Liquid argon time projection chamber, 40 kt fiducial, on-axis

will need to significantly increase the size of FD samples and improve constraints of systematic uncertainties to take advantage of the increased statistics. DUNE (Deep Underground Neutrino Experiment) (18–20) is the next step in the evolution of the Fermilab program, and Hyper-K (Hyper-Kamiokande) (21, 22) builds on the highly successful Super-K and T2K experiments. Hyper-K will use a 1.3-MW upgraded beam from J-PARC and proven WCD technology to rapidly accumulate the statistics necessary for precision measurements. DUNE is designed as a precision experiment, viewing a new broadband neutrino beam with a liquid argon time projection chamber (LArTPC) FD, which allows precision imaging of all final-state particles. Both experiments plan to use an ensemble of off-axis ND measurements to reduce systematic uncertainty from neutrino interaction modeling, as described in Section 7. References 18 and 22 provide the most up-to-date evaluations of expected experimental sensitivity for DUNE and Hyper-K.

The basic experimental details of these experiments are summarized in **Table 2**. While the focus of this article is to describe general principles of long-baseline experimental techniques, in many cases it is illustrative to discuss specific examples from current and future experiments. For this reason, additional details of the design and analyses for each of these experiments are included in the following sections.

3. NEUTRINO BEAMS

Conventional neutrino beams are produced using protons from an accelerator complex that impinge on a target to produce secondary hadrons that are then focused into a nearly parallel beam using magnetic focusing horns. The focused hadrons, primarily charged pions, decay into final states that include muon neutrinos (e.g., $\pi^+ \rightarrow \mu^+\nu_\mu$) and form a beam of neutrinos with an energy spectrum determined primarily by the kinematics of the focused hadrons and their decays. The sign of the current in the electromagnetic horns determines whether positive or negative particles are focused and, thus, whether the beam is predominantly composed of neutrinos or antineutrinos; for this reason, neutrino mode is sometimes referred to as forward horn current (FHC) mode and antineutrino mode as reverse horn current (RHC) mode. As a result of neutrinos from imperfectly focused or unfocused particles and from less common decay modes, the beam will contain some wrong-sign contamination (e.g., neutrinos in antineutrino mode) as well as an intrinsic ν_e component. The intrinsic ν_e component is an irreducible background to ν_e appearance in long-baseline experiments and is typically at the percent level. The wrong-sign component is larger in antineutrino mode than in neutrino mode because the focused mesons are produced from a positively charged initial state (proton collisions with the target), and thus

positively charged particles are more abundant than negatively charged particles. Optimization of the neutrino beam design can minimize both of these impurities in the beam, but they cannot be eliminated entirely using a horn-focused system.

For a given beam configuration, the number of neutrinos produced during a run period is directly tied to the number of protons delivered to the production target. Therefore, it is common for beam facilities and experiments to report their protons-on-target (POT) count. The instantaneous rate of protons can also be cast in terms of the proton beam's power, and thus this power is often reported as well. The horn-focused neutrino beams currently in use for long-baseline experiments are the NuMI beam at Fermilab (23) and the T2K experiment's neutrino beam to Kamioka from J-PARC (24). The NuMI beam has approached 900-kW power, and as of late 2022 the complex has delivered 3.9×10^{21} POT to the NOvA experiment, plus additional protons during prior experimental operators. Two-thirds of this NOvA running has so far been in neutrino (FHC) mode. The J-PARC beam has ramped up to approximately 500 kW so far and has also delivered 3.9×10^{21} POT for the T2K experiment, 56% of which has been in neutrino mode. The beam is going to be gradually upgraded up to 1.3 MW for the start of the future experiment in Japan, Hyper-K.

DUNE in the United States will depend on the construction of the initial phase of the LBNF (Long-Baseline Neutrino Facility) neutrino beamline at Fermilab (25). J-PARC and LBNF Phase I have planned proton beam powers of 1.3 and 1.2 MW, respectively. A beam power upgrade of the LBNF beamline from 1.2 to 2.4 MW in the second phase of DUNE is also planned.

For a detector placed directly downstream along the beam axis, a broad range of neutrino energies, proportional to the energies of the decaying pions, will be observed. A detector placed somewhat off-axis from the beam, in contrast, will observe neutrinos in a narrower range of energies since the higher the parent pion energy is, the lower the neutrino energy is along a given off-axis trajectory from the decay (see **Figure 4**). The off-axis angle also has a small impact on the flavor content of the beam.

Some experiments, such as MINOS (Main Injector Neutrino Oscillation Search) and DUNE, have chosen to place their detectors on-axis to maximize the neutrino flux and observe a broad

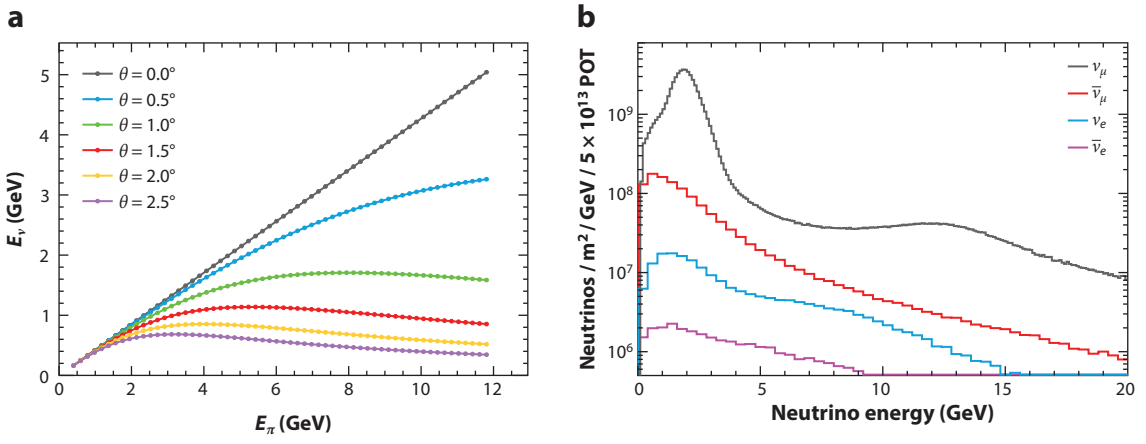

Figure 4

(*a*) Neutrino energy as a function of parent pion energy for different angles off-axis from the beam. Panel adapted with permission from Reference 26. (*b*) Predicted composition of the NuMI beam at the NOvA near detector with the horns focusing on positively charged hadrons (note the logarithmic vertical scale). The muon neutrino flux peaks at 1.9 GeV. Abbreviation: POT, protons on target. Panel adapted from Reference 27 (CC BY 4.0).

range of neutrino energies. This allows for measurements of neutrino oscillations across a full oscillation period or more, explicitly mapping out the oscillation behavior and allowing for checks of consistency with expectations from three-flavor oscillations. For other experiments, a choice of off-axis angle is made such that the observed neutrinos will peak at a desired energy with minimal background contributions from misreconstructed events feeding down from higher energies. In NOvA and T2K, the off-axis angle (0.8 and 2.5°, respectively) has been chosen to closely match the first oscillation maximum, enhancing the observable appearance and disappearance effects. The DUNE and Hyper-K Collaborations plan to further exploit this feature of neutrino beams by making ND measurements at multiple off-axis angles to better constrain systematic uncertainties; this technique is called PRISM (precision reaction-independent spectrum measurement) and is discussed further in Section 7.

Uncertainties in the prediction of the neutrino flux relate to the production of the secondary hadrons (hadron species, rates, and momenta) and to the details of the neutrino beamline such as the horn or other beamline elements' alignment and positioning. The hadron production uncertainties are typically larger than the beamline uncertainties. However, in long-baseline experiments, where the near-to-far flux ratio is more important than the absolute flux, focusing uncertainties may be the dominant source of flux-related uncertainty in measuring oscillation parameters. The total a priori flux uncertainty is typically around 10%.

Measurements of neutrino interactions in a detector are influenced by both the flux and cross sections, and it can be difficult to study either of these independently of the other in situ with, for instance, an experiment's ND data. However, there are several techniques that help. Fully leptonic neutrino interactions, such as neutrino–electron elastic scattering ($\nu e \rightarrow \nu e$), have cross sections that can be precisely calculated because they involve only weak interactions of fundamental leptons, so the neutrino flux can be measured without significant uncertainty from the cross section. For νe scattering events, the signal is a single, very forward-going electron with primary backgrounds coming from ν_e CC interactions or misidentified neutral pion decays. The MINERvA (Main Injector Neutrino Experiment to Study ν-A Interactions) experiment has used this technique to reduce the total uncertainty on the NuMI ν_μ flux from 9% to 6% (28). Given the relatively low cross section for νe scattering, the MINERvA measurement was statistically limited; in DUNE, this sample could be used to reduce the ν_μ flux uncertainty in the peak of the LBNF beam to 1–3% (29).

In another approach, neutrino-on-nucleus scattering events with very low energy transfer ν are selected.[3] For such low-ν events, the interaction cross section is nearly constant across neutrino energy, and thus the shape of the neutrino energy spectrum can be determined (30–32). The low-ν method has been used recently in the MINERvA experiment to constrain uncertainties on flux shape in the NuMI beam to well below the normalization uncertainties (33). Additionally, the high-energy part of the ν_μ spectrum is dominated by kaon decays, and this portion of the spectrum may be used to constrain the kaon component of the beam (34).

4. NEUTRINO–NUCLEUS INTERACTIONS

In long-baseline oscillation experiments with neutrino energies from ∼0.1 to ∼10 GeV, a variety of neutrino interaction channels may be relevant. CC channels are particularly important because they include the final-state charged leptons that are required for flavor tagging, and thus they allow for reasonably complete neutrino energy estimation. In contrast, NC channels leave

[3]Lowercase nu (ν) is the symbol most commonly used for this energy transfer. Note that the same symbol is used to refer to neutrinos.

no flavor information behind and typically have significant missing energy since the final-state neutrino exits the detector invisibly. Thus, NC channels are a source of backgrounds in oscillation measurements that must be removed through careful event selections, although these events do have other experimental uses (e.g., providing particles for detector calibration or allowing for certain new physics searches).

CC interactions in long-baseline experiments are often discussed and simulated in four categories: In quasielastic (QE) scattering, the outgoing particles are the charged lepton and a single nucleon; in meson exchange current (MEC) interactions, momentum is transferred to a quasi-bound dinucleon within the nucleus, resulting in typically two ejected nucleons; in resonant scattering (RES), the struck nucleon becomes excited into a resonant baryon state, which then decays typically into a nucleon and a pion; and in deep inelastic scattering (DIS), the neutrino momentum is high enough to resolve and to scatter off partons within the nuclei, leading typically to multihadron final states. Coherent interactions, wherein the entire nucleus recoils coherently, are also present but at much lower rates than the above at these energies. **Figure 5** shows the neutrino and antineutrino cross sections as a function of neutrino energy for QE, RES, and DIS channels. There remains significant uncertainty on the MEC component, and modern oscillation data analyses still constrain the MEC component empirically and with significant systematic uncertainties.

The initial state of the struck nucleon is affected by its nuclear environment. Fermi motion within the nucleus or, more generally, all the intranuclear interactions influence the final observed state. Additionally, the particles leaving the primary interaction have to travel out of the nuclear medium, and so-called final-state interactions can occur during this period, modifying the momenta or particle content of the final state. Both of these nuclear complications are significant and must be included in a complete description of neutrino scattering in long-baseline experiments.

Like most particle physics data analyses, long-baseline oscillation analyses use detailed simulations of relevant particle processes and detector elements to develop analysis algorithms and form predictions for experimental observables. The simulation chain used in these neutrino experiments

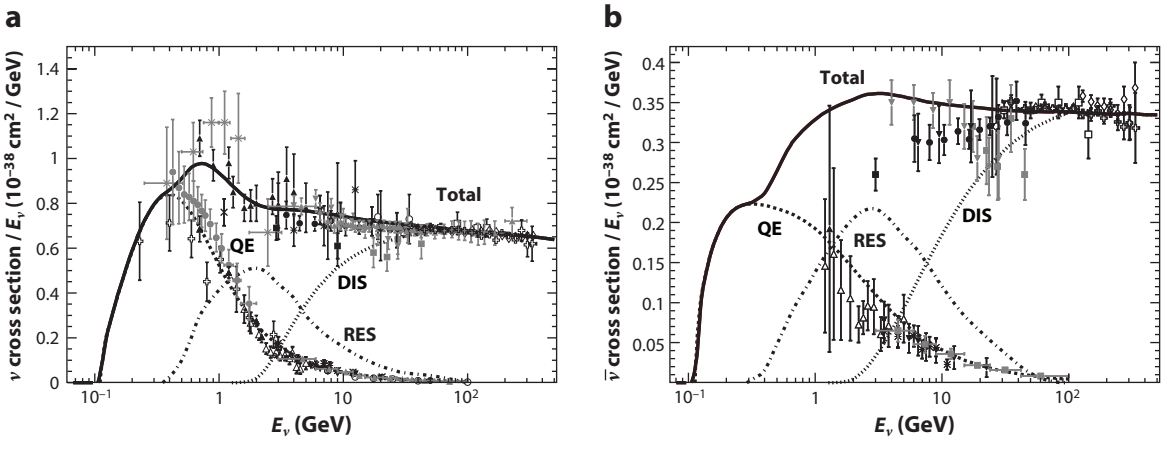

Figure 5

Total (*a*) neutrino and (*b*) antineutrino per nucleon charged-current cross sections for an isoscalar target divided by neutrino energy and plotted as a function of energy. Also shown are the various contributing processes. These contributions include quasielastic (QE) scattering (*dashed lines*), resonant scattering (RES; *dotted–dashed lines*), and deep inelastic scattering (DIS; *dotted lines*). Example predictions for each are provided by the NUANCE generator (35). See Reference 36 for detailed explanations of the data sets represented in this figure. Figure adapted with permission from Reference 36.

includes a detailed model of neutrino interactions, following the general picture above. Neutrino event generators such as GENIE (37), GIBUU (38), NEUT (39), and NUWRO (40) attempt to implement as full a description of neutrino–nucleus interactions as possible by combining the relevant processes into a cohesive model and tuning the parameters to best fit existing neutrino data. References 9, 41, and 42 provide additional information on the current status of neutrino interaction measurements and modeling efforts. To date, no model adequately describes all neutrino scattering data, a fact that reflects the complexity of neutrino-on-nucleus interactions. A complication in tuning to data is that scaling from one target nucleus to another is nontrivial, and measurements of neutrino interactions using detectors with different target nuclei are not directly comparable.

Typical a priori uncertainties in neutrino interaction models are 10–20%, and a large number of model parameters are required. NDs are designed to constrain these uncertainties by making high-statistics, high-precision measurements of neutrino interactions that can be used to constrain model parameters, to identify shortcomings of the model, or to provide data-driven predictions of FD spectra to minimize model dependence directly. NDs are often designed with the same target nucleus as the FD to eliminate uncertainty from differing target nuclei, but this is not always possible for practical reasons. Exclusive final-state samples, such as a QE-like sample or samples defined by the number of final-state pions, may be selected to separately constrain uncertainties from individual components of the model. Samples of neutrino interactions taken at different off-axis angles—and thus with different neutrino spectra as described in Section 3—may be used to decouple interaction model and flux model uncertainties. Section 7 discusses the various ND strategies in more detail. Even after significant uncertainty reduction using ND measurements, neutrino interaction modeling is one of the leading sources of systematic uncertainties in long-baseline oscillation experiments.

5. FAR DETECTORS

FDs must provide a large target mass to accumulate sufficient statistics for precision measurements, and thus they must be reasonably cost-effective at large scales. They must be capable of effectively reconstructing and selecting ν_e and ν_μ CC interactions and, thus, both muons and electrons. They must accurately identify the vertex of the neutrino interaction to ensure a well-defined fiducial volume. They must be able to measure the kinematics of the event well enough to infer the energy of the incoming neutrino since oscillations depend on this quantity. Most FDs cannot differentiate neutrinos from antineutrinos on an event-by-event basis (though some usable statistical separation is often possible), so these experiments benefit from the purity of the neutrino and antineutrino beams.

Detection and reconstruction strategies vary among experiments; major experimental programs may be based on ring-imaging WCDs (Super-K, Hyper-K), different types of scintillator detectors (MINOS, NOvA), or LArTPCs (DUNE). In each case, the FDs are chosen to be well matched to the details of the experiments, as discussed below. **Figure 6** shows examples of how a neutrino interaction looks in each type of detector.

5.1. Water Cherenkov Detectors

Ring-imaging WCDs work by observing the Cherenkov light produced when charged particles move faster than the phase velocity of light in the medium. This light is emitted conically in a well-defined way such that photosensors placed on the wall of a tank of water will record rings of light from which the position, energy, and type of particle may be inferred.

Cherenkov light is emitted at an angle $\theta_c = 1/\beta n$, where n is the refractive index of the medium. Time of flight to the photosensors provides information on the position and trajectory of the

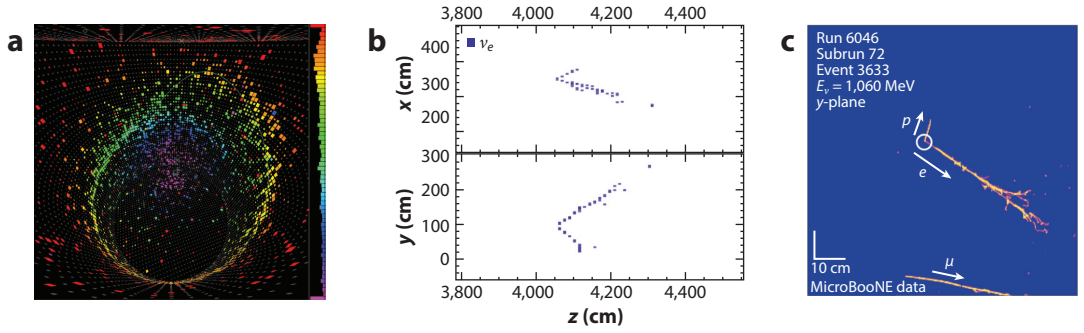

Figure 6

Typical ν_e interaction in the (*a*) Hyper-K water Cherenkov detector, (*b*) NOvA scintillator detector, and (*c*) MicroBooNE liquid argon time projection chamber. In each of the event displays, the colors correspond to the amount of energy observed in the detector. Panel *a* reproduced with permission from Reference 43. Panel *b* adapted with permission from Reference 44. Panel *c* adapted from Reference 45 (CC BY 4.0).

particle within the water volume, and the spatial pattern of the ring provides information on particle species and energy. These inferences are made using likelihood fits of various particle hypotheses to the observed light patterns. Muons and electrons produce distinctive ring shapes due to the differences in their propagation: The long straight muon tracks produce sharp rings, and the messy electromagnetic cascades from electrons lead to relatively fuzzy rings (see **Figure 6*a***). The Cherenkov threshold in Super-K is 0.73 MeV for electrons and 150 MeV for muons. Protons produced by the T2K neutrino beam are usually below the Cherenkov threshold and are not visible in the WCD. Because the hadronic side of the event is often unseen, the neutrino energy is inferred from the momentum of the observed lepton under an assumption of a two-body QE interaction off a stationary target, or a similar kinematic calculation for more complex final states.

WCDs may be enhanced by loading the water with gadolinium to increase the cross section for neutron capture. Neutron capture on gadolinium creates an 8-MeV cascade of γ-rays that allows identification of events with neutrons in the final state. In 2020, gadolinium was added to Super-K (46); while the primary motivation was for low-energy physics, adding gadolinium also enables neutrino–antineutrino discrimination in the long-baseline neutrino oscillation analysis because CC QE interactions have a proton in the final state for neutrinos and a neutron in the final state for antineutrinos.

A WCD has the advantage of allowing instrumentation of extremely large volumes (masses). Super-K has 50 kt of instrumented mass, and Hyper-K will be five times larger. For 40% photocoverage, typical of these detectors and sufficient for the long-baseline measurements described here, Hyper-K will require approximately 40,000 photomultiplier tubes.

5.2. Tracking Calorimeter Detectors

Muons at long-baseline energies are near minimum ionizing and have minimal radiative losses, giving rise to long straight tracks in the detectors. Tracking detectors can reconstruct such muon trajectories with excellent precision. Electrons at these energies produce electromagnetic cascades (called showers) with characteristic length scales given by the radiation length (X_0) of the material. Tracking detectors can produce detailed images of these showers from which one can infer the direction and energy of the original particles (for examples, see **Figure 6*b*,*c***). In general, reconstruction algorithms for tracking detectors attempt not to resolve individual particles within a shower but rather to reconstruct the shower as a single object. Showers initiated by a photon are

largely identical to those produced by a electron or positron, but they may be distinguished by a gap between the neutrino interaction vertex and the start of the shower, as a photon must travel a distance of order X_0 before initiating a shower. Another signature that differentiates electrons from photons is the energy deposition as a function of distance (dE/dx) in the early part of the shower, where the deposition corresponds either to a single minimum ionizing particle (MIP) in the case of an initial electron or positron or to two MIPs in the case of an electron–positron pair from an initial photon.

Muon energies in tracking calorimeters are best measured by range. Electron, positron, and photon energies are best measured calorimetrically via the total deposited energy in the electromagnetic shower. Other particles are reconstructed in various ways, though in practice energies of the rest of the event are often simply measured calorimetrically as well. For studying exclusive final states, dE/dx along a track provides a useful handle for picking out the highly ionizing protons, and track kinks or localized heavier depositions can be used to distinguish charged pions from the cleaner but otherwise similar-looking muons. Neutrons are difficult to detect as they may travel large distances in the detector before scattering, and they may deposit little visible energy when they do. Multiple Coulomb scattering can be used to make an estimate of the energy of muons that exit the detector, though the resolution is significantly worse than that for contained events measured by range.

5.2.1. Scintillator detectors. The NOvA detector (47) is a tracking calorimeter consisting of long cells of PVC filled with liquid scintillator. The cells are 3.9×6.6 cm in cross section, and they run the full transverse width or height of the detector; each layer of cells alternates between these orientations such that the position of a throughgoing track can be determined in three dimensions, as seen in **Figure 6*b***. Light produced by charged particles passing through a cell is collected using a wavelength-shifting fiber and read out by avalanche photodiodes at the edge of the detector. The spatial size of the cells was chosen to be much smaller than the radiation length (~50 cm) in the detector, allowing for excellent imaging of electromagnetic showers. Unique among current and planned long-baseline experiments, NOvA has functionally identical near and far detectors; the primary difference between the ND and FD is their size. This allows for additional analysis techniques when using the ND data to mitigate systematic uncertainties, as discussed below. The earlier MINOS experiment employed a similar strategy using magnetized steel-scintillator sampling calorimeters.

5.2.2. Liquid Argon Time Projection Chambers. LArTPCs consist of a volume (or array of similar volumes) of liquid argon with a strong and uniform electric field applied. When charged particles propagate through the liquid argon, they ionize the argon atoms to produce ionization electrons, which in turn drift in the electric field toward spatially granular readout instrumentation situated on one side of the volume. A two-dimensional image of the particle trajectories is thus recorded at the readout plane, and the third dimension is obtained from the relative arrival times of the drift electrons. Typical spacing of readout elements is 3–5 mm, with matching timing resolution for the drift dimension, so the spatial resolution of the resulting images is of a similar order. **Figure 6*c*** shows a ν_e interaction in an LArTPC. In addition to this imaging capability, the energy deposited by charged particles can be measured calorimetrically from the observed ionization charge.

Scintillation light is produced alongside the ionization signal, and this light can be detected promptly (relative to the much slower millisecond-scale drifting electrons) by a light detection system to determine the time (t_0) at which the interaction occurred. If the light collection system is suitably efficient, it can also be used for event calorimetry. The DUNE experiment uses the

LArTPC technology, as fits a longer-baseline experiment with correspondingly higher energies and higher-multiplicity final states.

6. NEAR DETECTORS

ND data allow significant reductions in systematic uncertainties from flux, cross sections, and detector response in long-baseline oscillation experiments. Various strategies for achieving these reductions are used, but all involve obtaining high-statistics ND samples of neutrino interactions differentiated in neutrino energy, flavor, interaction channel, and final-state kinematics. Because they are so close to the origin of high-power neutrino beams, NDs must be able to perform in a high-rate environment. Naive extrapolation of measurements from ND to FD is not possible because the flux at the two detectors is significantly different, mostly as a result of neutrino oscillations. These considerations mean that NDs must perform as well as, or often better than, their FD counterparts. Additionally, a broad suite of stand-alone measurements, including cross-section measurements, precision Standard Model measurements, and searches for new physics, is typically made using NDs in long-baseline experiments.

A given ND design is coupled strongly to the analysis techniques used. This section describes the detectors themselves, and Section 7 discusses the analysis techniques. Note that the categories need not be mutually exclusive.

6.1. Functionally Identical Near Detectors

When practical, experiments often include an ND or an ND component that is functionally identical to the FD. One example is the NOvA ND, which differs from the FD primarily in size (Section 5.2.1). This setup ensures that the target nuclei, event selection efficiencies, and energy resolutions are as similar as possible to those of the FD. The earlier MINOS experiment also used nearly identical ND and FD designs. DUNE's ND suite includes an LArTPC (see below) that functions similarly to the FD, though with more substantive differences than between NOvA's two detectors. In both the NOvA and DUNE cases, a downstream "muon catcher" is required at the ND to measure the momentum of muons that exit the back of the detector.

6.2. Multicomponent Detectors

In many cases, a single detector cannot meet all of the ND physics needs, so a suite of complementary detectors is employed. The T2K ND suite includes an on-axis iron-scintillator detector (INGRID) and an off-axis multicomponent detector (the ND280) consisting of an inner tracking region with time projection chambers (TPCs) interleaved with fine-grain detectors, an upstream π^0 detector, an electromagnetic calorimeter, and side muon detectors, all inside a magnetic field. This complex is currently being upgraded to reduce systematic uncertainty by approximately 30% relative to the T2K analysis for important flux and cross-section parameters (48). The DUNE ND suite consists of three components: the abovementioned LArTPC; a muon spectrometer, to be upgraded to a multipurpose detector such as a magnetized gaseous argon TPC (see Section 6.4); and a straw-tube tracker spectrometer (see Section 6.3).

6.3. Segmented Near Detectors

An important feature of the planned ND280 upgrade (48) is the Super-FGD, a fine-grained, fully active plastic scintillator detector. Super-FGD consists of approximately two million individual scintillator cubes, each $1 \times 1 \times 1$ cm^3 and coated in a reflecting layer for optical segmentation. Scintillation light from each cube is read out in three orthogonal directions using

wavelength-shifting fibers coupled to multipixel photon counters such that three-dimensional position reconstruction is possible. This design is expected to significantly improve capability to reconstruct low-energy particles and neutrons.

DUNE's SAND (System for On-Axis Neutrino Detection) (49) includes a straw-tube tracker consisting of low-mass straws (5 mm in diameter with 12-μm walls) arranged in alternating x and y layers. Thin, passive target layers are interspersed among the straws and account for most of the detector's mass, separating the target mass from the low-mass tracking system and allowing for the possibility of multiple target nuclei for additional cross-section studies. This inner tracker is surrounded by a magnet and calorimeter originally used in the KLOE (50) experiment. In addition to providing vital on-axis beam monitoring, SAND is capable of high-precision neutrino interaction measurements that may help to constrain models of neutrino interaction. One proposed study aims to infer CC interactions on a free proton by subtraction of samples interacting on the polypropylene (CH_2) and graphite (C) targets of the tracker (49).

6.4. Argon Near Detectors

DUNE's ND suite is designed around argon-based detectors to match the target nucleus of the FD. DUNE's ND-LAr (49) is an LArTPC with an operational principle similar to the FD's. Given the much higher event rates at the ND site, ND-LAr will consist of a number of smaller independent TPCs, reducing drift times and event pileup. Reconstruction (including $\mu^{\pm}$ differentiation) of muons that exit the back of ND-LAr will initially be provided by a downstream magnetized spectrometer. The proposed Phase II replacement for the spectrometer is ND-GAr (51), a magnetized high-pressure gaseous argon TPC surrounded by an electromagnetic calorimeter inside the magnet, and a muon system. While ND-GAr would serve as the muon catcher for ND-LAr, its low thresholds and low density should allow greatly improved particle tracking and identification capabilities for interactions on argon, enabling further constraints of neutrino scattering uncertainties and measurements of final states at a level of detail not possible with the higher-density LAr detector.

6.5. Movable Detectors

The two argon-based components of the DUNE ND will have the capability to move in tandem a distance of 30 m transverse to the beam, enabling measurements at off-axis angles ranging from 0 to 3.2° and facilitating the PRISM analysis described in Section 7.1.3. Similarly, the Hyper-K Collaboration plans to include a movable WCD as part of its ND suite. This detector will be located at an intermediate baseline of around 1 km and will cover an off-axis range of approximately 1–4°.

6.6. Additional Nuclear Targets

The π^0 detector (52) in T2K's ND280 detector is composed of layers of plastic scintillator alternating with water targets and brass sheets or lead sheets. The water targets can be filled with water or left empty. Data are collected in both modes; this allows for a subtraction analysis to infer the interaction rate on water alone, an example of which is described in Reference 53. Subtraction analyses are also proposed for DUNE's SAND.

7. ANALYSIS STRATEGIES

Neutrino oscillation parameters are extracted by comparing observed flavor-separated FD spectra to corresponding predicted spectra with varied assumptions for the parameters. In practice, this is done with detailed fitting algorithms designed to disentangle multiple oscillation and nuisance

parameters (e.g., from systematic uncertainties). Because detectors differ significantly from experiment to experiment, so too do the approaches for neutrino energy estimation, flavor tagging, and background rejection when forming the event spectra; however, the fundamental goals of these analysis components are the same across experiments. The core FD spectra of interest are ν_μ- and ν_e-selected events as a function of reconstructed neutrino energy. These samples may be further differentiated by final-state particle content, event classifier variables, or particle kinematics (e.g., the presence or absence of a charged pion in T2K's ν_e samples, or separation into quartiles of hadronic energy fraction in NOvA's ν_μ samples). These choices can lead to additional power in the oscillation parameter fits by reducing the statistical influence of backgrounds or of finite detector resolution or by recovering events otherwise rejected by a more basic approach.

7.1. Predicting Far Detector Spectra

Several conceptually different approaches have been developed for creating FD predictions. All use some mix of ND data and simulation tools to reduce the impact of systematic uncertainties to the greatest extent possible. Three leading approaches are sketched here. Next-generation experiments are being designed with a blend of approaches in mind.

7.1.1. Near detector fits. A direct way to incorporate ND constraints in the FD predictions is to explicitly fit a sophisticated simulation-based flux, cross section, and detector response model to high-statistics ND data samples. Taking T2K as an exemplar of this approach, the collaboration's most recent result uses 22 independent samples selected from the various ND280 subdetectors, separated into many exclusive final states (e.g., based on the presence of protons, pions, or photons) and binned across useful kinematic variables (e.g., lepton momentum and angle). This pool of data is fitted to constrain the experimental model parameters, and this constrained model with its correlated a posteriori systematic uncertainties is used in the fit to the oscillation-sensitive FD spectra. Additional systematic uncertainties not constrained by the ND fits are also present (e.g., FD-specific detector uncertainties). The T2K Collaboration has carried out analyses in which the ND and FD samples are fitted simultaneously as well as ones that take a two-step fitting approach. T2K's flux and cross-section uncertainties are reduced by approximately a factor of two in this process (54).

7.1.2. Extrapolation. The essentially identical construction of NOvA's ND and FD allows FD predictions to be formed rather directly from ND data. For example, the observed rates of ND ν_μ events in bins of neutrino energy, hadronic energy, and transverse lepton momentum are converted into expected rates for these events in the FD. The experiment's detailed flux, cross section, and detector response simulation is used to calculate this near-to-far translation, or "extrapolation," naturally folding in gross factors like distance from the beam source and detector mass as well as the many more subtle factors like reconstruction smearing, event containment differences, and the influence of neutrino oscillations themselves.

Unlike in the above ND fit procedure, systematic uncertainties on the simulation are not constrained by the ND data; they are left at their (large) a priori values. Instead, the influence of the uncertainties on the oscillation fit is significantly reduced simply because the data-based FD predictions are largely stationary under model variations since the model is used only to correct for near–far differences and not to form the FD predictions from scratch. The most recent NOvA oscillation analysis is described in Reference 17.

7.1.3. PRISM. As described in Section 3, the neutrino energy spectrum varies with off-axis angle: The spectrum becomes narrower and its peak energy becomes lower as the off-axis angle

increases. Data collected at different off-axis angles can therefore be used to disentangle or constrain flux and cross-section uncertainties in a powerful way.

In ND-fit-style analyses, off-axis samples can be included in the fits to provide better constraints and, importantly, to provide additional protection against inadequacies in the fitted model. For example, a cross-section model deficiency that might be unknowingly tuned away when fitting at a single detector position is unlikely to survive when confronted with data from multiple off-axis angles (i.e., data with distinct neutrino energy spectra). Tension in such ND fits would alert experimenters to model issues that could potentially influence the FD oscillation fits. As a proof of concept, DUNE conducted a sensitivity analysis assuming a hypothetical but plausible modeling error where the energy balance between recoiling protons and neutrons is incorrect. This ad hoc mismodeling was designed to keep on-axis ND data unchanged while inducing unacceptable oscillation fit biases. The off-axis data would, in this case, clearly reveal the presence of a modeling issue (49).

In an extrapolation-style analysis, the PRISM technique offers a novel and particularly effective reduction in model dependence. For example, the FD ν_μ energy spectrum looks very different from the corresponding ND spectrum due to oscillations. Thus, model uncertainties or deficiencies that influence energy reconstruction, such as the ad hoc proton–neutron energy variation described above, can bleed into extrapolation-based FD predictions. With PRISM, ND data sets are collected with multiple neutrino energy spectra. If these spectra span a wide enough range of energies and are sufficiently distinct from one another, it is possible to construct from them a new virtual ND neutrino energy spectrum that is similar to the oscillated spectrum at the FD for any choice of oscillation parameters by taking linear combinations of the off-axis spectra. This allows the creation of virtual ND data sets that have FD-like neutrino spectra, including the influence of oscillations.

In this way, FD predictions can be formed directly from ND data without suffering from spectral differences, eliminating the primary avenue by which energy-coupled model uncertainties enter the analysis. **Figure 7** shows this technique applied in the presence of the proton–neutron model variation described above.

7.2. Statistical Procedures

While details of statistical analysis procedures are beyond the scope of this article, we briefly describe common techniques used to interpret long-baseline oscillation data. The Particle Data Group review article on statistics (9) provides an excellent introduction to statistical analysis in particle physics.

Both frequentist and Bayesian statistical analysis procedures are used in the field for carrying out hypothesis tests or setting confidence or credible intervals for parameters of interest. Parameter fits are commonly performed using a binned Poisson likelihood ratio λ, typically cast in logarithmic form as follows:

$$-2\ln\lambda(\theta) = 2\sum_{i=1}^{N}\left[p_i(\theta) - d_i + d_i\log\frac{d_i}{p_i(\theta)}\right], \quad 5.$$

where the sum runs over the N data bins, d_i is the observed number of events in the ith bin, $p_i(\theta)$ is the predicted number of events in the ith bin, and θ is the set of free parameters in the fit, including any nuisance parameters related to systematic uncertainties. If any of the free parameters have a priori constraints, these constraints can be injected in a number of ways depending on the statistical methodology in use and the complexity of the prior distributions. In the simplest scenario, one can add penalty terms to the above likelihood ratio; an uncorrelated parameter θ_s with Gaussian prior uncertainty σ_s centered on mean $\hat{\theta}_s$ would lead to a penalty term of $(\theta_s - \hat{\theta}_s)^2/\sigma_s^2$.

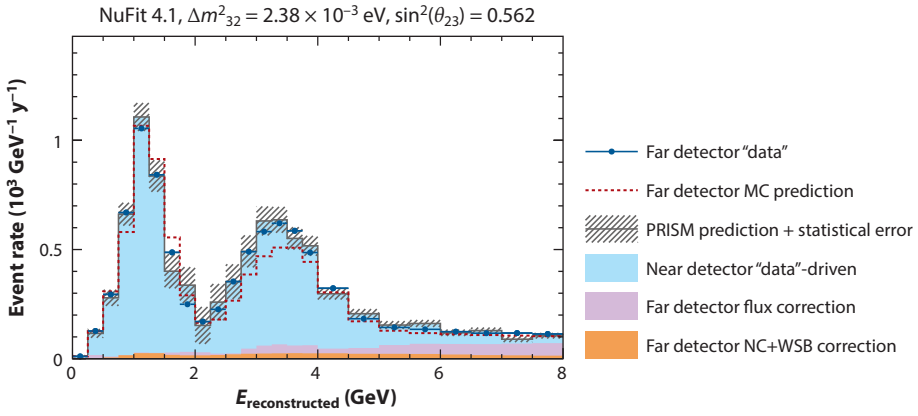

Figure 7

DUNE-PRISM ν_μ FD prediction for various oscillation parameters. The fake FD "data" (*solid dark blue lines* and *circles*) are generated using the proton–neutron model alteration described in Section 7.1.3. The unmodified, and thus incorrect, model prediction is labeled MC prediction (*dashed red lines*); it deviates clearly from the FD data. In contrast, the linear combination prediction (*solid gray lines* with *hatched error bars*) based on ND data correctly predicts the FD spectrum despite the presence of a large "unknown" cross-section modeling problem. The statistical error bars on the prediction reflect the limited size of the MC sample used in the study. The blue shaded region indicates the portion of the PRISM prediction that comes directly from the ND data, while the purple and orange shaded regions are additional corrections that complete the prediction. For technical details, see Reference 49. Abbreviations: FD, far detector; MC, Monte Carlo; NC, neutral current; ND, near detector; PRISM, precision reaction-independent spectrum measurement; WSB, wrong-sign background. Figure adapted from Reference 49 (CC BY 4.0).

If the conditions for Wilks's theorem (55) are met and the event counts are sufficiently large, then the minimum of $-2\ln\lambda$ follows a chi-squared distribution, which allows for practical simplifications in the statistical analysis. However, these conditions are very often not met in long-baseline experiments. In frequentist analyses, confidence intervals with correct coverage may be determined using the Feldman–Cousins unified approach (56), for which a large ensemble of pseudoexperiments is generated using Monte Carlo simulation to empirically determine the postfit likelihood ratio distributions across parameter space. This technique is computationally expensive but can often make substantial corrections relative to faster but approximate methods. In Bayesian analyses, Markov chain Monte Carlo methods are commonly used, in which an ensemble of points in parameter space is ultimately drawn according to the posterior probability density. Bayesian credible intervals are then built from this ensemble of points, as in, for example, Reference 57. Experimental collaborations sometimes report results from multiple styles of statistical analysis in parallel.

8. OTHER PHYSICS WITH LONG-BASELINE EXPERIMENTS

The detection capabilities of long-baseline neutrino oscillation experiments are well suited to a range of other physics topics, such that these experiments typically produce a broad range of physics results. While these physics topics are beyond the scope of this review, they greatly enhance the breadth of the physics programs in long-baseline experiments, so we mention them briefly here.

The ν_μ and ν_e CC samples that are used in the measurement of three-flavor oscillation parameters are, naturally, also sensitive to deviations from three-flavor oscillation due to BSM (beyond

the Standard Model) effects such as mixing with sterile neutrinos (10, 58), nonstandard interactions (59–61), and violations of *CPT* symmetry or Lorentz invariance (62–64). In some cases, these effects may modify the observed spectra in ways that are similar to variations of the three-flavor oscillation parameters; multiple long-baseline experiments with different baselines and energies can help resolve such potential degeneracies.

The massive FDs used in long-baseline experiments can provide significant sensitivity to baryon-number-violating processes such as proton decay (63, 65) and neutron–antineutron oscillation (66, 67), particularly when background is suppressed by locating the detectors deep underground. While such processes are predicted by many Grand Unified Theories, they have not yet been observed. The best limits for most nucleon decay modes have been set by the Super-K experiment (68, 69).

FDs in long-baseline experiments are also sensitive to neutrinos from natural sources, such as atmospheric and solar neutrinos. Atmospheric neutrinos (70), produced when cosmic rays interact in the Earth's atmosphere, have an energy spectrum that extends to very high energies and have a large range of azimuthal angles. The original discovery of neutrino oscillation was based in part on observation of atmospheric neutrinos, as described in Section 1. In future experiments, these samples are expected to provide complementary sensitivity to accelerator neutrinos for the parameters governing long-baseline oscillation; Hyper-K plans to use this sample for determination of the neutrino mass ordering. Study of atmospheric neutrinos is also a promising avenue for searches for BSM effects such as *CPT* and Lorentz violation.

Solar neutrinos (71) are sensitive to the two neutrino mixing parameters that are not accessible in long-baseline neutrino oscillation experiments and have historically played a large role in our understanding of neutrinos. In future experiments, solar neutrinos may be used to study the metallicity of our Sun, to provide input to models of stellar evolution, to probe the interaction of neutrinos with matter in the Sun, and to search for new physics effects, in addition to continuing the program of precision measurements of three-flavor neutrino oscillation.

Neutrinos from a core-collapse supernova (72, 73) act as messengers, providing a wealth of information about the physics of stellar collapse and explosion and, in turn, stellar evolution and nucleosynthesis, as well as providing information about neutrinos themselves. Next-generation long-baseline oscillation experiments would detect thousands of supernova burst neutrinos if such an event occurs in our Galaxy while the experiments are operational, with highly complementary samples detected by DUNE and Hyper-K. DUNE will be more sensitive to electron neutrinos via CC absorption of electron neutrinos on ^{40}Ar, while Hyper-K will be more sensitive to electron antineutrinos via inverse β decay. Neutrino experiments play an important role in multimessenger astronomy, participating in the early warning system for supernova bursts—SNEWS (74)—as neutrinos will be the first terrestrial sign of a supernova burst.

While NDs for long-baseline experiments are designed primarily to constrain systematic uncertainties for the oscillation analysis, these detectors also provide rich physics programs, including BSM searches and Standard Model measurements (49, 64, 75).

9. CONCLUSION

Long-baseline accelerator-based neutrino oscillation experiments are able to measure θ_{23}, Δm^2_{32}, θ_{13}, and δ_{CP} in a single experiment. These experiments observe an initial beam of predominantly muon neutrinos or antineutrinos at two detector sites, one near the beam source before significant oscillations have occurred and the other at a large FD sited hundreds of kilometers downstream. The oscillation channels most readily observed are $\nu_\mu/\bar{\nu}_\mu$ disappearance and $\nu_e/\bar{\nu}_e$ appearance, with sensitivity to θ_{13}, δ_{CP}, and the θ_{23} octant coming from the appearance channel.

To accumulate sufficient statistics at the FD, high-power beams and massive detectors are required. To make precision oscillation measurements, sophisticated NDs that are well matched to their FD counterparts are required to constrain systematic uncertainties from flux, neutrino interaction models, and detector effects. Modeling the complex interactions of neutrinos with nuclei, and understanding the impact of imperfections in this model, is a critical and challenging aspect of these measurements that requires careful design of detectors and analysis techniques as well as dedicated theoretical and experimental effort.

Currently operating experiments—NOvA and T2K—have developed sophisticated analysis techniques, have observed both disappearance and appearance signals, and have published measurements of Δm^2_{32}, $\sin^2\theta_{23}$, and δ_{CP}. However, the next generation of experiments—Hyper-K and DUNE—is necessary to definitively determine the neutrino mass ordering, to provide discovery sensitivity to *CP* violation, and to make precision measurements of the parameters governing neutrino masses and mixing. Beyond long-baseline oscillations, these experiments have broad physics programs with connections to low-energy neutrinos from astrophysical sources, BSM physics, and precision neutrino interaction physics.

SUMMARY POINTS

1. Long-baseline neutrino oscillation experiments detect $\nu_e/\bar{\nu}_e$ appearance and $\nu_\mu/\bar{\nu}_\mu$ disappearance in $\nu_\mu/\bar{\nu}_\mu$ beams, providing sensitivity to three-flavor oscillation parameters governing $\nu_1\nu_3$ and $\nu_2\nu_3$ mixing. The neutrino mass ordering, the value of the *CP*-violating phase δ_{CP}, and the θ_{23} octant are unknown and are being pursued by current and next-generation experiments.
2. Long-baseline oscillation measurements rely on an understanding of the neutrino flux, neutrino scattering cross sections, and detector effects that, a priori, have large associated systematic uncertainties. Sophisticated detector designs and analysis techniques are required to sufficiently mitigate systematic uncertainties.
3. Detectors for long-baseline oscillation experiments must provide sufficiently large target mass to collect the large numbers of neutrino events needed for precision measurements and must be sophisticated enough to allow detailed reconstruction for event classification and energy estimation. Details of detection, reconstruction, and analysis strategies vary among experiments and must be well matched to the neutrino flux at each location.
4. Near detectors are critical for constraining systematic uncertainties, and these detectors typically have more stringent performance requirements than far detectors.
5. Long-baseline neutrino oscillation experiments have broad physics programs beyond three-flavor oscillation measurements, including study of solar and supernova burst neutrinos, searches for physics beyond the Standard Model (BSM) (including baryon-number-violating processes), and neutrino–nucleus interaction measurements.

FUTURE ISSUES

1. Next-generation long-baseline oscillation experiments (Hyper-Kamiokande and DUNE) are being built in Japan and the United States. These experiments are highly

complementary in design with different flux, baseline, target nucleus, and detector technologies. Commitment from the international community is essential to realize both of these important experimental programs.

2. Joint fits among long-baseline experiments will be powerful in leveraging experimental complementarity and in differentiating between three-flavor and BSM signatures. As statistical uncertainties decrease with ever-growing data sets, significant effort will be required to fully understand and treat potential correlations in systematic uncertainties across experiments with differing designs and experimental details.
3. Neutrino–nucleus interaction models and related simulation tools are important for long-baseline neutrino oscillation experiments. Continued development of these models and tools, folding in current and future neutrino or other (e.g., electron) scattering data, requires collaboration among nuclear and particle theorists and experimentalists.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

R.B.P. and E.W. are supported by the US Department of Energy. Thanks to the NOvA, T2K, MicroBooNE, Hyper-K, and DUNE Collaborations for use of materials from their publications.

LITERATURE CITED

1. Esteban I, et al. *J. High Energy Phys.* 2009:178 (2020)
2. Fukuda Y, et al. (Super-Kamiokande Collab.) *Phys. Rev. Lett.* 81:1562 (1998)
3. Ahmad QR, et al. (SNO Collab.) *Phys. Rev. Lett.* 87:071301 (2001)
4. Ahmad QR, et al. (SNO Collab.) *Phys. Rev. Lett.* 89:011301 (2002)
5. Pontecorvo B. *Sov. Phys. JETP* 6:429 (1957)
6. Maki Z, Nakagawa M, Sakata S. *Prog. Theor. Phys.* 28:870 (1962)
7. Mena O, Parke S. *Phys. Rev. D* 69:117301 (2004)
8. Smirnov AY. *Phys. Scripta T* 121:57 (2005)
9. Workman RL, et al. (Part. Data Group) *PTEP* 2022:083C01 (2022)
10. Giunti C, Lasserre T. *Annu. Rev. Nucl. Part. Sci.* 69:163 (2019)
11. Agafonova N, et al. (OPERA Collab.) *Phys. Rev. D* 100:051301 (2019)
12. Jarlskog C. *Z. Phys. C* 29:491 (1985)
13. Adamson P, et al. (MINOS Collab.) *Phys. Rev. Lett.* 107:011802 (2011)
14. Acero MA, et al. (NOvA Collab.) *Phys. Rev. Lett.* 127:201801 (2021)
15. Abe K, et al. (T2K Collab.) *Phys. Rev. D* 99:071103 (2019)
16. Abe K, et al. (T2K Collab.) *Phys. Rev. D* 91:072010 (2015)
17. Acero MA, et al. (NOvA Collab.) *Phys. Rev. D* 106:032004 (2022)
18. Abud AA, et al. (DUNE Collab.) arXiv:2203.06100 [hep-ex] (2022)
19. Abi B, et al. (DUNE Collab.) *Eur. Phys. J. C* 80:978 (2020)
20. Abi B, et al. (DUNE Collab.) *J. Instrum.* 15:T08008 (2020)
21. Bian J, et al. (Hyper-Kamiokande Collab.) arXiv:2203.02029 [hep-ex] (2022)
22. Abe K, et al. (Hyper-Kamiokande Collab.) arXiv:1805.04163 [physics.ins-det] (2018)
23. Adamson P, et al. *Nucl. Instrum. Methods A* 806:279 (2016)
24. High-Intensity Proton Accel. Proj. Team. JAERI-Tech 2003-044/KEK Report 2002-13, Jpn. At. Energy Res. Inst., T kai, Ibaraki (2003)

25. Rakhno I, Mokhov N, Tropin I, Striganov S. Report FERMILAB-CONF-22-710-AD, Fermi National Accel. Lab., Batavia, IL (2022)
26. Duffy KE. *First measurement of neutrino and antineutrino oscillation at T2K*. PhD Thesis, Oxford Univ., Oxford, UK (2016)
27. Acero MA, et al. (NOvA Collab.) *Phys. Rev. D* 98:032012 (2018)
28. Park J, et al. (MINERvA Collab.) *Phys. Rev. D* 93:112007 (2016)
29. Marshall CM, McFarland KS, Wilkinson C. *Phys. Rev. D* 101:032002 (2020)
30. Mishra S. In *Proceedings of the Workshop on Hadron Structure Functions and Parton Distributions*, ed. D Geesaman, J Morfin, C Sazama, WK Tung, pp. 84–123. Singapore: World Scientific (1990)
31. Seligman WG. *A next-to-leading order QCD analysis of neutrino-iron structure functions at the Tevatron*. PhD Thesis, Nevis Lab., Columbia Univ., Irvington-on-Hudson, NY (1997)
32. Bodek A, Sarica U, Naples D, Ren L. *Eur. Phys. J. C* 72:1973 (2012)
33. Devan J, et al. (MINERvA Collab.) *Phys. Rev. D* 94:112007 (2016)
34. Sepulveda-Quiroz JA. *Measurement of the kaon production normalization in the NuMI target using uncontained charged-current muon neutrino interactions in the NOvA far detector*. PhD Thesis, Iowa State Univ., Ames (2018)
35. Casper D. *Nucl. Phys. B Proc. Suppl.* 112:161 (2022)
36. Formaggio JA, Zeller GP. *Rev. Mod. Phys.* 84:1307 (2012)
37. Andreopoulos C, et al. *Nucl. Instrum. Methods A* 614:87 (2010)
38. GiBUU. *The Giessen Boltzmann-Uehling-Uhlenbeck Project*. **https://gibuu.hepforge.org/** (2023)
39. Hayato Y, Pickering L. *Eur. Phys. J. ST* 230:4469 (2021)
40. Golan T, Juszczak C, Sobczyk JT. *Phys. Rev. C* 86:015505 (2012)
41. Mosel U. *Annu. Rev. Nucl. Part. Sci.* 66:171 (2016)
42. Campbell JM, et al. arXiv:2203.11110 [hep-ph] (2022)
43. Super-Kamiokande Collab. Research history. *Institute for Cosmic Ray Research*, *University of Tokyo*. **https://www-sk.icrr.u-tokyo.ac.jp/en/sk/about/history**
44. Calvez S, Catano-Mur E. *Fermilab News*, Nov. 7. **https://news.fnal.gov/2019/11/gotta-catch-em-all-new-nova-results-with-neutrinos-and-antineutrinos/** (2019)
45. Abratenko P, et al. (MicroBooNE Collab.) *Phys. Rev. D* 105:112003 (2022)
46. Abe K, et al. (Super-Kamiokande Collab.) *Nucl. Instrum. Methods A* 1027:166248 (2022)
47. Ayres DS, et al. (NOvA Collab.) Report FERMILAB-DESIGN-2007-01, Fermi National Accel. Lab., Batavia, IL (2007)
48. Abe K, et al. (T2L Collab.) Report CERN-SPSC-2019-001/SPSC-TDR-006, CERN, Geneva (2019)
49. Abed Abud A, et al. (DUNE Collab.) *Instruments* 5:31 (2021)
50. Sciascia B. (KLOE Collab.) *Nuovo Cim. C* 033N5:221 (2010)
51. Abed Abud A, et al. (DUNE Collab.) arXiv:2203.06281 [hep-ex] (2022)
52. Assylbekov S, et al. *Nucl. Instrum. Methods A* 686:48 (2012)
53. Abe K, et al. (T2K Collab.) *Phys. Rev. D* 91:112010 (2015)
54. Abe K, et al. (T2K Collab.) *Phys. Rev. Lett.* 124:161802 (2020)
55. Wilks SS. *Ann. Math. Stat.* 9:60 (1938)
56. Feldman GJ, Cousins RD. *Phys. Rev. D* 57:3873 (1998)
57. Abe K, et al. (T2K Collab.) *Phys. Rev. D* 96:092006 (2017). Erratum. *Phys. Rev. D* 98:019902 (2018)
58. Gariazzo S, et al. *J. Phys. G* 43:033001 (2016)
59. Davidson S, Pena-Garay C, Rius N, Santamaria A. *J. High Energy Phys.* 0303:011 (2003)
60. Gonzalez-Garcia MC, Maltoni M. *Phys. Rep.* 460:1 (2008)
61. Biggio C, Blennow M, Fernandez-Martinez E. *J. High Energy Phys.* 0908:090 (2009)
62. Argüelles CA, et al. arXiv:2203.10811 [hep-ph] (2022)
63. Dev PSB, et al. arXiv:2203.08771 [hep-ex] (2022)
64. Martin-Albo J. (DUNE Collab.) *Proc. Sci.* NuFact2021:147 (2022)
65. Nath P, Fileviez Perez P. *Phys. Rep.* 441:191 (2007)
66. Mohapatra RN, Marshak RE. *Phys. Rev. Lett.* 44:1316 (1980). Erratum. *Phys. Rev. Lett.* 44:1643 (1980)
67. Phillips DG II, et al. *Phys. Rep.* 612:1 (2016)

68. Takenaka A, et al. (Super-Kamiokande Collab.) *Phys. Rev. D* 102:112011 (2020)
69. Matsumoto R, et al. (Super-Kamiokande Collab.) *Phys. Rev. D* 106:072003 (2022)
70. Kajita T. *Annu. Rev. Nucl. Part. Sci.* 64:343 (2014)
71. Gann GDO, Zuber K, Bemmerer D, Serenelli A. *Annu. Rev. Nucl. Part. Sci.* 71:491 (2021)
72. Mirizzi A, et al. *Riv. Nuovo Cim.* 39:1 (2016)
73. Horiuchi S, Kneller JP. *J. Phys. G* 45:043002 (2018)
74. Antonioli P, et al. *New J. Phys.* 6:114 (2004)
75. Branca A, et al. *Symmetry* 13:1625 (2021)

Annual Review of Nuclear and Particle Science

Volume 73, 2023

Contents

Lepton Flavor Violation and Lepton Flavor Universality Violation in *b* and *c* Decays
Diego Guadagnoli and Patrick Koppenburg 1

New Solutions to the Gauge Hierarchy Problem
Anson Hook 23

COHERENT at the Spallation Neutron Source
P.S. Barbeau, Yu. Efremenko, and K. Scholberg 41

Experimental Considerations in Long-Baseline Neutrino Oscillation Measurements
Francesca Di Lodovico, Ryan B. Patterson, Masato Shiozawa, and Elizabeth Worcester 69

Detection and Calibration of Low-Energy Nuclear Recoils for Dark Matter and Neutrino Scattering Experiments
Jingke Xu, P.S. Barbeau, and Ziqing Hong 95

Implications of Large-N_c QCD for the *NN* Interaction
Thomas R. Richardson, Matthias R. Schindler, and Roxanne P. Springer 123

The Hubble Tension and Early Dark Energy
Marc Kamionkowski and Adam G. Riess 153

High-Energy to Ultrahigh-Energy Neutrino Interactions
Mary Hall Reno 181

Medium Response and Jet–Hadron Correlations in Relativistic Heavy-Ion Collisions
Shanshan Cao and Guang-You Qin 205

Boson–Boson Interactions at the LHC
J. Manjarrés Ramos and Guillelmo Gómez-Ceballos 231

Physics of the Top Quark at the LHC: An Appraisal and Outlook of the Road Ahead
P. Ferreira da Silva 255

Recent Progress in Leptonic and Semileptonic Decays of Charmed Hadrons
Bai-Cian Ke, Jonna Koponen, Hai-Bo Li, and Yangheng Zheng 285

The *s* Process and Beyond
Maria Lugaro, Marco Pignatari, René Reifarth, and Michael Wiescher 315

Ultra-High-Energy Gamma-Ray Astronomy
Zhen Cao, Songzhan Chen, Ruoyu Liu, and Ruizhi Yang 341

Deep-Sea and Lunar Radioisotopes from Nearby Astrophysical Explosions
Brian D. Fields and Anton Wallner 365

Physics Beyond the Standard Model Associated with the Top Quark
Roberto Franceschini 397

A Guide to Hunting Long-Lived Particles at the LHC
Simon Knapen and Steven Lowette 421

Errata

An online log of corrections to *Annual Review of Nuclear and Particle Science* articles may be found at http://www.annualreviews.org/errata/nucl