




Enabling Neutrino & Antineutrino Appearance Observation Measurements with HPC Facilities

N. Buchanan¹, S. Calvez¹, P.F. Ding², D. Doyle¹, C. Green², A. Himmel², B. Holzman², J. Kowalkowski², A. Norman², M. Paterno², S. Sehrish², A. Sousa³, B. White²

¹Colorado State University, ²Fermi National Accelerator Laboratory, ³University of Cincinnati

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.



NOvA Experiment & Neutrino Oscillation Physics

NOvA is an accelerator-based long-baseline neutrino oscillation experiment designed to measure both ν_e appearance and ν_μ disappearance in order to determine the neutrino mass hierarchy, improve knowledge of the neutrino mixing structure, and probe the CP violating phase δ_{CP} [1].

The experiment uses a multidetector design where both NOvA detectors use a functionally identical design and differ only in the their size and locations.

The experiment features:

- A 300 ton near detector located underground, approximately 1 km from the neutrino beam production target, positioned 14 mrad off the beam axis.
- A massive 14 kt far detector located in Ash River, Minnesota, 810 km and 14 mrad off axis from the primary target station.
- The detectors sample the 700 kW NuMI neutrino beam generated at Fermilab using 120 GeV protons.




FIG. 1: Candidate neutrino interaction event in the NOvA detector. The event was classified as a ν_μ charged-current interaction.



Neutrino and Antineutrino Oscillations

- Neutrinos have been observed to change, or oscillate, between their quantum mechanical states of lepton flavor.
- The process through which this occurs is described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing formalism and the associated mixing matrices:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix} \begin{pmatrix} \cos \theta_{13} & 0 & \sin \theta_{13} e^{-i\delta} \\ 0 & 1 & 0 \\ \sin \theta_{13} e^{i\delta} & 0 & \cos \theta_{13} \end{pmatrix} \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} & 0 \\ -\sin \theta_{12} & \cos \theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}$$

- NOvA has extensive data on neutrino mixing for ν_e and ν_μ channels
- In 2017 and 2018 NOvA took data with **antineutrinos** to study the effect of mixing on the antiparticle states [2,3].
- Combining neutrino and antineutrino data provides strong constraints on the mass hierarchy and mixing parameters.
- Complicated fits to the data are required to extract the physics content.




FIG. 2: Neutrino/antineutrino oscillations bi-probability ellipses for NOvA data with 2018 best fit solution. Blue/red ellipses are for the Normal/Inverted neutrino mass hierarchy. Solutions for the upper ($\theta_{23} > \pi/4$) and lower octant ($\theta_{23} < \pi/4$) are shown.

Neutrino Oscillation Measurements




FIG. 3: Simulation of the unoscillated reconstructed neutrino energy spectrum (green) and of the observed muon neutrino spectrum in the NOvA far detector (blue). Comparison between the spectra allows extraction of oscillation parameters.

- Use measured near detector neutrino spectra to predict the far detector spectra before and after convolution with an oscillation model.
- Fit for the oscillation model parameters using a likelihood method and determine the confidence intervals for the measurement.




FIG. 4: Rejection significance under the Gaussian statistical assumption for different neutrino mass hierarchy solutions as a function of $\sin^2 \theta_{23}$. Profiling is performed over other oscillation and nuisance parameters.

Statistical Corrections




FIG. 5: Sensitivity contours under Gaussian statistics assumptions compared to a Feldman-Cousins corrected computation. The corrected contours can reveal large islands in the parameter space where sensitivity is greatly improved.




FIG. 6: Feldman-Cousins correction surfaces in units of $\Delta\chi^2$ for the 2σ (95%) confidence bands in δ_{CP} and Δm^2_{32} . The 1, 2, and 3 σ Gaussian contours (dashed lines) are shown along with the statistically corrected contours (solid lines).

Multi-Universe Methods & Feldman-Cousins

Why do we need statistical corrections?

- Normally we can use a χ^2 distribution to draw confidence intervals using Wilk's theorem for statistical ensembles of measurements.
- Neutrino oscillation measurements violate several assumptions of Wilks theorem (low statistics and bounded parameters).

What is Feldman-Cousins [4]?

- A non-parametric Monte Carlo method for drawing confidence intervals under non-Gaussian situations.
- "Pseudoexperiments" are generated and fit to build up empirical distributions that replace the standard χ^2 .
- This is done at every point in the physics model's parameter space.




FIG. 7: Comparison of the χ^2 distribution under Gaussian assumptions for 2 degrees of freedom with the ensemble of pseudoexperiments for one bin of the NOvA oscillation measurements [4].

- The Feldman-Cousins procedure can both increase and decrease a given measurement's sensitivity to different parts of the physics parameter space.




FIG. 8: Feldman-Cousins correction surfaces for 90% C.I. comparing $\chi^2(2)$ distribution and NOvA pseudo-experiments multi-universes in Δm^2_{32} and $\sin^2 \theta_{23}$. Red values indicate higher sensitivity and blue indicate lower sensitivity deviation.

- Each point in parameter space requires thousands of pseudoexperiments and fits, which makes this procedure **extremely computationally intensive**.




FIG. 9: $\Delta\chi^2$ distribution as a function of $\sin^2 \theta_{23}$ for over 6 million individual pseudo-experiments. P-value contours for 0.68 and 0.95 are shown in blue/red.

NOvA Fully-Corrected Analysis Results




FIG. 11: The Feldman-Cousins corrected results for NOvA provide world-leading constraints on the oscillation model for neutrinos and antineutrinos. They favor the Normal mass hierarchy and the θ_{23} upper octant solution of PMNS mixing [6].

The Role of High-Performance Computing

- Feldman-Cousins fits/corrections can only be performed after all other stages of the analysis.
- This has been the bottleneck for previous NOvA physics results.
- Antineutrino data increased the dimensionality and complexity to our parameter estimations.
- Running the 2018 fits and corrections with traditional grid computing at FNAL, even with resources from the Open Science Grid, would have taken 5+ months to complete.

NERSC Supercomputing Center

The CORI-1, CORI-2 and Edison supercomputers at NERSC were critical to the computation of the NOvA results. Over 33 million hours of computation were performed on these machines during two dedicated runs of 16 & 36 hours on May 7th and 25-26th, 2018. Results were in review less than 6 hours later and were shown at **Neutrino 2018** on June 4th, 2018.

HEPCloud: A window to HPC

The **HEPCloud** systems allows physicists to transparently provision resources across diverse sets of facilities.

NOvA used HEPCloud [7] to provision the largest HTCondor pool ever run, at over 1 million CPU cores. These seamlessly-established queues kept all the cores at NERSC busy with the different sets of NOvA fits and computations, and minimized the development needed to port the NOvA workflows to the HPC machine architectures.




FIG. 10: Over a million cores on CORI and Edison were provisioned and run through the **HEPCloud** facility in two analysis runs in May 2018.

Advanced Domain Decomposition with MPI

- Have transitioned decomposition of Feldman-Cousins scan and fitting domains to use the DIY block-parallel environments and tools [8].

- Allows for automated mapping of DIY blocks into MPI ranks for maximum parallelism in HPC setting.

- Internal to each block have begun development of multi-threaded fitting framework to minimize memory footprint per fit and maximize usage of cores/threads on KNL platforms.

- Requires moving away from ROOT framework to combination of DIY/MPI and parallelism optimized C++ fitting libraries. This will enable us to take full advantage of vectorization hardware and KNL platforms.




FIG. 12: Blockwise decomposition of the fitting domain hyperspace for one of the NOvA correction surfaces, using the DIY scheme. The DIY software maps these blocks into MPI ranks which are then run on the HPC platform.

Changing the HEP Data and Analysis Model

The future of computing focuses on leveraging the HPC centers. However, current HEP data models and analysis workflows are not designed to run optimally in a massive parallel environment.

Through the SciDAC-4 HEP Data Analytics on HPC project, we have modified both our **data model** and **analysis workflows** to take full advantage of HPC systems and modern libraries with implicit parallelism. Our new data/analysis model switches from using the HEP standard ROOT IO and TTree model [9] to using a HDF5 [10] table representations of the data.

In this HDF5 model, we have:

- Minimized *reading, communication* and *synchronization* between processes.
- Implicit Parallelism**: user-written code looks just like serial code.
- All data processed for a given "slice" (neutrino interaction) in a single MPI rank.
- Enabled organizing the data into a single HDF5 file, containing many different tables, some *regular* (one entry per slice), others *irregular*.




FIG. 13: Blockwise organization of NOvA data into regular (left) and irregular (right) length records. This permits the indexing of information across different neutrino events (shown by different colors) and the mapping of the data into the different ranks of the calculation.

Distributing and reading information

- Eliminates communication between ranks during "slice-by-slice" selection.
- Results in perfect data parallelism in running the full selection code.
- Each rank first reads its "fair share" of index info from a regular table.
 - Reads are overlapped to identify event boundaries
 - Assignment of ranks to events for even computational balance
- Event "ownership" information distributed to all ranks.
- Index information is read from *irregular* tables. Allows calculation of entries to be read by each rank, based on event "ownership".
- Each rank reads *only* relevant rows and columns from relevant tables.
 - ALL relevant data read by some rank.
 - NO rank reads the same data as another rank.

Demonstrated Perfect Strong Scaling

To demonstrate the IO performance of HDF5 for neutrino data analysis, we have measured the scaling and speed of reading and analyzing liquid argon Time Projection Chamber (TPC) data that was acquired by the LArIAT experiment at Fermilab [11].

- In these tests we merged all the data taken from the LArIAT experiment into a single HDF5 file. This represented 42 TiB of LAr-TPC which was compressed to a 4 TiB disk-resident file.
- Using a Python/MPI program on Cori Phase II at NERSC, we read and performed a Fast Fourier Transform (FFT) based noise reduction algorithm on the data.
- We observed perfect strong scaling up to 76,800 ranks**
- When reading from the Cori Burst Buffer, the program read and decompressed the entire dataset in less than 20 seconds




FIG. 14: Scaling behavior for event processing speed (left) and read plus decompression speed (right) observed using HDF5 representation of LArIAT TPC data. Blue/magenta circles show performance with 5/10 TiB burst buffer allocation. Results show perfect scaling up to 76,800 ranks in the MPI calculations.

Example tables from the HDF5 file

run	subRun	event	slice	distallpntop	...35 more...
16433	61	356124	35	nan	
16433	61	356124	36	-0.74013746	
16433	61	356124	37	nan	
16433	61	356125	1	nan	
16433	61	356125	2	423.6337	
16433	61	356125	3	-2.849864	

run	subRun	event	slice	vtxid	tpng3d	...6 more...
16433	61	356124	35	0	0	
16433	61	356124	36	0	1	
16433	61	356124	36	1	1	
16433	61	356124	36	2	5	
16433	61	356125	1	0	1	
16433	61	356125	3	0	0	

TABLE 1: Organization of NOvA data records into regular table (top) with one entry per neutrino slice and into an irregular table (bottom) with one entry per identified vertex in the interaction topology.

User Analysis Code Examples

```
def kNuSecondAnaContainment(tables):
    df = tables['sel_nuecosrej']
    return (df.distallpntop > 63.0) & \
           (df.distallpntop > 12.0) & \
           (df.distallpntop > 12.0) & \
           (df.distallpntop > 12.0) & \
           (df.distallpntop > 18.0) & \
           (df.distallpntop > 18.0)
```

```
def vtxelasticCut(tables):
    df = tables['vtx_elastic']
    df['good'] =
        (df.vtxid == 0) & (df.pnng3d > 0)
    KL = ['run', 'subRun', 'event', 'slice']
    return df.groupby(KL).agg(np.any)
```

LISTING. 1 Example analysis code demonstrating the selection of neutrino events within the detector's fiducial volume using regular table structure (top) and the selection of neutrino events with at least one 3-dimensional vertex in the interaction topology using an irregular table structure (bottom). Both types of selection are intrinsically parallel and run over the entire NOvA candidate event data with a maximum number of ranks determined by the number of candidate interactions. This gives a potential scaling to $O(10^{12})$ ranks.




FIG. 15: Event display of charged particle interactions in the LArIAT detector, courtesy of the LArIAT collaboration.




FIG. 16: File size distribution for all files at NERSC. From "Planning for the Future of Data, Storage and I/O at NERSC", Glenn Lockwood (2018).

The Future of HEP Neutrino Analysis

The tools being developed in collaboration with NOvA and the SciDAC-4 HEP on HPC project are designed to enable a new paradigm for neutrino analysis that is focused on current HPC facilities and targeted for future exascale systems.

This new model enables ALL stages of the neutrino analysis to occur within the HPC environment, and enables for the first time the capability to iteratively "tune" analyses and see the impact of changes on the final physics sensitivities.

This model leverages our HDF5 neutrino data format to combine machine learning, massive parallel event selection, advanced fitting, and automated domain decomposition to optimize and shorten the time to obtain physics results.

Full utilization of the next generation of HPC machines will enable physicists to probe with greater precision and higher sensitivity the hidden secrets of the neutrino world.




FIG. 17: Schematic flow of future neutrino analysis workflow on HPC platforms, permitting iterative exploration of complex parameter spaces.

References:

- [1] Ayres, D. et al. "The NOvA Technical Design Report" (2007) doi: [10.2172/935497](https://doi.org/10.2172/935497).
- [2] P. Adamson et al. (NOvA Collaboration) Phys. Rev. Lett. 118 (2017)
- [3] Acero, M.A. et al. (NOvA Collaboration) Phys. Rev. D98 (2018).
- [4] G. Feldman, R. Cousins. doi:[10.1103/PhysRevD.57.3873](https://doi.org/10.1103/PhysRevD.57.3873)
- [5] Sousa, A. et al. CHEP 2018 (2018)
- [6] Sanchez, M. (2018) Zenodo. doi:[10.5281/zenodo.1286758](https://doi.org/10.5281/zenodo.1286758).
- [7] Fuess, S. et al. J. Phys. Conf. Ser. 898 (2017).
- [8] Peterka, T. et al. in proceedings of LNAV '11 (2011)
- [9] R. Brun, F. Rademakers. doi:[10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X).
- [10] HDF5 Group, "The HDF5 Proposal", ESE-RFC-007 (2006)
- [11] Paley, J. et al. arXiv:1406.5560 (2014).



U.S. DEPARTMENT OF
ENERGY

Office of
Science


SciDAC
Scientific Discovery
through
Advanced Computing



Colorado
State
University



University of
CINCINNATI



SC18
Dallas, TX
hpc
inspires