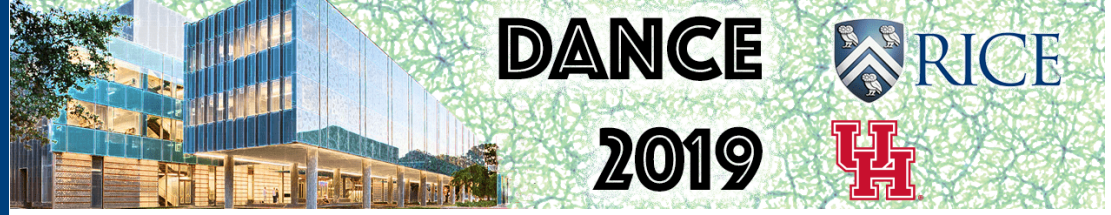

NOvA Software Infrastructure

Gavin S. Davies

THE UNIVERSITY of MISSISSIPPI

October 28, 2019

NOvA experiment

NuMI **O**ff-axis ν_e **A**ppearance experiment
NuMI = Neutrinos at the Main Injector

- Start with **world's most powerful neutrino beam**
- NuMI ν_μ beam at Fermilab

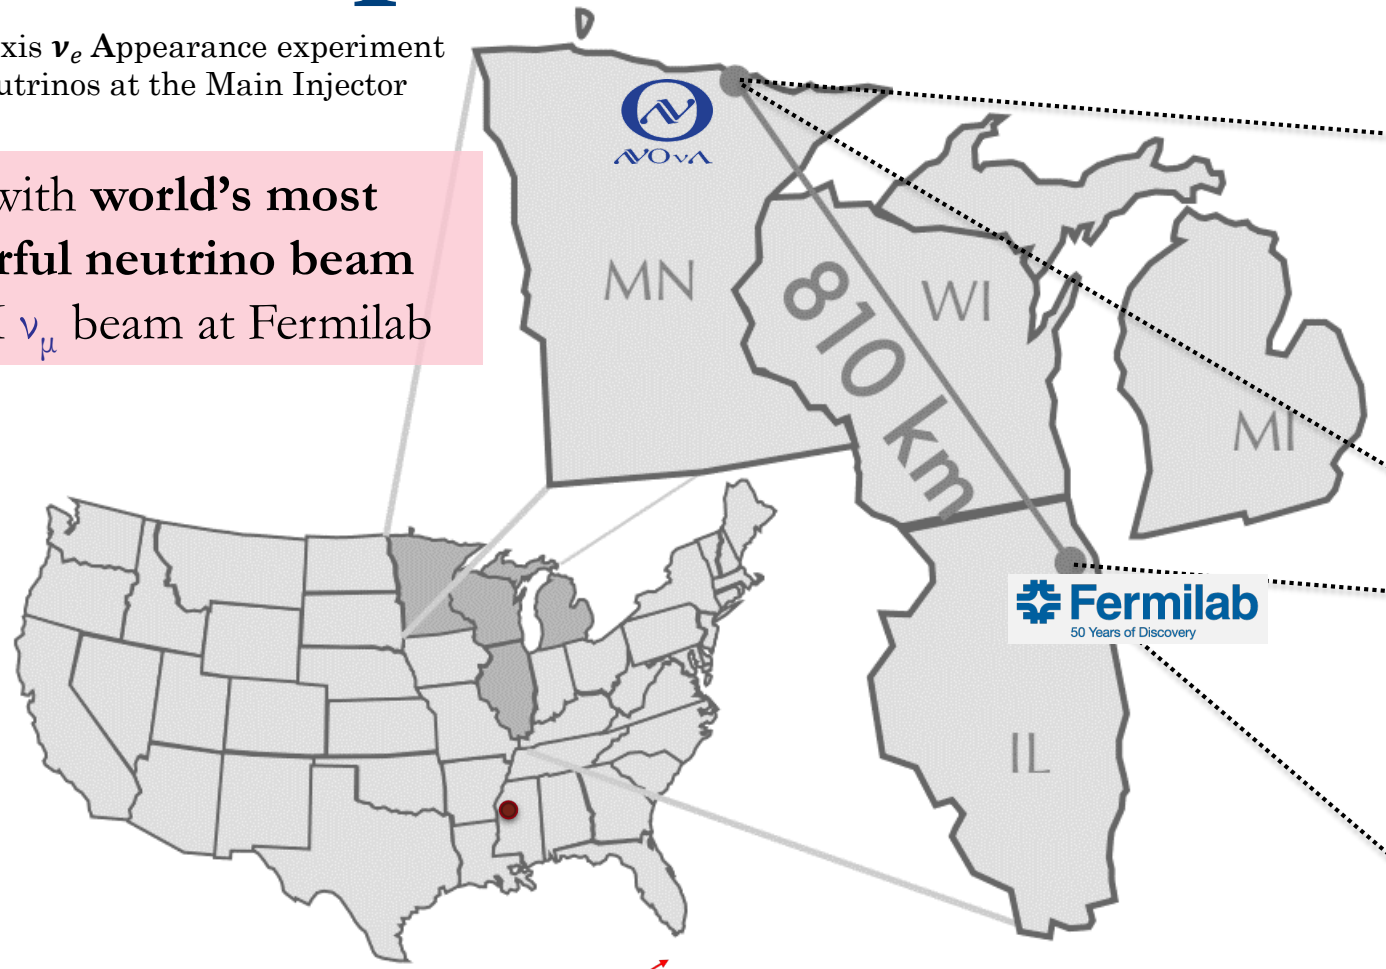

Far Detector

Near Detector

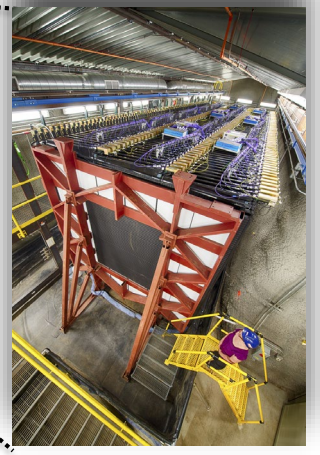

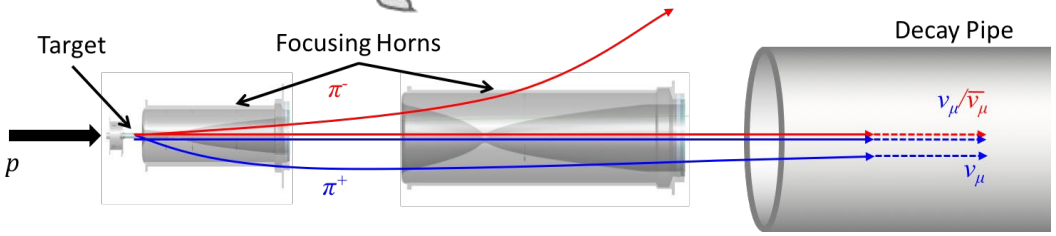

http://novaexperiment.fnal.gov

NOvA Physics

a long-baseline neutrino oscillation experiment

ν_μ CC Disappearance
- $\nu_\mu \rightarrow \nu_\mu$ & $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$
- Precision measurements of:
  $\sin^2(\theta_{23})$ & $|\Delta m^2_{32}|$

ν_e CC Appearance
- $\nu_\mu \rightarrow \nu_e$ & $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$
- Neutrino mass hierarchy
- CP violating phase
  θ_{13} & θ_{23} & δ_{CP}

First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA
Phys. Rev. Lett. 123, 151803 (Oct. 2019)

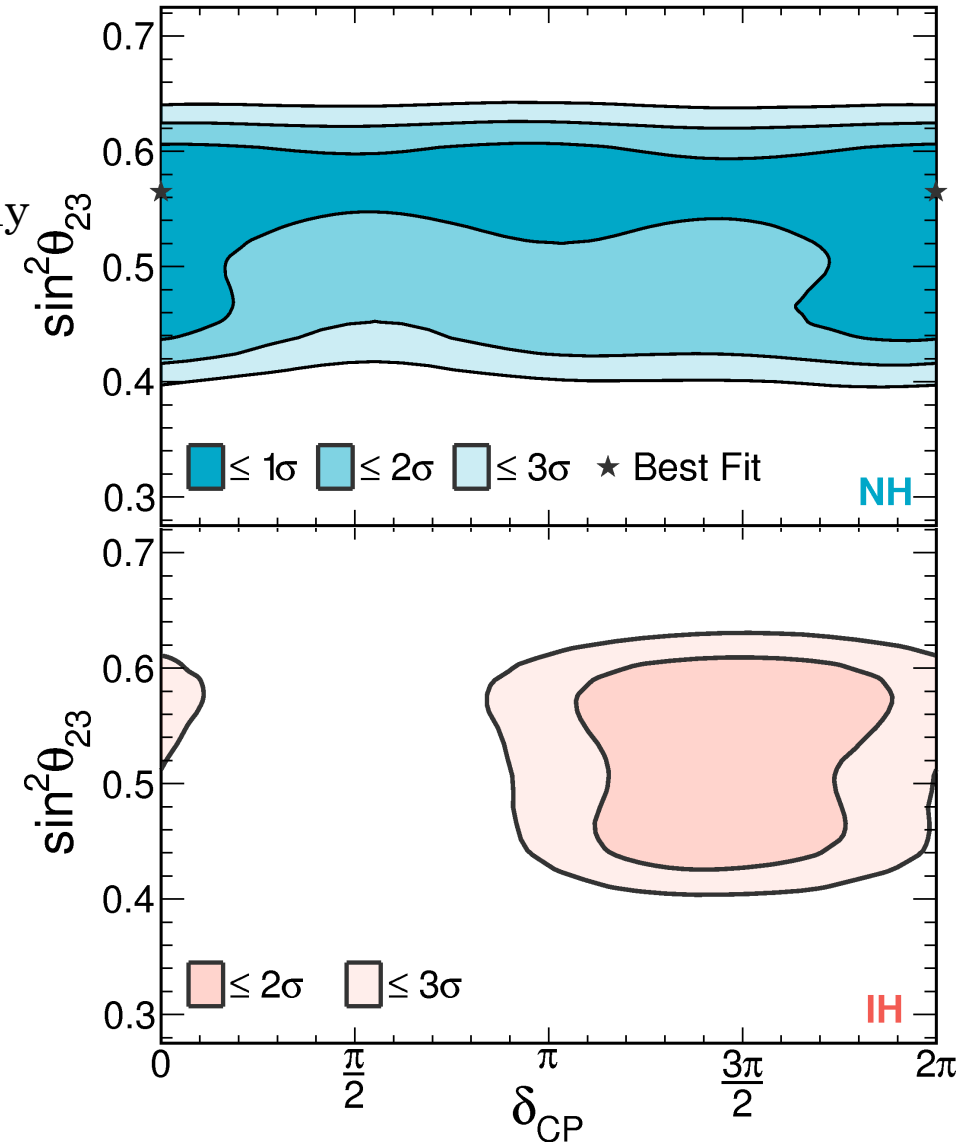

-Neutrino scattering/cross-sections
(new:arXiv:1902.00558)
-Sterile ν searches
-Seasonal cosmic ray studies
(new: arXiv:1904.12975)
-Supernova ν's
-Exotic searches:
magnetic monopoles, GW multi-messenger. n-n̄ oscillations, …

NOvA Computing I

- Virtual Machines
  - 15 VMs for interactive use on BlueArc (NFS-based Network Attached Storage [NAS])
  - Several hundred TB of BlueArc "short-term" storage
    - Originally everything via BlueArc – not scalable
- **CVMFS**
  - NOvASoft maintained/published on OSG server (/cvmfs/nova.opensciencegrid.org)
  - Scratch and Persistent areas
  - File streaming via **xrootd**

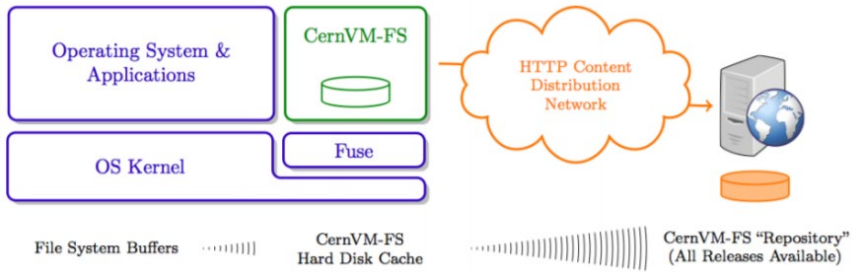

- Tape/Data Storage
  - Long-term storage
  - Cache management: Frontend ~4 PB **dCache** shared across IF experiments
  - **SAM** (Sequential data Access Metadata) data cataloguing system
  - **FTS** (File Transfer Service) daemon interfaces w/ SAM and automates file transfers
- Databases:
  - Several (PostgreSQL), required for online and offline operations – Accessed via http for offsite usage

NOvA Computing II

NOvASoft

- built using Fermilab's custom version of **Software Release Tools (SRT)**
- a UNIX based software management system
- Allows for asynchronous development
- Concept of a "test release" – single/multi package dev. versus base
  - Simple, fast build times
- SVN (offline), CVS (online) repositories – git for ML tools [see F. Psihas talk Tuesday]
- Alternatively, **Multi Repository Build (MRB) system** using **Cmake**

art

- event-processing framework [**J.Phys.Conf.Ser. 396 (2012) 022020**]
  - developed and maintained by Fermilab software engineers
  - Provides a robust, professionally maintained foundation

UPS/UPD

- access to external products is provided by a Fermilab-developed product-management package called Unix Product Support
- UPD (UNIX Product Distribution) provides the functionality for uploading/downloading products between local systems and product distribution servers
- Support waning – Fermilab developing **Spack**-based package management tool

NOvA Computing III

Jenkins

- open source automation server – supports building, deploying and automating projects
- Nightly builds
  - "Early warning" system for commits
- Continuous Integration

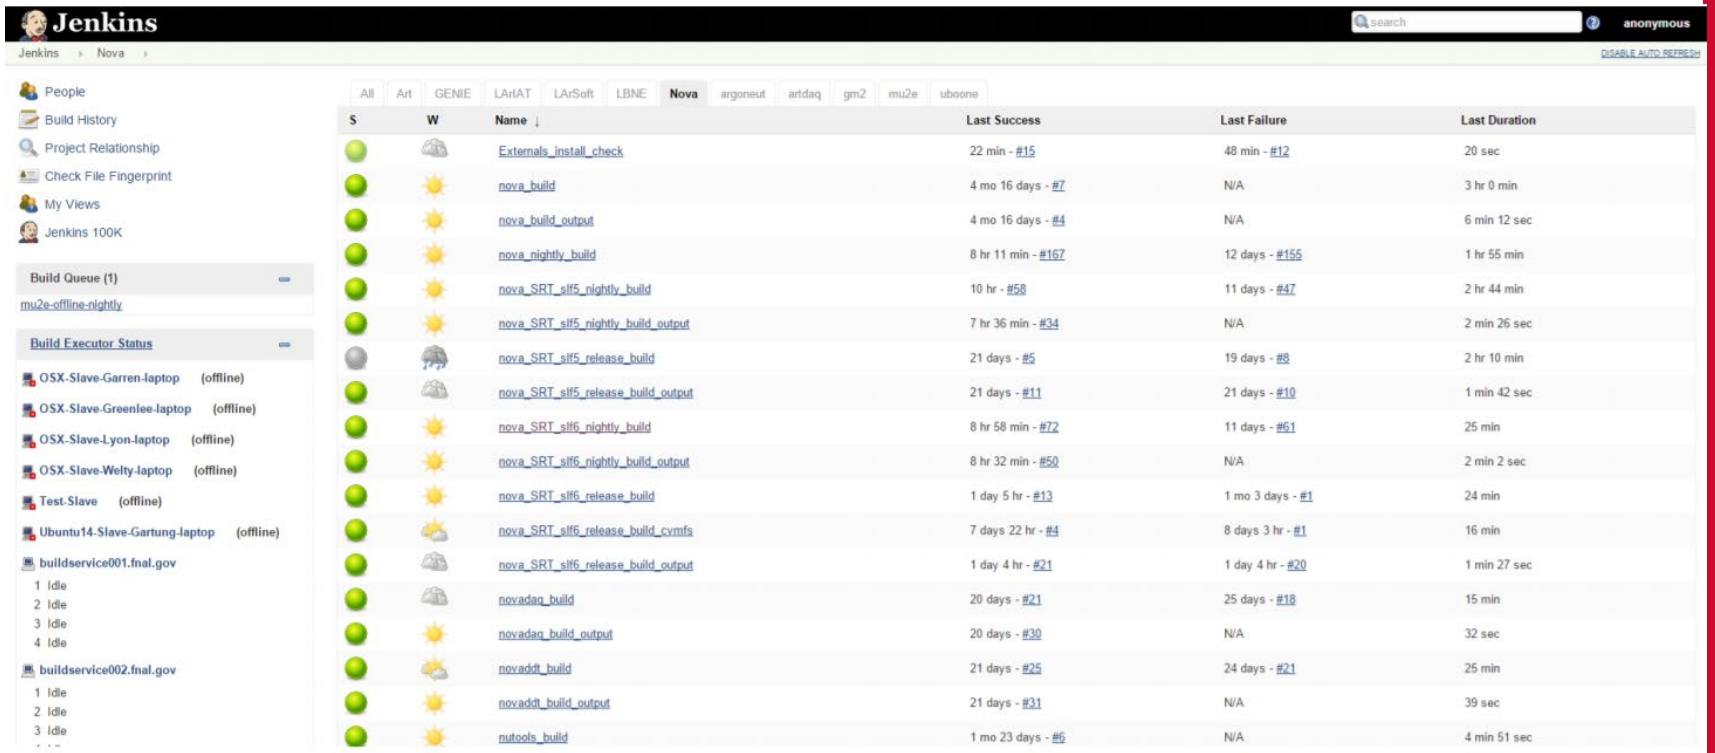

NOvA Simulation

Step	CHEP-2015	CHEP-2018	Products
Beam simulation	FLUKA / FLUGG	GEANT / External Data (MINERvA)	Simulates the incoming neutrino spices and spectrum from π, K, μ decays
Neutrino interactions	GENIE	GENIE + Cross-section tuning	Produces particle lists and kinematics to be propogated through the detector
Cosmic rays	CRY	Data triggers + Overlay	
Detector simulation	GEANT 4	GEANT 4 upgrade (neutron simulation) + revision in geometry	Propagates particles through the detector and produces energy deposits in active materials
Readout electronics and DAQ	Custom simulation routines	Updates of the light model and electronics setup	Converts energy deposits into scintillation light, its transport to APD and simulates the readout response, produces output like raw data

11

NOvA Computing IV

Batch

- Access via OSG to multiple clients and specific NOvA-dedicated offsite clients
  - includes Harvard University, Southern Methodist University, Institute of Physics of the Academy of Sciences (Prague, Czech Republic), and Ohio Supercomputing
- Local batch cluster: ~40 nodes
  - Grid slots at Fermilab for NOvA: 1300 node quota
    - opportunistic slots also available
  - Remote batch queues: thousands of additional slots

NOvA HPC

NERSC

- In collaboration with DOE SciDAC-4 and the HEPCloud program
  - Scientific Discovery through Advanced Computing
- Feldman-Cousins corrections in timely fashion
- Implemented with **Docker** image
- Estimated 50x speedup as compared to Fermigrid

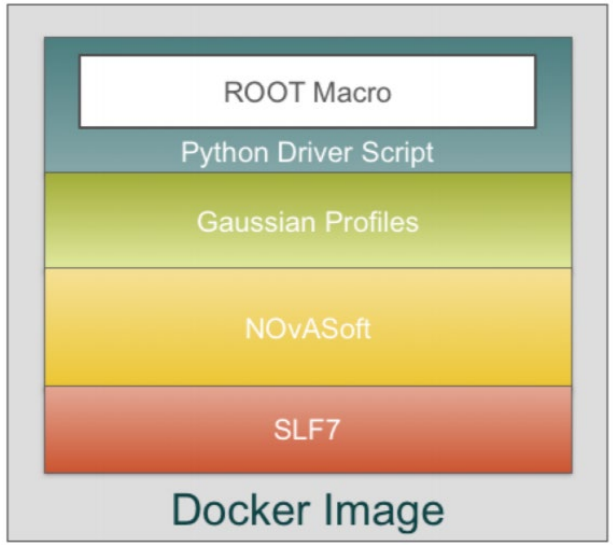

D. Doyle, CSU

Requires $\mathcal{O}(10^7)$ $\Delta\chi^2$ calculations!

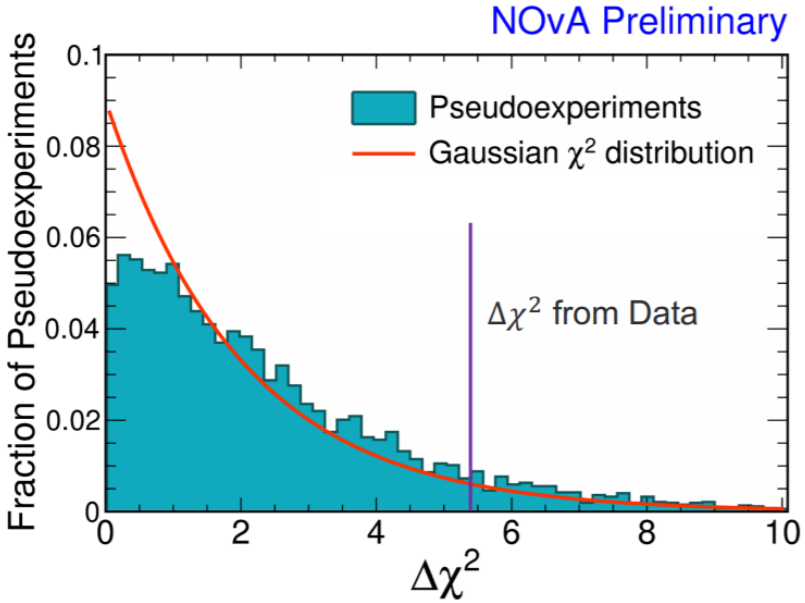

https://news.fnal.gov/2018/07/fermilab-computing-experts-bolster-nova-evidence-1-million-cores-consumed/

Needs/Challenges

- Scalability

- Knowledge transfer (and retention)
  - Includes training future considering proliferation of tools
  - Community sharing

Thank you for the invitation!

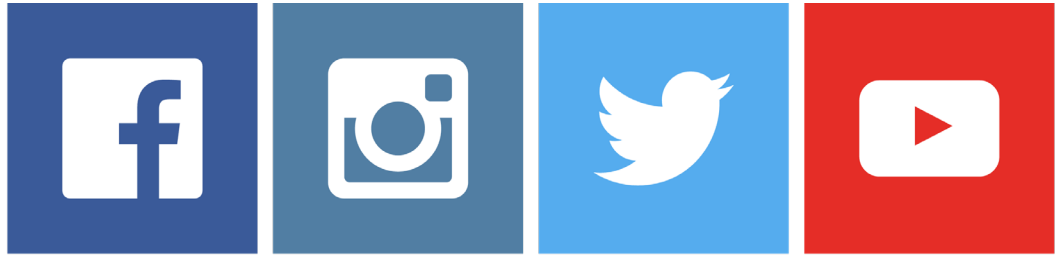

http://novaexperiment.fnal.gov