# Analyzing NOvA Neutrino Data with the Perlmutter Supercomputer


N. Buchanan[1], Steven Calvez[1], **Derek Doyle**[1], V Hewes[2], Alexander Himmel[3], James Kowalkowski[3], Andrew Norman[3], Marc Paterno[3], Tom Peterka[4], Saba Sehrish[3], Alexandre Sousa[2], **Tarak Thakore**[2], Orcun Yildiz[4]

[1]Colorado State University, Fort Collins, CO USA. [2]University of Cincinnati, OH USA. [3]Fermi National Accelerator Laboratory, Batavia IL USA., Batavia IL USA. [5]Argonne National Laboratory, Argonne, IL USA.


## Neutrino Physics

Neutrinos are a fundamental particle of the universe. They are emitted via beta-decay by everyday sources like bananas, intergalactic stellar sources, and are a byproduct of uranium enrichment. Since their discovery in 1956, we have confirmed their non-zero mass by observing a quantum mechanical phenomenon called "oscillation", and have developed a number of practical applications including multi-messenger astronomy and nuclear anti-proliferation.

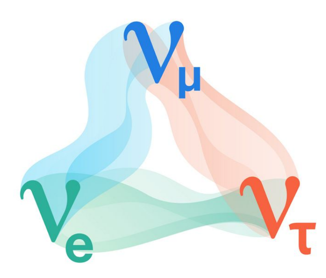


### Oscillations

$$1 - P(\nu_\mu \to \nu_s) \approx 1 - \cos^4\theta_{14}\cos^2\theta_{34}\sin^2 2\theta_{24}\sin^2\Delta_{41} - \sin^2\theta_{34}\sin^2 2\theta_{23}\sin^2\Delta_{31} + \frac{1}{2}\sin\delta_{24}\sin\theta_{24}\sin 2\theta_{23}\sin\Delta_{31}$$

$$P(\nu_\mu \to \nu_\mu) \approx 1 - \sin^2 2\theta_{24}\sin^2\Delta_{41} + 2\sin^2 2\theta_{23}\sin^2\theta_{24}\sin^2\Delta_{31} - \sin^2 2\theta_{23}\sin^2\Delta_{31}$$

An extension to the PMNS model, a 3+1 flavor state, describes the "Beyond the Standard Model" (BSM) scenario where oscillations occur between an additional "sterile" flavor state, which cannot interact with matter. This scenario is currently one of the most intensely studied area of BSM physics.

- Neutrinos are found in one of three states that can interact with matter: electron flavor, muon flavor, or tau flavor.
- As they travel through space, neutrinos oscillate between states based on distance traveled and mass.
- The 3-flavor PMNS formalism describes this oscillation as rotations in a three-dimensional space.

### Cross Sections

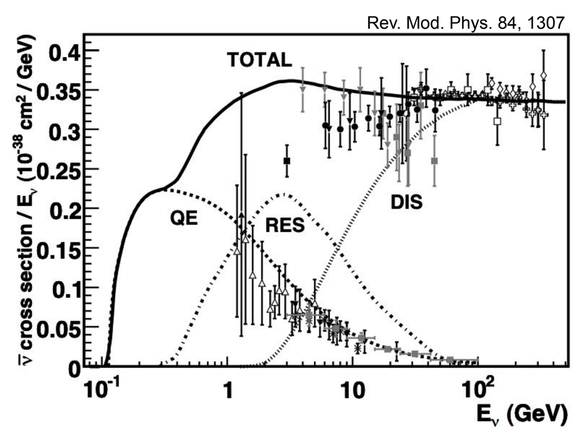


- These three active neutrino flavor states interact via the Weak force at a very low rate
  - Over 100 billion neutrinos from the sun are passing through the tip of your thumb every second
  - Requiring large-volume detectors and intense neutrino beams to induce interactions
- As neutrino beams become more powerful, nuclear effects play an increasingly important role.
- Understanding these effects over a wide range of neutrino energies is an active and growing area of neutrino research.

## The NOvA Experiment

### Physics Goals

- **Nu**MI **O**ff-axis $\nu_e$ **A**ppearance
- Measures oscillation of neutrinos and antineutrinos
- Designed to maximize sensitivity to $\nu_e$ appearance
- Search for beyond standard model such as sterile neutrinos and Non-Standard Interactions

### Massive Imaging Detectors

| Far Detector (FD) | Near Detector (ND) |
|---|---|
| Ash River, MN | Fermilab, IL |
| 15m x 15m x 60m | 4m x 4m x 15m |
| 14 kT | 200 ton |

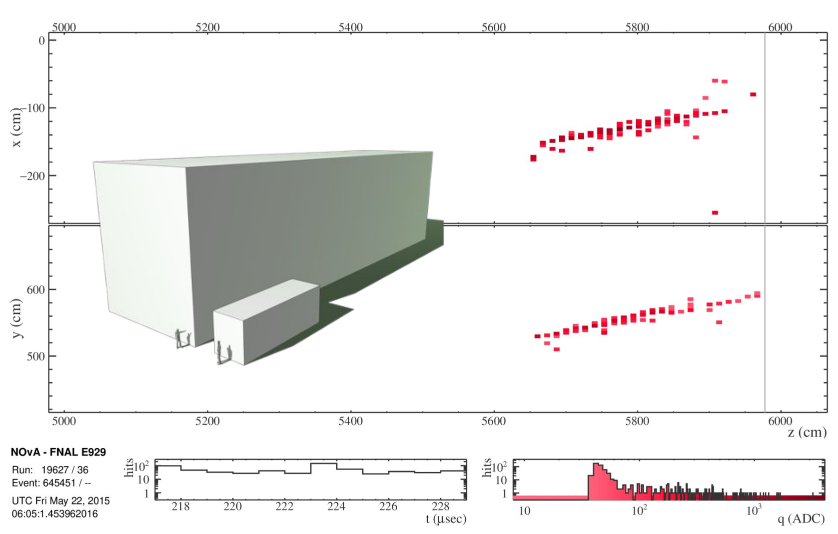

## NOvA Computing

### Traditional

**Computing infrastructure**

- Fermilab grid compute system for bulk analysis work
- Few-core login nodes for development and interactive analysis work

**Data ingestion**

- ~40 PB of data on tape storage
- Terabytes cached for FTP to grid nodes
- 100 MB average file size

**Analysis infrastructure**

- C++ (ROOT)
- Deeply-nested ntuple event loop
- Histogramming, optimization, statistical tools

### HPC

**Advantages**

- Enable increasingly complicated physics analyses with large pool of high speed computing resources – **NOvA Feldman-Cousins**
- Leverage massively parallel systems for improved computational efficiency and perform large-scale interactive analysis work – **NOvA Notebook Analysis Workflow**

**Getting there**

- Leverage distributed file systems and parallel IO with HDF5
- Support current analysis software with containerization
- Develop python analysis infrastructure for HDF5 and Jupyter

## NOvA Feldman-Cousins Procedure

**Summary of Feldman-Cousins Procedure**

- The best fit oscillation parameters ($\boldsymbol{\theta}_{best}$) are inferred by comparing the numbers of neutrino events in observed data and the Monte Carlo expectation through minimizing a $\chi 2$ function with respect to large number of parameters.
- Since neutrino measurements violate conditions for the Wilks theorem (low statistics, bounded parameters), which enable analytic approximation, confidence intervals need be estimated through a brute-force method.
- Thousands of statistically fluctuated pseudo-experiments are generated from Monte Carlo predictions for each point($\boldsymbol{\theta}_i$) and fitted to build empirical test statistics distribution $\lambda_i = \chi 2(\boldsymbol{\theta}_i) - \chi 2(\boldsymbol{\theta}_{best})$.
- Feldman-Cousins (FC) corrected confidence intervals are obtained by comparing the observed $\Delta\chi 2$ at each point and the empirical test statistic distributions.

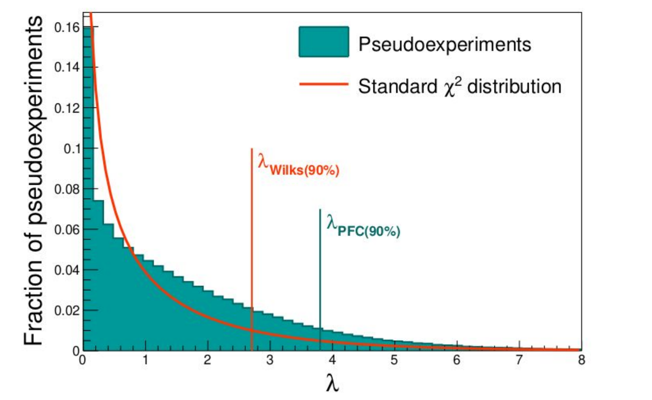


Fig: This plot shows an example of pseudo-experiment test statistics distribution for one degrees of freedom. Critical test statistics value to obtain 90% C.L. are marked as vertical lines with the PFC method and Wilk's approximation.

**Computational Challenge**

- Fit-time/pseudo experiment ~ O(1 hour)
- Total pseudo-experiments/plot ~ O(10 million)
- Total computing cost ~ O(20 million) CPU hours
- Computations need to be completed just in a few weeks before major conferences.

## Feldman-Cousins Implementation and Performance on NERSC Supercomputers

Fig: (right) 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. (left) NERSC Perlmutter system

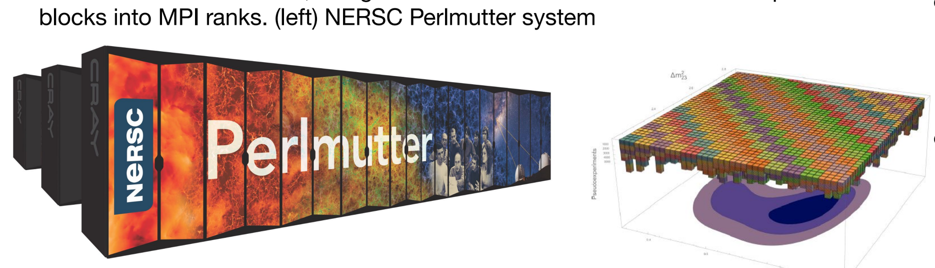


**NERSC Supercomputing Facility**

- Premier HPC center for science research, serving 7000 scientists, currently operating Cori and Perlmutter supercomputers

**Perlmutter**

- Two phase deployment since 2021
- It's indented to power AI applications through GPUs
- Runs on AMD Milan EPYC CPUs, which provide over 10x performance improvement over the Cori-KNL CPUs

- NOvA analysis framework CAFAna and the input files are stored in SL7 container to maintain software compatibility and fast parallel file access.
- The parameter space is decomposed through user-friendly DIY libraries over specified MPI ranks.
- Pseudo-experiment fit-time distribution is obtained through many small-scale jobs, which is used to determined large-scale job configurations.
- The 2022 NOvA sterile neutrino FC-correction analysis was performed on Cori-KNL nodes and for the first time on Perlmutter CPU nodes.
- Perlmutter-CPU nodes provided about 13x throughput over the Cori-KNL nodes to complete same amount of computations..

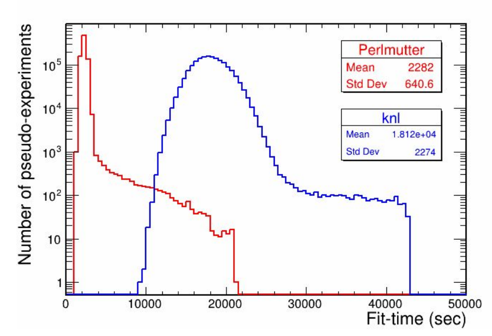


Fig : Pseudo-experiment fit-time distributions for Cori-KNL and Perlmutter for the plot $\sin^2\theta_{34}$-$\Delta m^2_{41}$

| Machine | Nodes /job | Ranks/ node | Job Duration | Total CPU Hours Utilized |
|---|---|---|---|---|
| Cori-KNL | 1682 | 68 | 12:00 hours | 20 M |
| Perlmutter-CPU | 200 | 128 | 6:00 hours | 1.4 M |

Table: Cori and Perlmutter systems computing resources used. Both parameter spaces considered in this analysis were equally computationally intensive in benchmark tests on the same machine.

- Low RAM footprint in the analysis code allowed usage of all available physical CPU cores on a node, resulting in maximum computing performance.
- It would be greatly beneficial to implement a dynamic load balancing mechanism that would assign new pseudo-experiment fits to the DIY blocks depending on which blocks completed their work before others and the remaining job time.
- Considerable effort was required to build new system libraries in SL7 container to be compatible with the new MPI stack on Perlmutter. In future, it would be desirable to find a better solution to maintain such software compatibility, as it's difficult for experiments to upgrade containers.

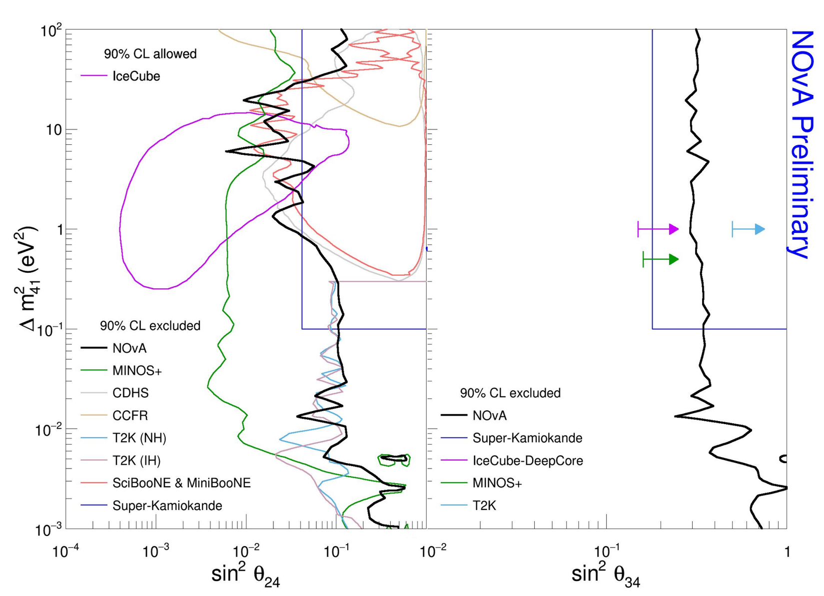


Fig: FC-corrected limits (black contour) for NOvA and comparison with other experiments for two parameter spaces presented at the Neutrino-2022 conference. The space to the right of the contour is excluded.

## Future Prospects for the FC Implementation

- Dynamic load balancing will enable more efficient computing resource utilization and flexible job configurations.
- Parameter sampling algorithm using Gaussian Process Regression will improve efficiency of confidence interval construction
- Accelerate objective function evaluation with GPUs

## NOvA Notebook Analysis Workflow

HPC platforms support large-scale on-demand and interactive data analysis. We have developed a Python-based columnar data analysis ecosystem with *implicit parallelism*

### HDF5 Data Organization

- HDF5 chunked-based I/O enables analysis of large files
- Concatenated almost 3TB of uncompressed data on NERSC Cori systems using only 985 CPU hours.
- Achieved 11x data compression
- Demonstrated significant speedup (10-20x) even with few ranks
- Enables analyses on local machines and improves flexibility of parallel event selection

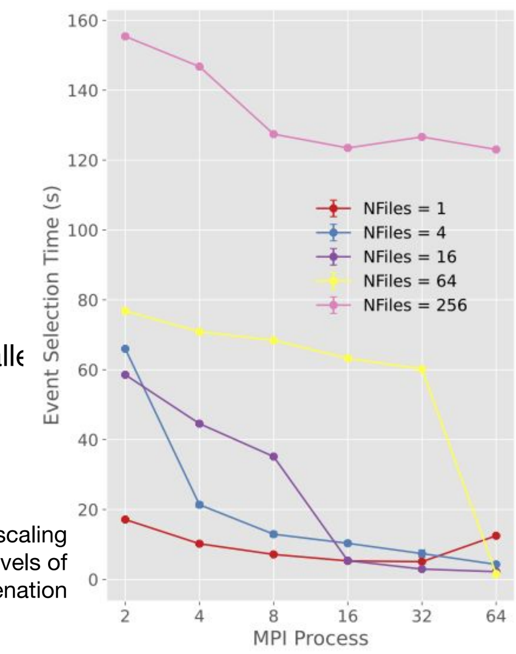


Fig: Event selection time scaling for increasing levels of concatenation

### PandAna

- Implicitly-parallel python columnar data analysis framework
- Interprets data as pandas Dataframe
- Parallel I/O with MPI

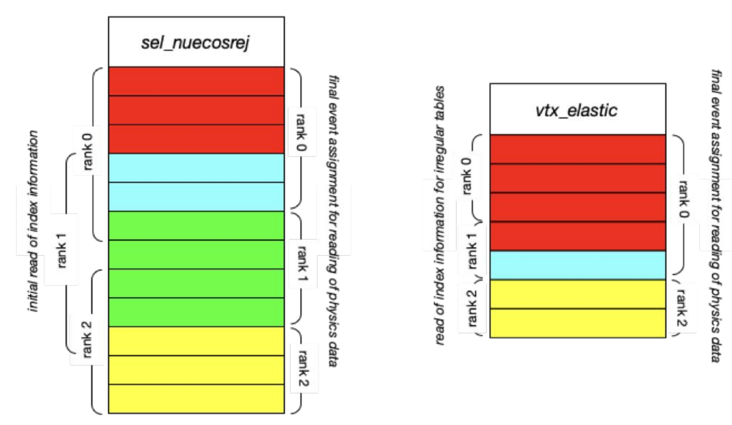


Fig: Schematic of PandAna data-parallel workload distribution using event indexing information. Indices enable cross-table association and supports irregular table sizes

```python
@pandana.Cut
def kFiducialCut(tables: pandana.Table) -> pd.DataFrame:
        df = tables['rec.vtx.elastic'] # returns a DataFrame
        df = (df['vtx.x'] > -100) & \
                 (df['vtx.x'] <  160) & \
                 (df['vtx.y'] > -160) & \
                 (df['vtx.y'] <  100) & \
                 (df['vtx.z'] >  150) & \
                 (df['vtx.z'] <  900)
        return df.groupby(level=['run', 'subrun', 'event']).first()

kNeutrinoEnergy = pandana.Var(lambda tables: tables['rec.mc.nu']['E'])
```

Fig: Example user code implementing a PandAna Cut (filter) and Var (transformation) demonstrating implicit parallelism

### Llama

- Histogram package for data-parallel analysis
- Wraps boost_histogram for performance
- Implicitly parallel data aggregation
- Implements Unified Histogram Interface for easy plotting

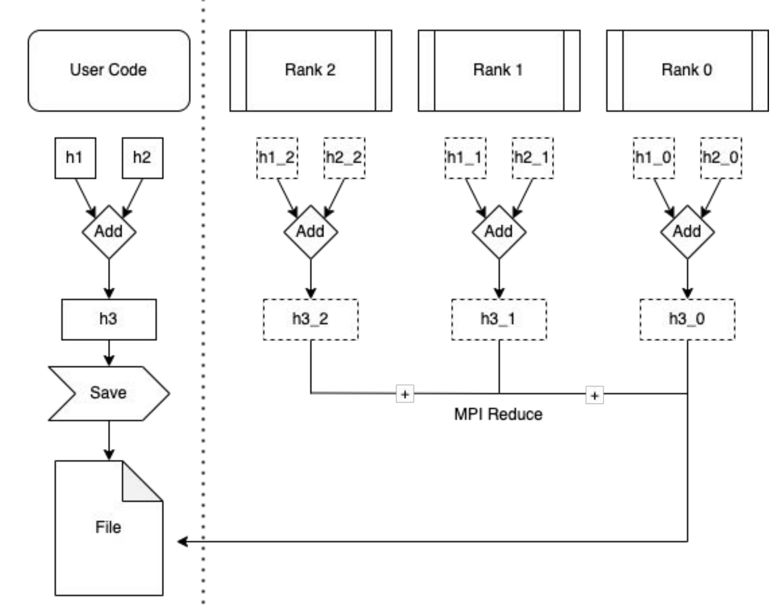


Fig: Schematic of Llama's lazy aggregation of histogrammed data based on arithmetic properties of histogram operations

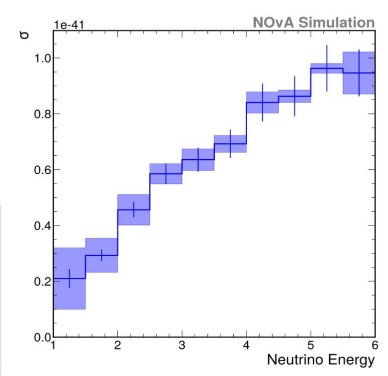


### Neutrino Cross Section Measurement

*We demonstrate near-perfect scalability of our implicitly-parallel workflow on a realistic sample size*

- NERSC's Perlmutter system
- Processed 600 GB of data
- Number of unique H5groups accessed = 16
- Number of unique H5datasets read = 54

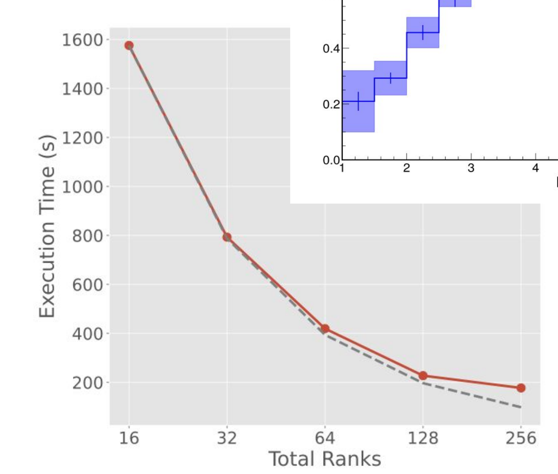


Fig. Near perfect scaling of our implicitly parallel analysis framework is demonstrated with an analysis of realistic scale

## Conclusions

- HPC platforms will enable more robust physics results on reasonable timescales
- Two physics applications successfully implemented on the Perlmutter
- World-leading limits on sterile neutrino oscillation parameters enabled by NERSC systems
- Demonstrated near perfect scaling of python-based implicitly-parallel analysis workflow
- Potential for further acceleration with Perlmutter GPU system

### References


[1] *Phys. Rev. Lett.* 123, 151803, 2019.
[2] *Nucl. Instrum. Meth.*, vol A389, pp. 81-86, 1997
[3] http://hdf.ncsa.uiuc.edu/HDF5
[4] arXiv:2203.13768, 2022.
[5] *EPJ Web Conf.* 251, 03033. 2021.
[6] *Phys. Rev. Lett.* 127, 201801, 2021.
[7] arXiv:2207.14353, 2022.
[8] *2011 IEEE Symposium on LDAV*: 105-112, 2011.
[9] J. Hartnell, "New Results from the NOvA Experiment". NEUTRINO 2022 Conference
[10] *Phys. Rev. D* 101, no.1, 012001, 2020.


### Acknowledgments


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. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists, Office of Science Graduate Student Research (SCGSR) program. The SCGSR program is administered by the Oak Ridge Institute for Science and Education for the DOE under contract number DE-SC0014664. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory.


http://novaexperiment.fnal.gov