

"Freight Train" production model on the NOvA experiment and NOvA efforts at Argonne Leadership Computing Facility

Andrew Dye

University of Mississippi

Presented at Fermilab New Perspectives 2023

June 26, 2023



THE UNIVERSITY of
MISSISSIPPI



Fermilab



The NOvA experiment - the beam

- ▶ Long baseline, high energy experiment to study neutrino oscillations
- ▶ NuMI Off-axis ν_e Appearance
 - ▶ NuMI, Neutrinos at the Main Injector
 - ▶ Nominally 700kW beam fired at graphite target
 - ▶ Pions from collision selected utilizing focusing horns
 - ▶ Pions decay into muon and muon neutrinos
 - ▶ Muon neutrinos continue from near to far detector, ideally oscillating into electron neutrino
 - ▶ Detectors placed 14mrad off center of beamline to maximize electron neutrino appearance




The NOvA experiment - the detectors

- ▶ Near detector
 - ▶ Located at Fermilab
 - ▶ 105m underground
 - ▶ 0.3 kiloton with 20k detection channels
- ▶ Far detector
 - ▶ Located 810km away in Ash River, Minnesota
 - ▶ On the surface
 - ▶ 14 kiloton with 344k detection channels



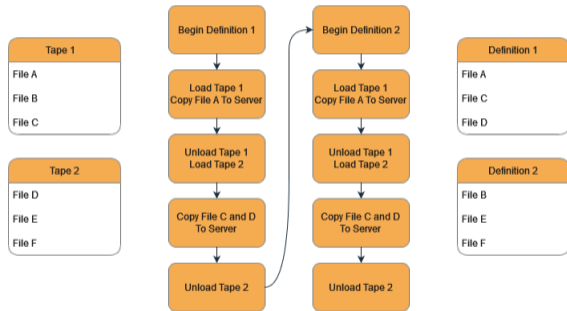
Data management for the NOvA experiment at Fermilab

- ▶ The NOvA experiment needs to store large amounts of data
 - ▶ On the order of 10's of petabytes
- ▶ This data is stored on physical tape at Fermilab
 - ▶ Must be fetched by robot and loaded
 - ▶ Raw data from many trigger streams stored on each tape
 - ▶ Fermilab wrote a short [article](#) about this process
- ▶ Getting data to the server, known as prestaging, takes time
 - ▶ Prestaging is a large bottleneck for production at NOvA



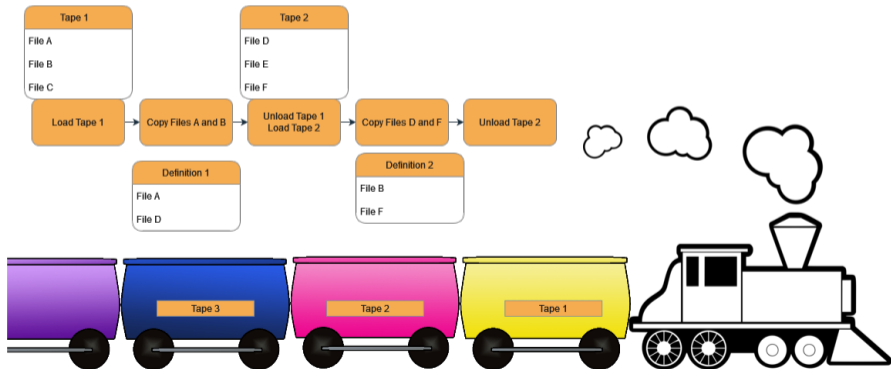
How NOvA used to request data

- ▶ Create groupings of files based on need
 - ▶ Called "Definitions"
- ▶ Process each Definition one by one
- ▶ Potentially load a tape multiple times
 - ▶ A major contributor to the bottleneck

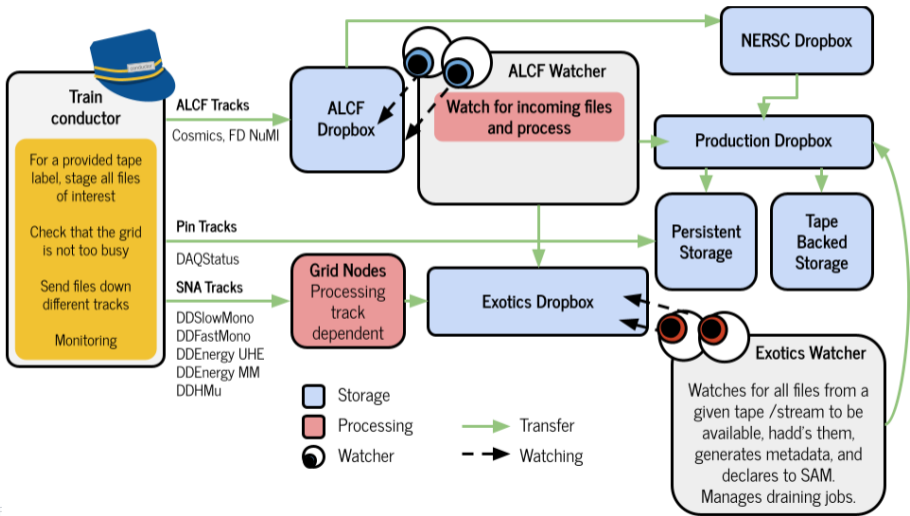


Freight Train Model – Simple

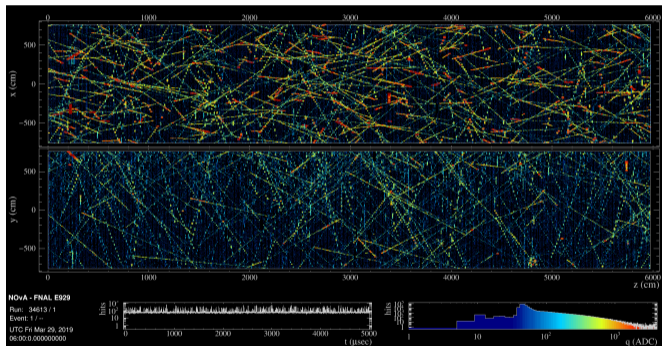
- ▶ Still create Definitions as normal
- ▶ Go through each Definition and arrange files by what tape it lives on
- ▶ Move sequentially through tapes, pulling listed files for every Definition



Freight Train Model – Full




- ▶ Far detector – on the surface – large amount of cosmic rays

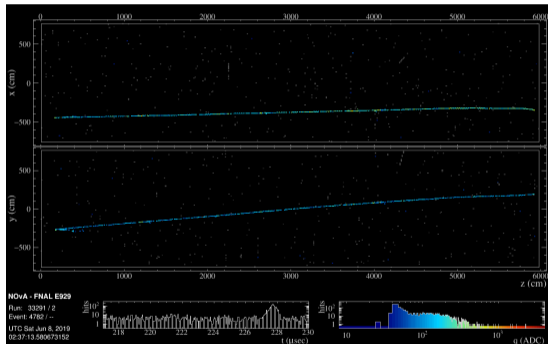


Cosmic Filtering

- ▶ Filtering necessary since there could be candidates for the oscillation analysis in the background
- ▶ Amount of recorded cosmic data coupled with probability of finding candidate seems overwhelming
 - ▶ To fully process all recorded cosmic data currently would take over 2 years of continuous processing
 - ▶ Only 1 in 10^7 cosmics will be candidates



- ▶ Cosmic filtering – computationally expensive – 10^6 rejection power




NOvA at Argonne Leadership Computing Facility (ALCF)

- ▶ For efficient filtering, we require a reliable access to large number of GPUs
 - ▶ Process involves running event display images through a neural network
 - ▶ Image based algorithms benefit greatly from GPU processing
- ▶ ALCF provides this resource via its ThetaGPU farm
 - ▶ A High-Performance Computing Facility with a few very powerful GPU nodes




State of Cosmic Filtering at ALCF


- ▶ Currently have 4 different streams of jobs running on ThetaGPU
 - ▶ Each stream able to run 128 jobs simultaneously
 - ▶ 512 jobs at once!



Summary and Moving Forward



- ▶ Freight Train production model set to increase overall throughput of data
 - ▶ Moving ahead, only need to make sure the trains don't derail!
- ▶ Cosmic filtering at ALCF on track and is a valuable asset to the NOvA Collaboration
 - ▶ Keep aware of any other jobs that could benefit from GPU processing



ENGINEERING TIP:
WHEN YOU DO A TASK BY HAND,
YOU CAN TECHNICALLY SAY YOU
TRAINED A NEURAL NET TO DO IT.





Thank you

Questions?

