

Enabling Neutrino and Antineutrino Appearance Observation Measurements with HPC Facilities

Norm Buchanan

Department of Physics
Colorado State Univeristy
Fort Collins, CO USA
norm.buchanan@colostate.edu

Steven Calvez

Department of Physics
Colorado State Univeristy
Fort Collins, CO USA
steven.calvez@colostate.edu

Pengfei Ding

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
dingpf@fnal.gov

Derek Doyle

Department of Physics
Colorado State Univeristy
Fort Collins, CO USA
dddoyle@colostate.edu

Christopher Green

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
greenc@fnal.gov

Alexander Himmel

Neutrino Division
Fermi National Accelerator Laboratory
Batavia, IL USA
ahimmel@fnal.gov

Burt Holzman

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
burt@fnal.gov

James Kowalkowski

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
jbk@fnal.gov

Andrew Norman

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
anorman@fnal.gov

Marc Paterno

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
paterno@fnal.gov

Saba Sehrish

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
ssehrish@fnal.gov

Alexandre Sousa

Department of Physics
University of Cincinnati
Cincinnati, OH USA
sousaa@ucmail.uc.edu

Brandon White

Scientific Computing Division
Fermi National Accelerator Laboratory
Batavia, IL USA
bjwhite@fnal.gov

Abstract—When fitting to data with low statistics and near physical boundaries, extra measures need to be taken to ensure proper statistical coverage. The method $\text{NO}\nu\text{A}$ uses is called the Feldman-Cousins procedure, which entails fitting thousands of independent pseudoexperiments to generate acceptance intervals that are then used to correct our fits. The scale required by the Feldman-Cousins procedure makes it extremely computationally intensive. In past analyses, it has taken up to several weeks to complete, bottlenecking our final results. We will present recent work by members of the NOA experiment and the SciDAC4 collaboration that utilized the super-computing facilities at NERSC to process our Feldman-Cousins corrections over 50x faster, allowing us to perform a greater number of studies while increasing the precision of our fits.

I. NEUTRINO AND ANTINEUTRINO OSCILLATIONS

Neutrinos have been observed to change, or oscillate, between their quantum mechanical states of lepton flavor, governed by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing formalism. $\text{NO}\nu\text{A}$, a long-baseline neutrino oscillation experiment designed to measure both ν_e appearance and ν_μ disappearance, has collected extensive data on both neutrino (antineutrino) mixing for ν_e ($\bar{\nu}_e$) and ν_μ ($\bar{\nu}_\mu$) channels. Com-

binning these data provides strong constraints on the mass hierarchy and mixing parameters of the PMNS matrix.

To perform these measurements immense, peta-scale, datasets taken by the $\text{NO}\nu\text{A}$ detector need to be examined and specific candidate neutrino interactions within the data are selected according to a complex set of computations and criteria. After selecting the neutrino interactions and examining their properties, we determine the PMNS mixing parameters. Data-driven predictions under different parameter sets that maximize a likelihood function are compared with the observed data to determine confidence intervals for the observation.

II. $\text{NO}\nu\text{A}$ EVENT SELECTION

The first steps of these analyses involve the identification and selection of candidate electron neutrino (ν_e) and muon neutrino (ν_μ) events that are observed in the detectors. These events need to be chosen from the massive parent datasets that include all of the temporal and geometric reconstructions of particle interactions captured by the detector readout systems. These data are represented as time series, which are “sliced” into time windows of interest, and are referred to in the

experiment as *slices*. These regions are identified through an algorithmic procedure akin to image recognition, but tuned to the specific topologies of neutrinos interacting with different nuclear structures. For the data used in the 2017/2018 long baseline oscillation analyses, the size of this refined “candidate neutrino” data set is ≈ 25 TiB. This represents a size approximately 400 times smaller than the parent dataset, yet one that still contains the large and complex body of information regarding each candidate event that is needed for analysis.

We have reorganized these data from their original ROOT [1] formatted n-tuple representation, into a single HDF5 [2] file containing ≈ 80 tables. As a *slice* is a fundamental unit of processing in the NO ν A analysis, we present a mapping of data to computation that ensures that each slice is processed by a single MPI rank. We have provided a library-level solution to enable such a mapping in the read operation. After each rank has mapped and read the slice-based data it needs, there is no additional communication needed among the ranks.

The computations that are carried out on each rank are implemented using numpy [3] and pandas [4]. Using HPC resources at Argonne Leadership Computing Facility (ALCF) and National Energy Research Scientific Computing (NERSC), we will compare the performance and scalability of our implementation against the current implementation that uses thousands of small files and the ROOT analysis framework that is prevalent in the HEP community.

III. FELDMAN-COUSINS METHODS

After the selection of the different neutrino interaction events, we create and fit the resulting neutrino energy spectra and determine the PMNS mixing parameters along with their associated significances. Calculation of the significance of the measurement requires knowledge of the distribution describing the relevant test statistic. NO ν A uses a log-likelihood ratio ($\Delta\chi^2$) that should follow a χ^2 distribution under certain circumstances, namely sufficiently high statistics and unbounded parameters [5]. However, the elusive nature of the neutrino makes the observation of neutrino events very rare, and the oscillation parameters described in the PMNS formalism are inherently bounded. Consequently, NO ν A must generate its own empirical $\Delta\chi^2$ distributions following the Feldman-Cousins procedure [6] to calculate significances.

The Feldman-Cousins procedure is a non-parametric Monte Carlo method for drawing confidence intervals under non-Gaussian situations. It involves generating “pseudoexperiments” and performing fits to statistically fluctuated fake data sets. Repeating this process many times at a given set of parameters builds up a distribution of $\Delta\chi^2$ with which a significance of observation can be calculated. This is done at every point in the physics model’s parameter space so that confidence intervals can be constructed with the correct statistical coverage, but it also means that this process is *extremely* computationally intensive.

IV. THE ROLE OF HIGH PERFORMANCE COMPUTING

The Feldman-Cousins procedure can only be performed after all other stages of the analysis, making it the bottleneck for previous NO ν A physics results. Furthermore, the addition of antineutrino data increased the dimensionality and complexity of our fitting, demanding more computing resources than ever before. If NO ν A was to use traditional means to process these corrections on the Fermigrid, the latest analysis would have taken more than five months to complete. Instead, we worked to port the existing infrastructure onto high performance computing platforms like those at NERSC.

The machines we used at NERSC, CORI-1 and CORI-2, were critical to the computation of NO ν A results. Over 33 million hours of computation were performed on these machines during two dedicated runs of 16 and 36 hours on May 7th and 25-26th, 2018. The use of NERSC HPC facilities enabled us to review the results less than six hours, which were then shown at Neutrino 2018 on June 4th, 2018, a feat that would not have been possible otherwise.

V. ADVANCED DOMAIN DECOMPOSITION WITH MPI

Working toward our future goal of increasing the efficiency and degree to which we can parallelize the Feldman-Cousins procedure, we have transitioned the decomposition of the Feldman-Cousins fitting domain to use the DIY block-parallel environment and tools. The move to DIY allows for automated mapping of DIY blocks into MPI ranks for maximum parallelism in HPC settings. We have also begun development of a multi-threaded fitting framework to minimize the memory footprint per fit and maximize usage of cores and threads on vector platforms like KNL. This requires moving away from the ROOT framework to a combination of DIY/MPI and parallelism optimized C++ fitting libraries, enabling us to take full advantage of vectorization. Full utilization of these HPC machines will allow us to investigate the detailed effects of event selection on sensitivities with greater precision and enable greater physics opportunities on NO ν A and other neutrino experiments.

REFERENCES

- [1] R. Brun and F. Rademakers, “ROOT: An object oriented data analysis framework,” *Nucl. Instrum. Meth.*, vol. A389, pp. 81–86, 1997.
- [2] “HDF5,” <http://hdf.ncsa.uiuc.edu/HDF5>.
- [3] “Numpy,” <http://www.numpy.org>.
- [4] W. McKinney, “pandas: a Foundational Python Library for Data Analysis and Statistics,”
- [5] S. S. Wilks *Ann. Math. Statist.*, 1938.
- [6] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Physical Review D*, 1998.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program, grant 1013935. This material is based upon work supported by the Fermilab Laboratory Directed Research and Development program, LDRD2016-010.