IceCube ML

Benedikt Riedel
(Presented by Josh Peterson)
UW-Madison

A3D3 MMA Meeting
10 July 2023

ICECUBE
SOUTH POLE NEUTRINO OBSERVATORY

IceCube

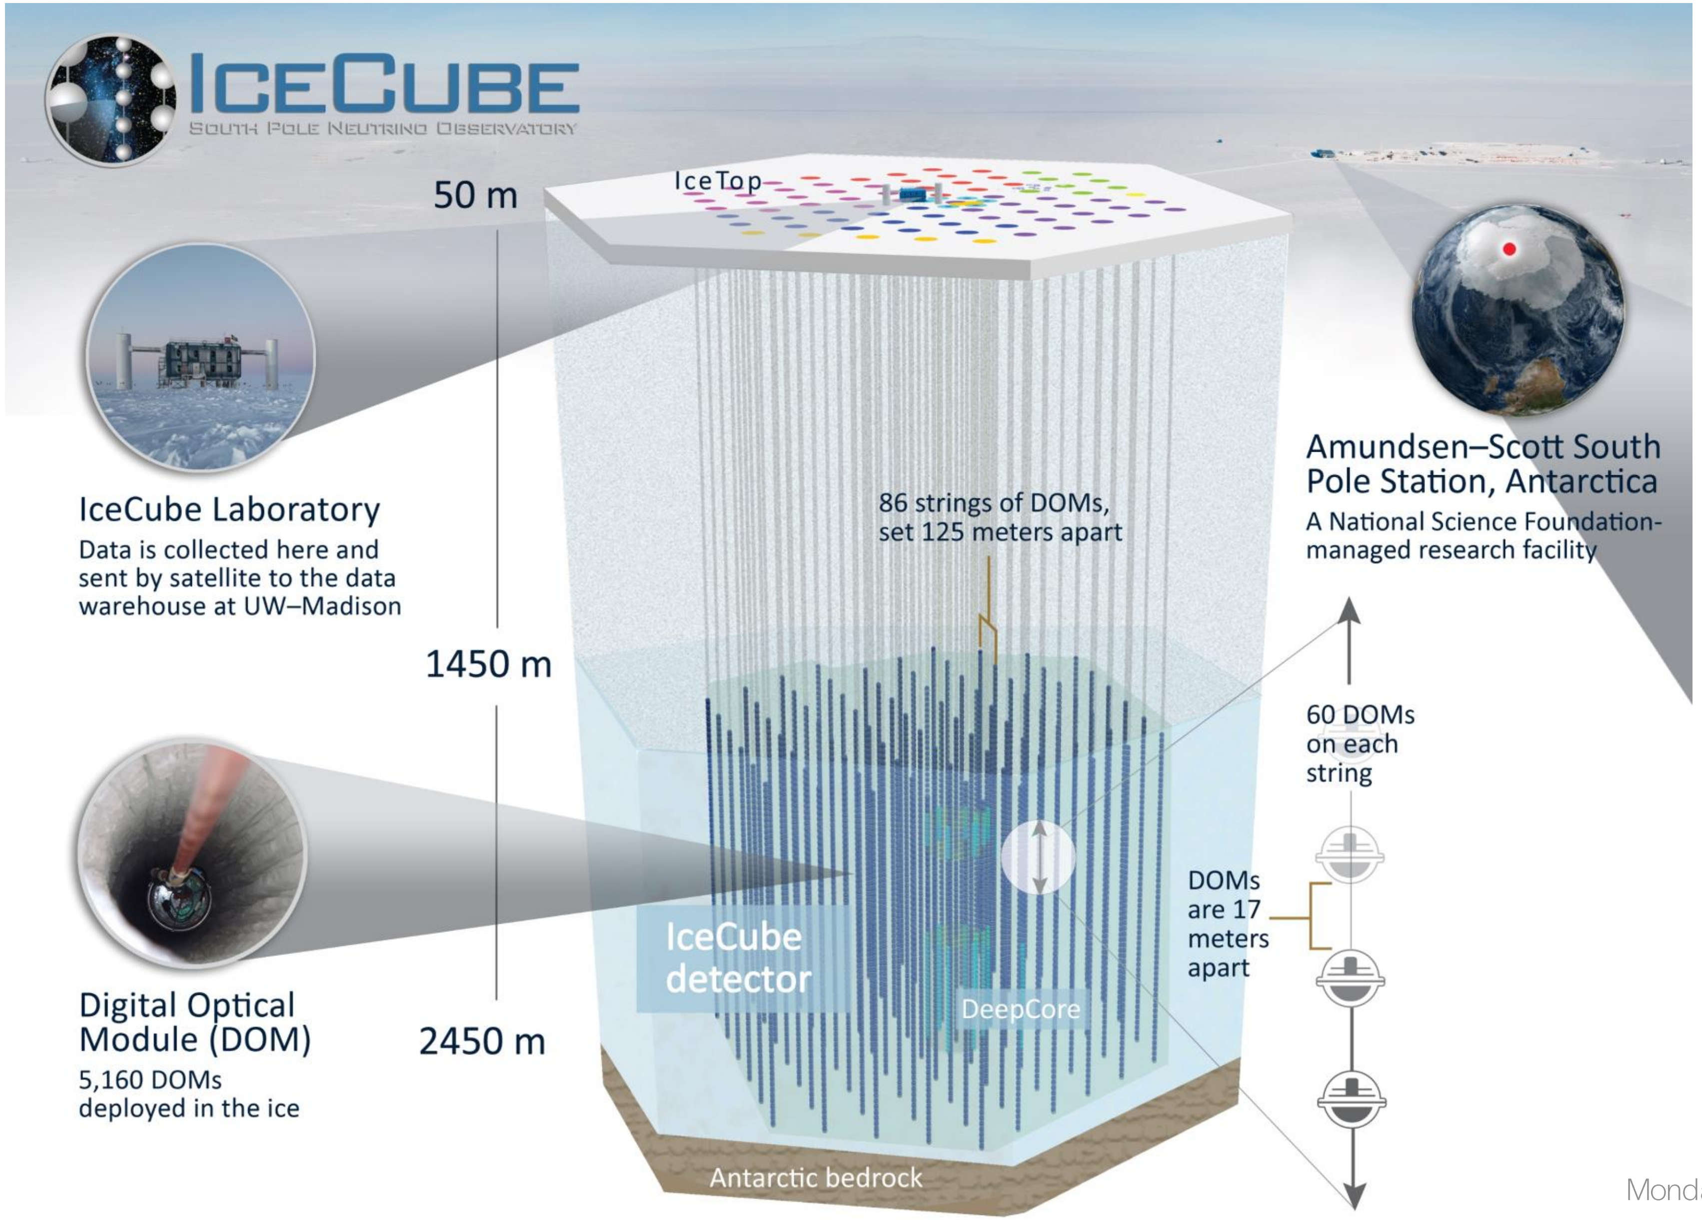

ICECUBE
SOUTH POLE NEUTRINO OBSERVATORY
50 m
IceTop
IceCube Laboratory
Data is collected here and sent by satellite to the data warehouse at UW–Madison
1450 m
86 strings of DOMs, set 125 meters apart
Amundsen–Scott South Pole Station, Antarctica
A National Science Foundation-managed research facility
60 DOMs on each string
DOMs are 17 meters apart
Digital Optical Module (DOM)
5,160 DOMs deployed in the ice
2450 m
IceCube detector
DeepCore
Antarctic bedrock

IceCube Science

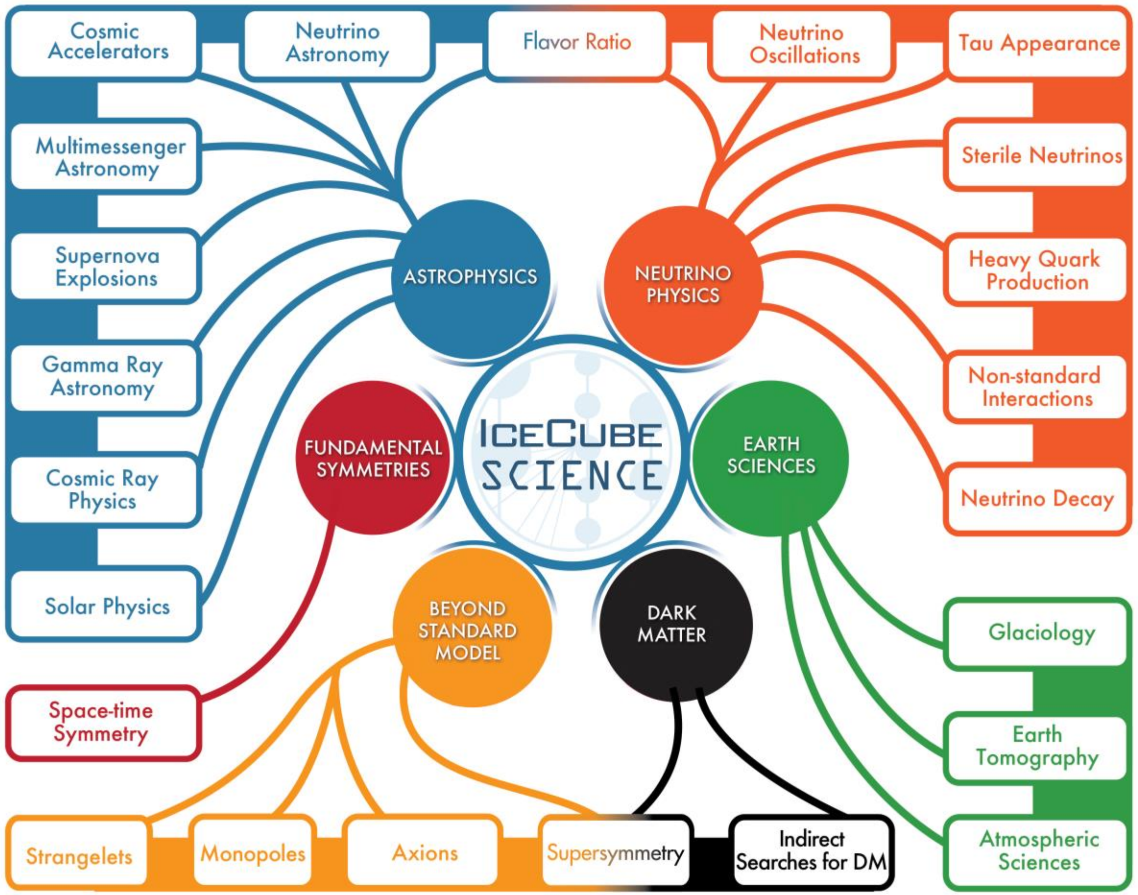

- Novel instrument in multiple fields
- Broad science abilities across 9+ orders of magnitude
  - MeV neutrinos: Core-Collapse Supernovae
  - GeV – TeV: Atmospheric neutrinos to study neutrino and BSM studies
  - > TeV: Neutrinos for VHE/UHE neutrino astronomy
- Lots of data that needs to be processed in different ways
- Lots of simulation that needs to be generated

IceCube Science – Why neutrinos?

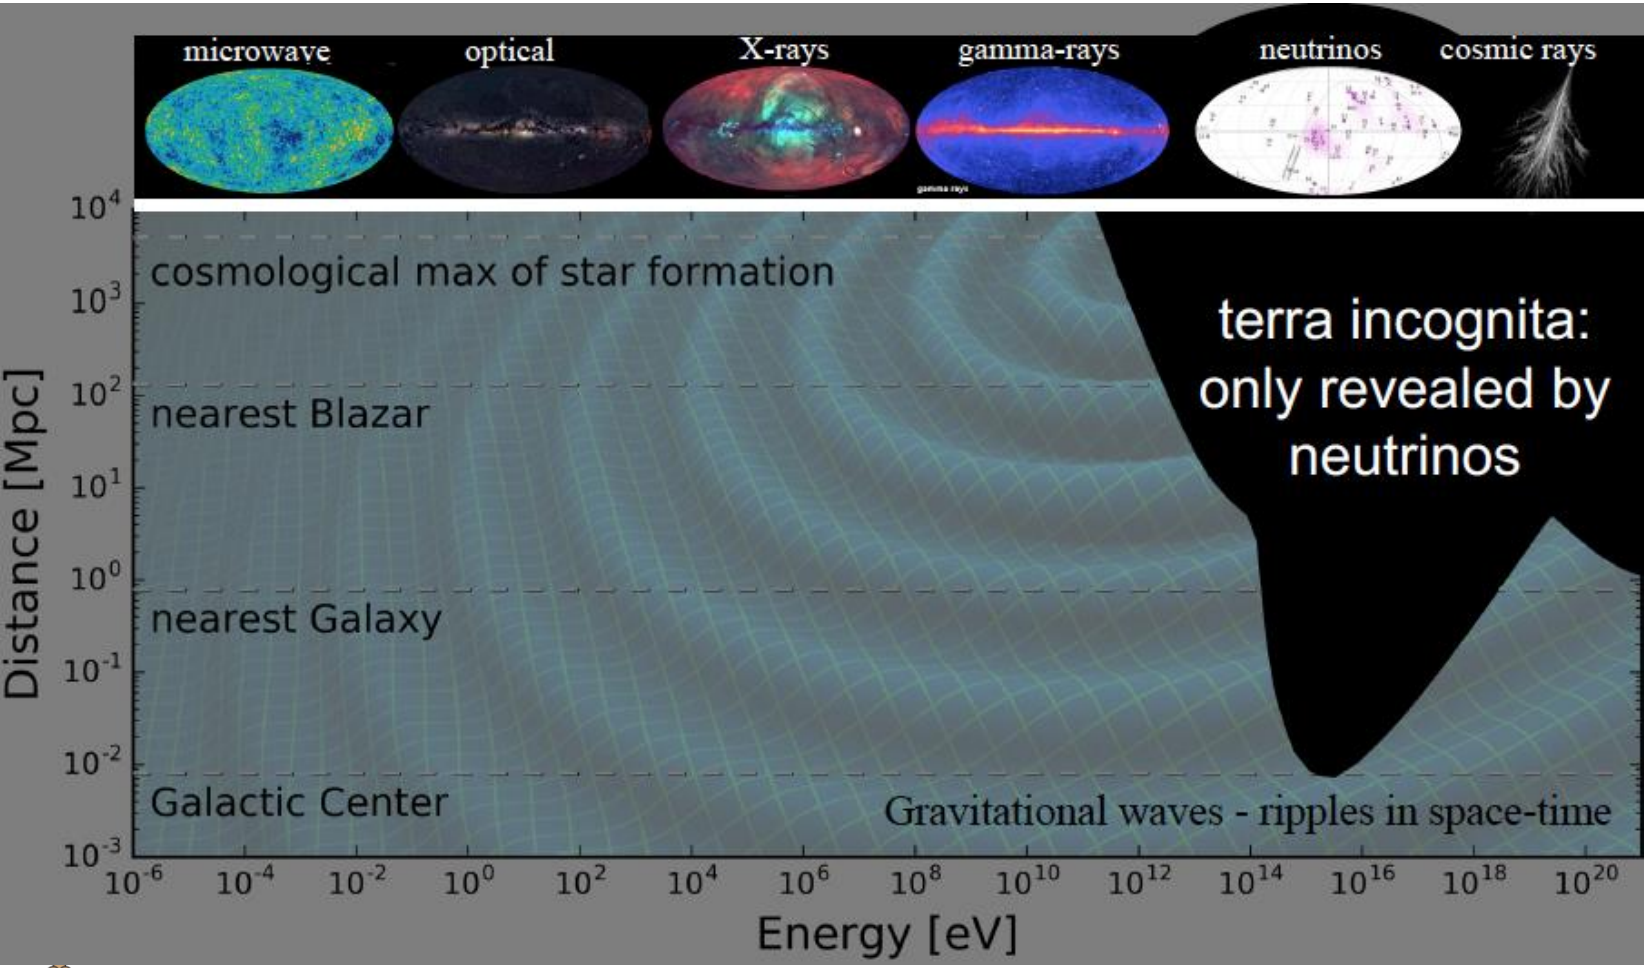

- 20% of universe is dark to “traditional” astronomy, i.e. using electromagnetic waves/light
- Need a new set of “messengers” – Gravitational Waves and Neutrinos

IceCube Science – How does it work?

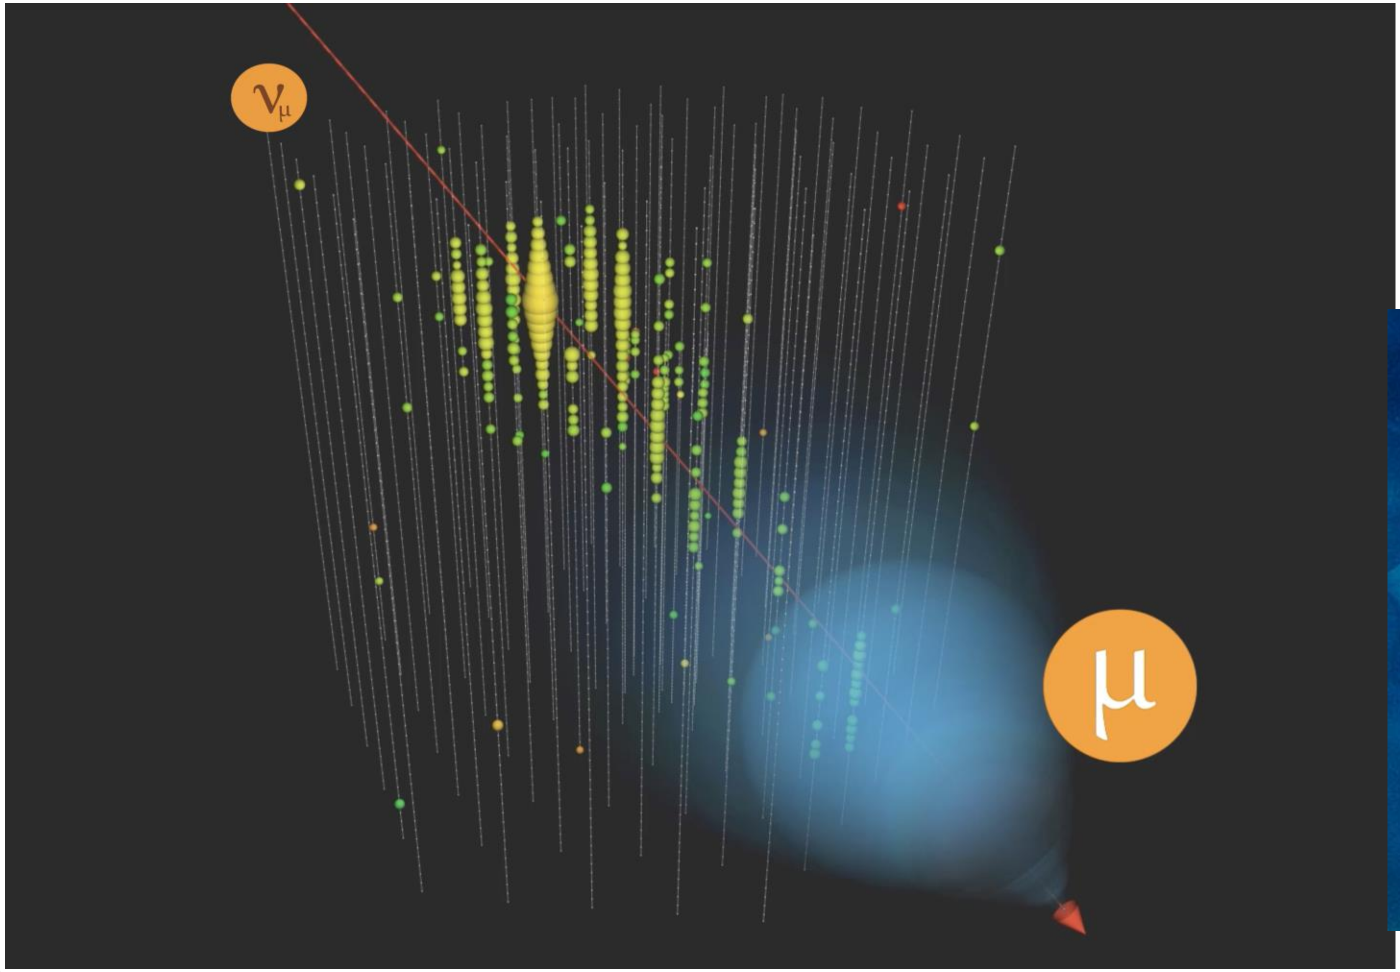

- Cherenkov light - Sonic boom with light
- Cherenkov light appears when a charged particle travels through matter faster than light can

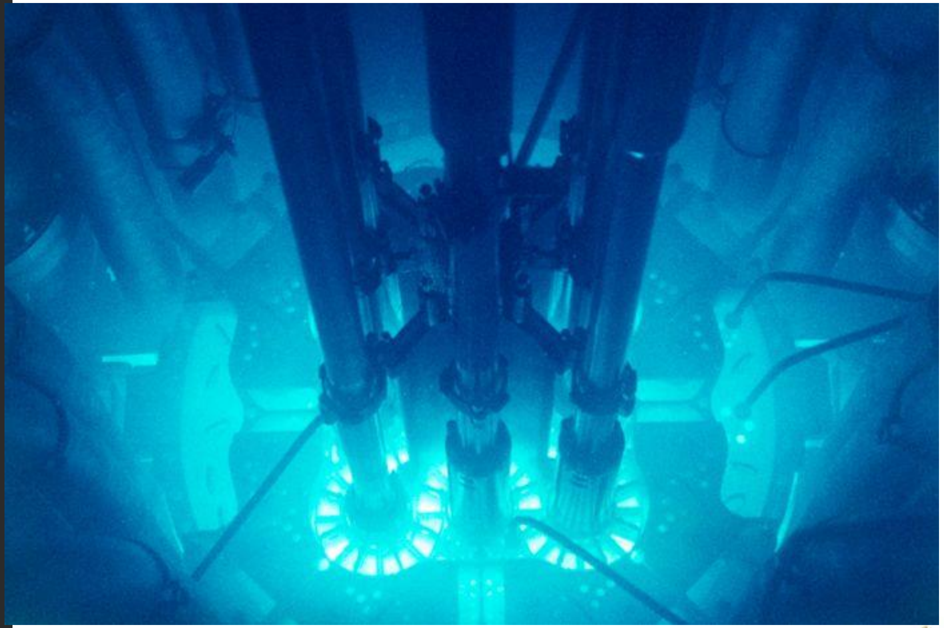

IceCube's Future – Upgrade

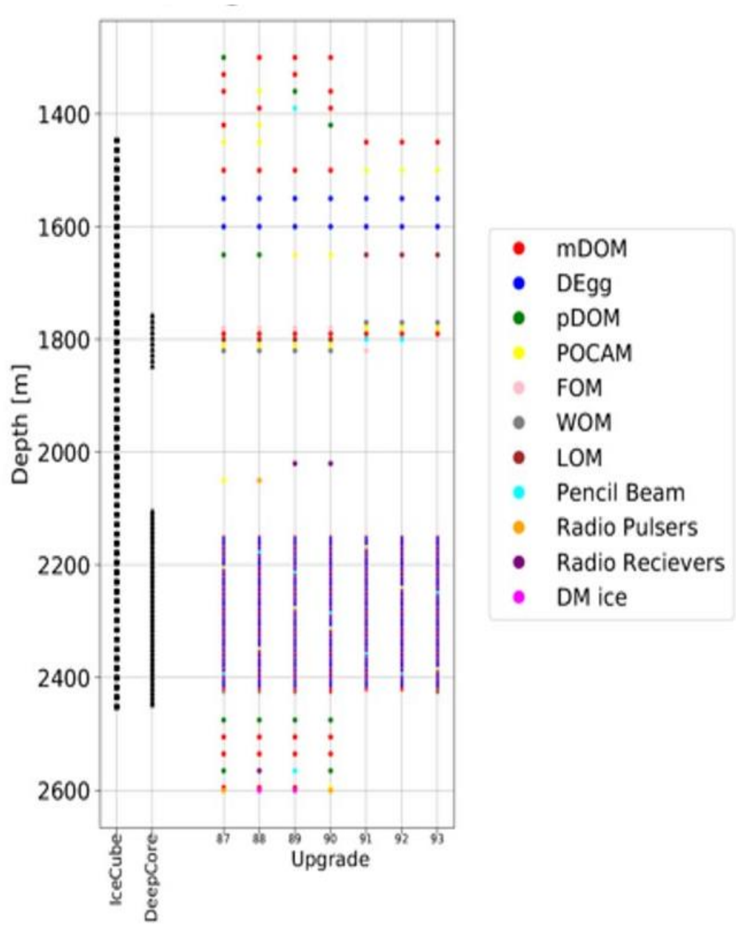

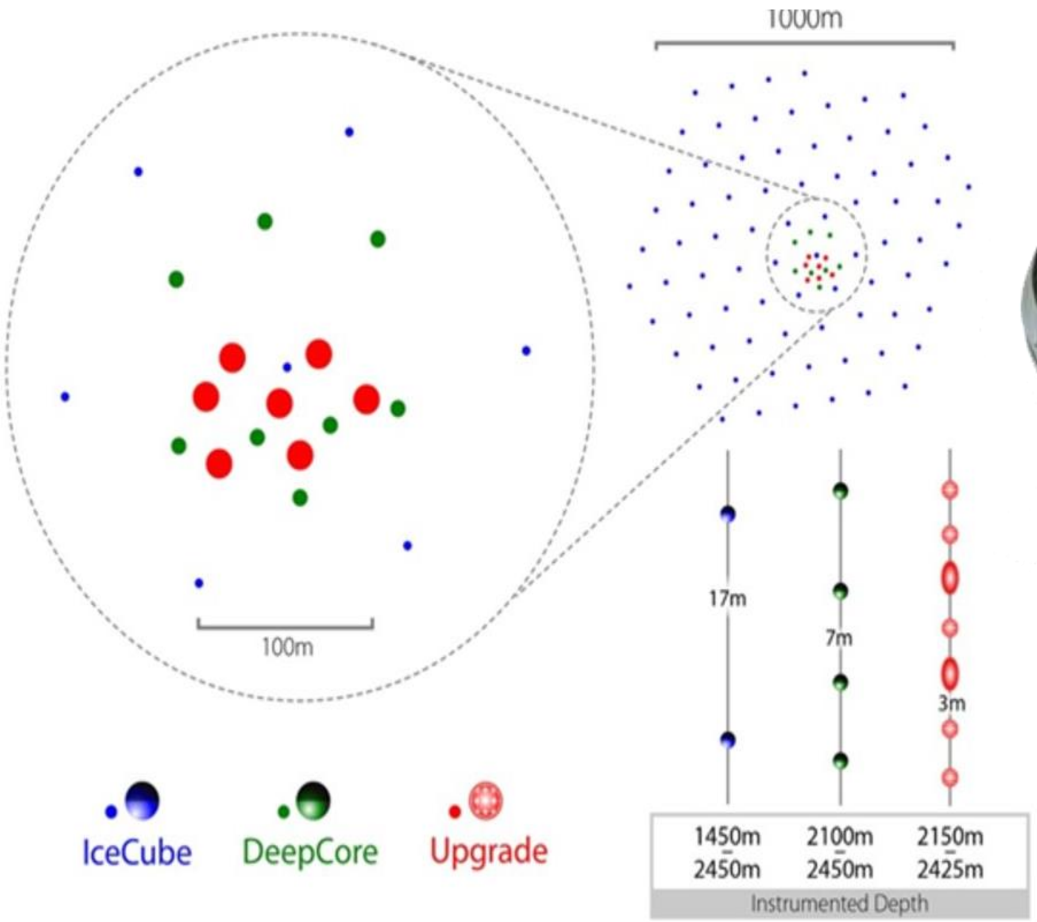

IceCube's Future – Gen2

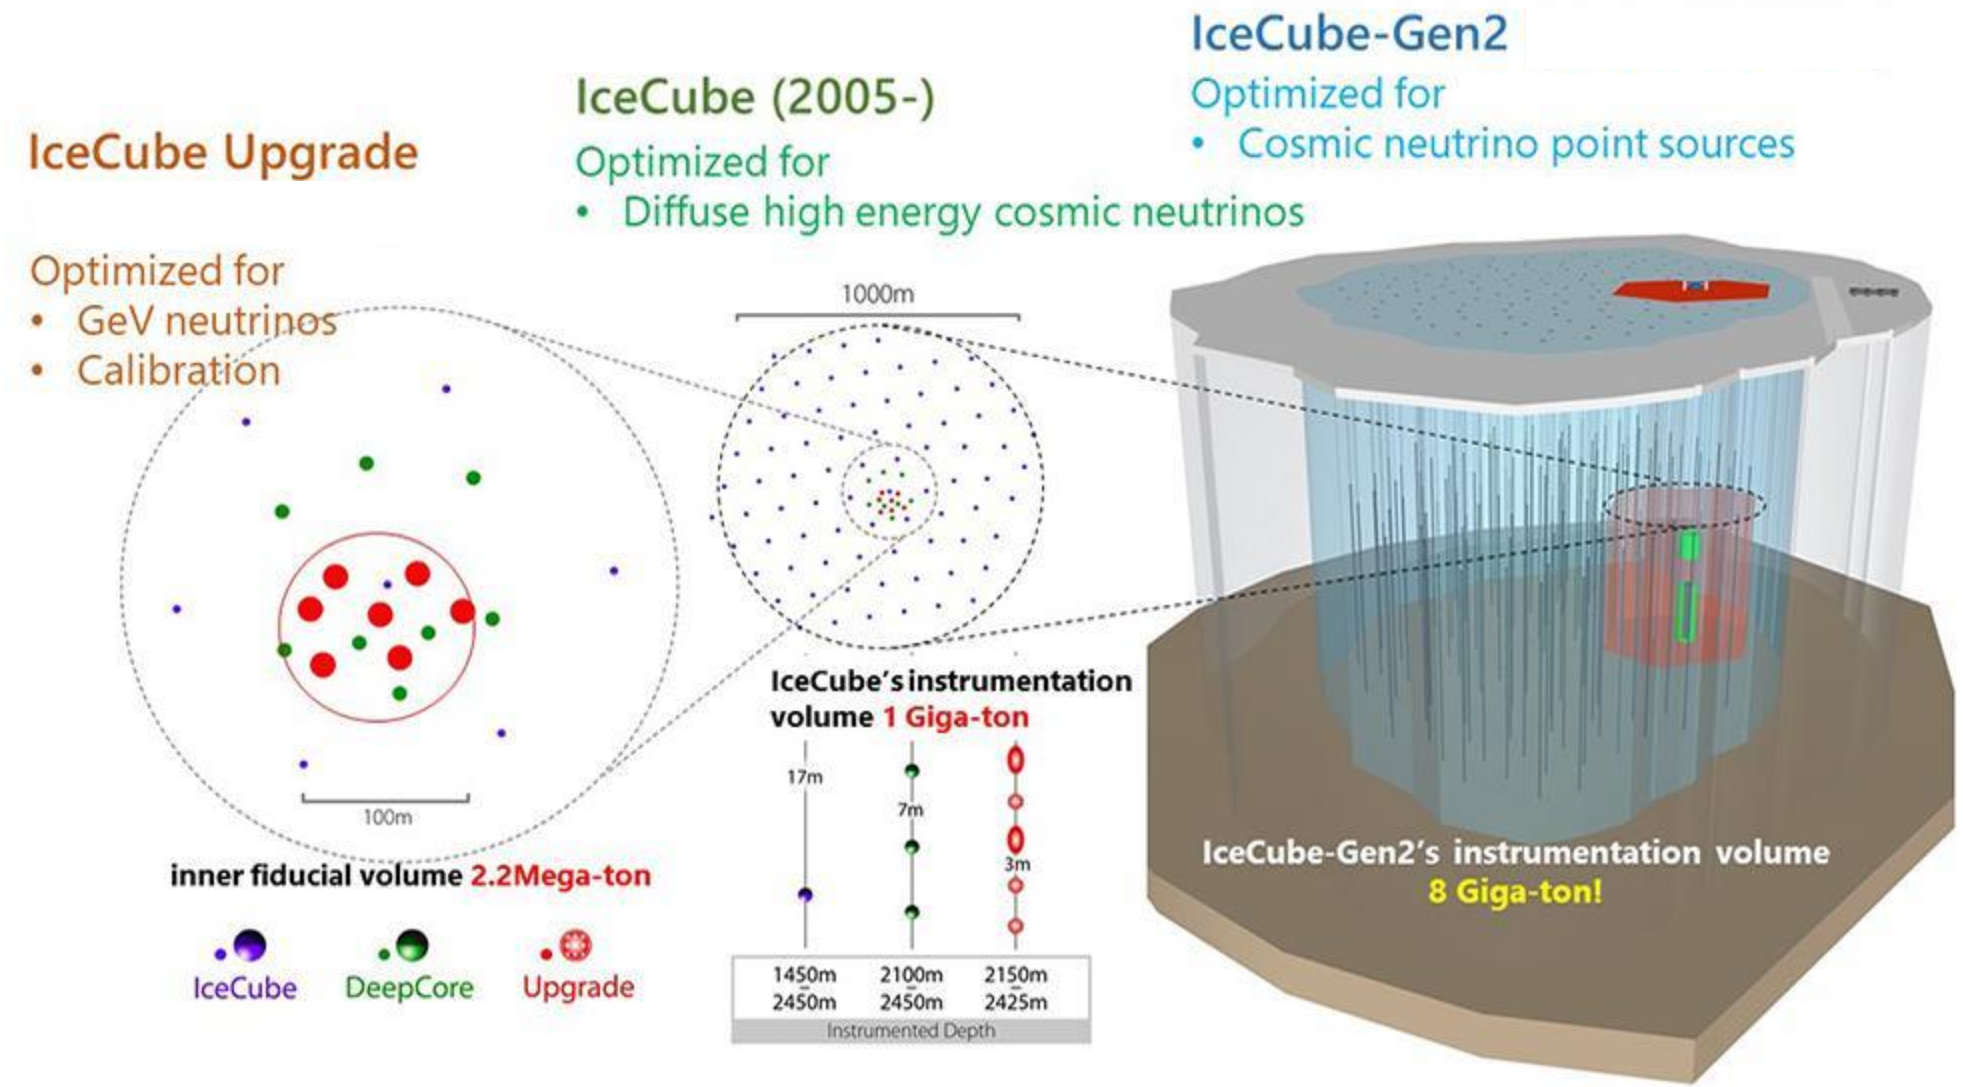

IceCube Events

- Data is a three-dimensional movie or time-evolving point cloud of photon hits
- Different neutrino flavors produce different event topologies

Track – Muon Neutrino

Good Pointing, Poor Energy

Cascade – Electron/Tau Neutrino, NC Interactions

Poor pointing, Good Energy*

Double Bang – High Energy Tau Neutrino*

Good Pointing, Good Energy

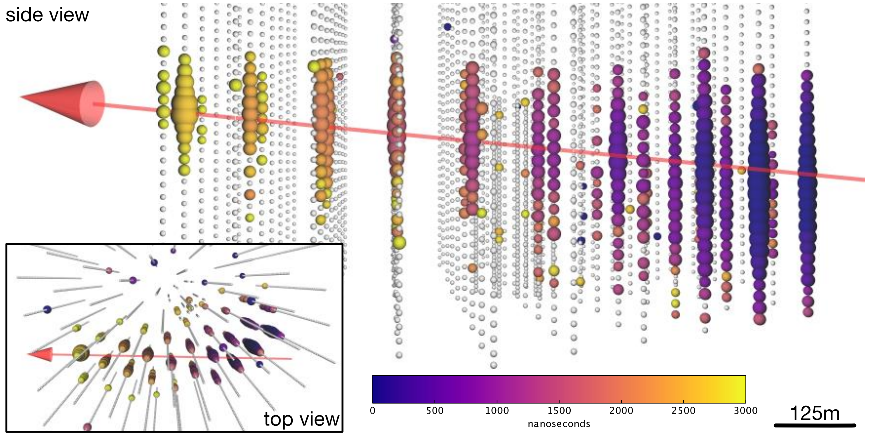

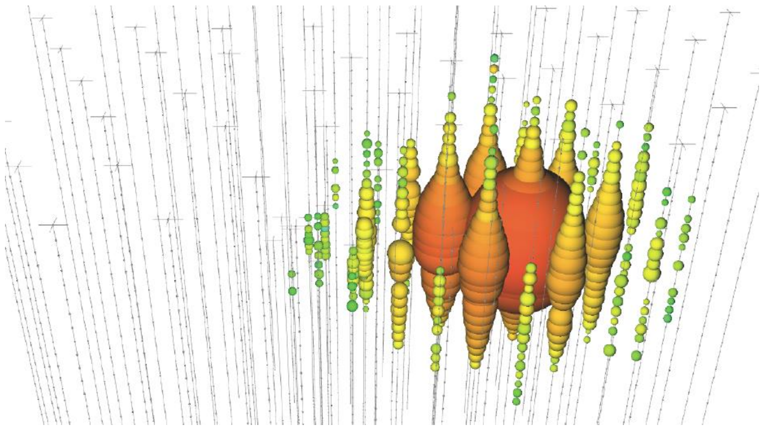

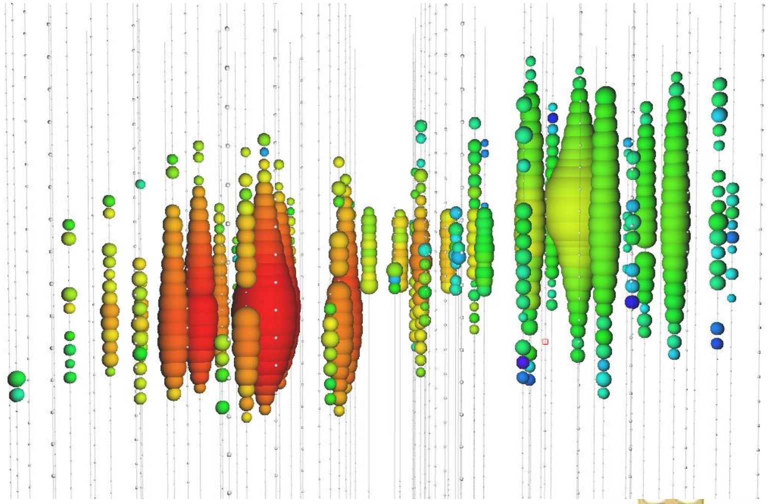

* We will come back to this

Applications – Particle ID and Classifcation

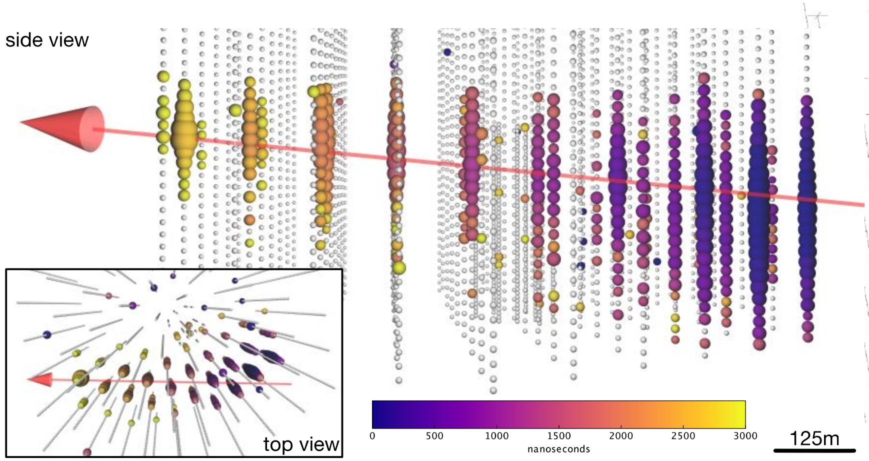

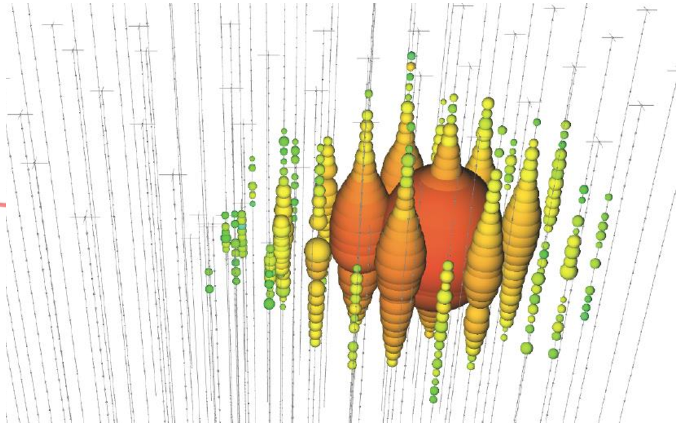

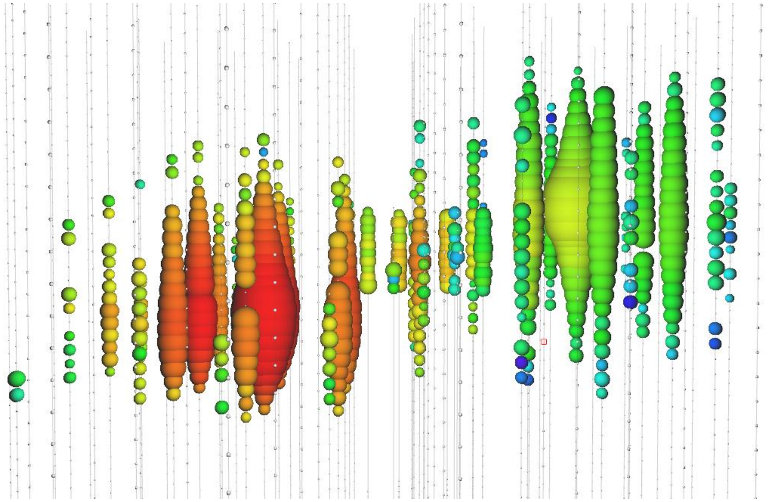

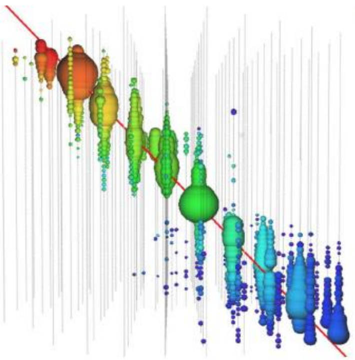

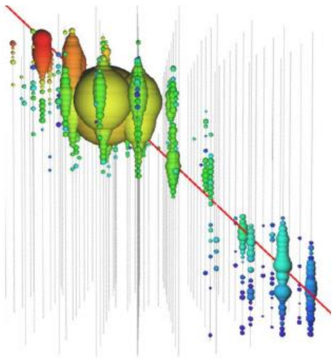

Track with energy of 26 GeV

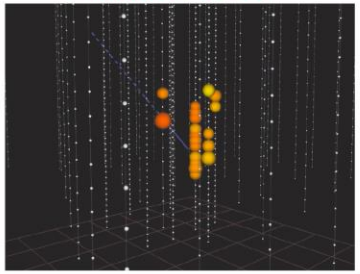

Cascade with energy of 30 GeV

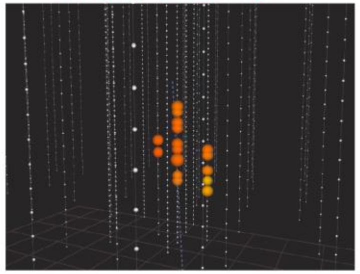

- Same particles can "look" different
  - Lower energy = Less Information
  - Higher Energy
    - Cataphoric/Stochastic energy losses
    - Relativistic Boosting "delays" particle decays – High energy Tau
  - Different ice models
- Subtle differences can identify the particle

Astrophysical Tau Neutrino Search

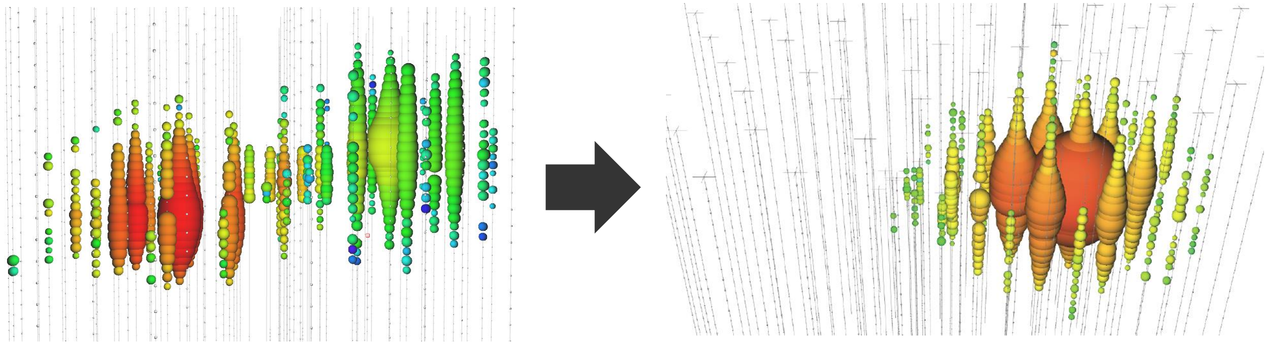

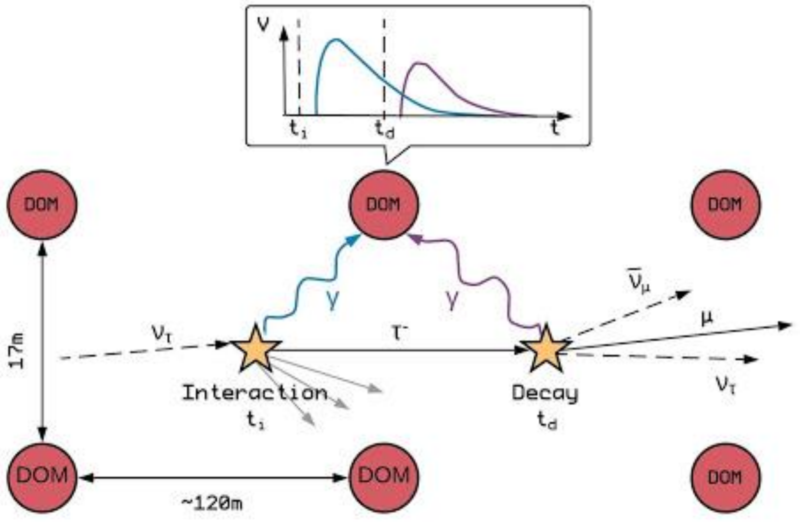

- TeV – O(1) PeV Tau neutrinos look like Electron neutrinos due to sparse instrumentation
- Differentiation by shape of waveform in a given module, i.e. two waveforms in the same module offset by a certain quantity
- Create an image (2D histogram) of the of the charge distribution in time along a string
- CNN used to find the subtle difference in waveform shapes

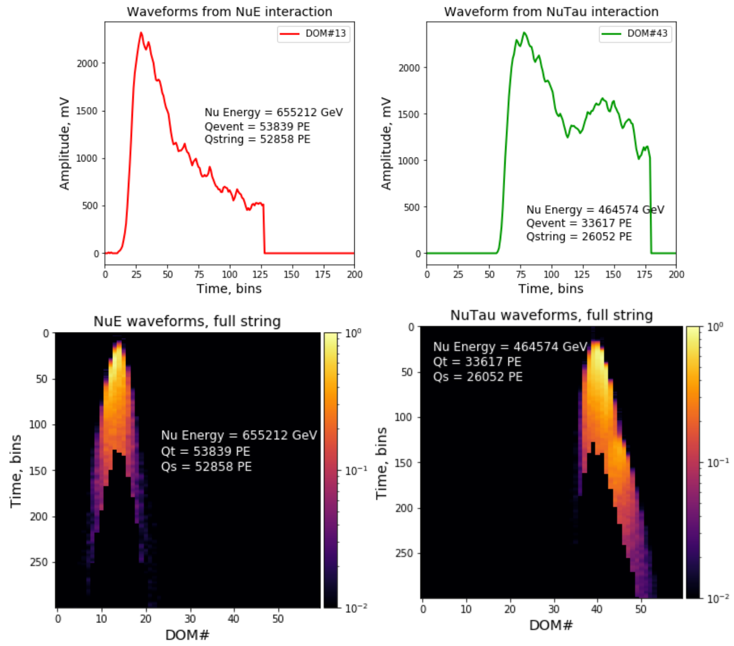

Combining Maximum Likelihood and NN

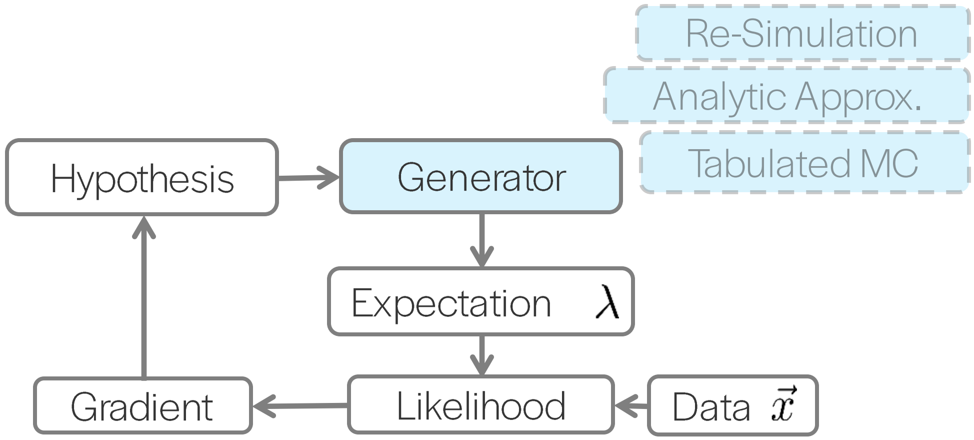

- Learn the PDF/likelihood
- Use Gradient Descent as a minimizer – More information than traditional minimizers
- Still have the errors from the maximum likelihood, but can encode the details that tabulated light yield can't
- Doing this in MMA is a next step

Southern Sky Neutrino Point Source

- Southern Hemisphere is astrophysically-speaking more interesting
  - More individual sources
  - More larger sources
- ...But there are issues for IceCube
  - Energy spectrum E^N, where N is -2ish, at >= TeV flux is tiny
  - No earth as a shield from atmospheric background for the Southern sky – Extremely high background that looks exactly like the background
  - Estimating track-like event energy is difficult, so can't cut on energy
  - Flux of >= 10 PeV tau-like events (good pointing and energy) is extremely low
  - Cascade-like events have poor pointing resolution, but are the only realistic option

Southern Sky Neutrino Point Source

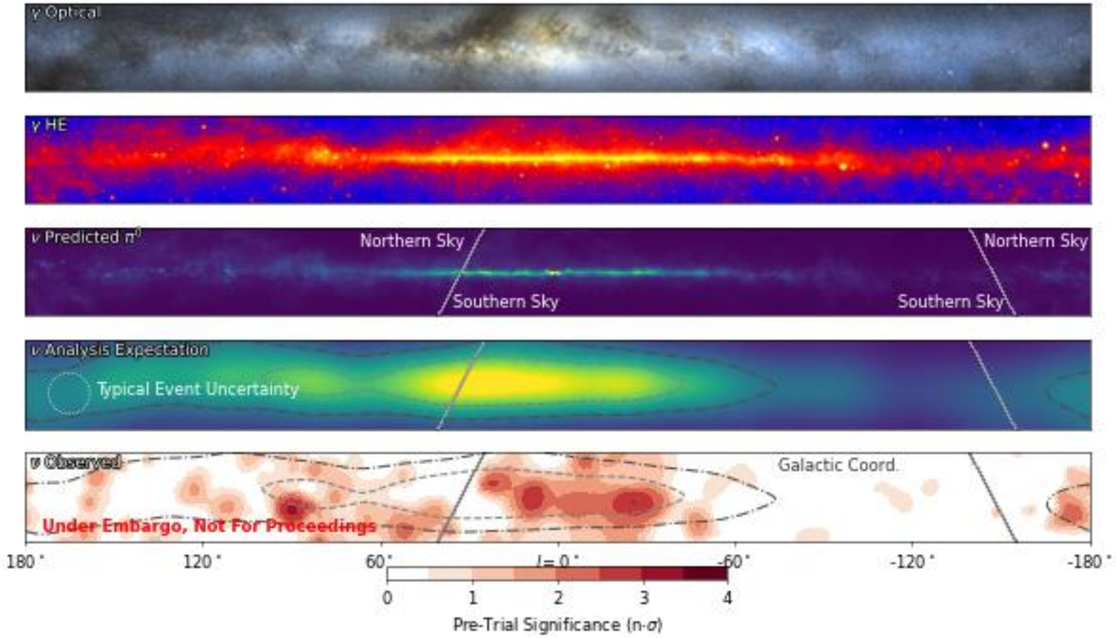

- Using Cascade events for “large” (e.g. Milky Way) or isolated sources a good option for Southern Sky point source search – Clear differentiation from background
- Maximum-Likelihood method for cascade pointing insufficient to find a source – Using BDT and CNN to find and reconstruct cascade events
- Third High Energy Neutrino Point Source! – Galactic Plane

Kaggle Competition

- Kaggle competition finished: https://www.kaggle.com/competitions/icecube-neutrinos-in-deep-ice
  - Top 3 Results all use some form of attention:
    - https://www.kaggle.com/competitions/icecube-neutrinos-in-deep-ice/discussion/402976
    - https://www.kaggle.com/competitions/icecube-neutrinos-in-deep-ice/discussion/402882
    - https://www.kaggle.com/competitions/icecube-neutrinos-in-deep-ice/discussion/402888
  - Better results than current state of the art!
  - Large Language Model technology gets best results
- Data is public!

IceCube Science – Multi-Messenger Astrophysics

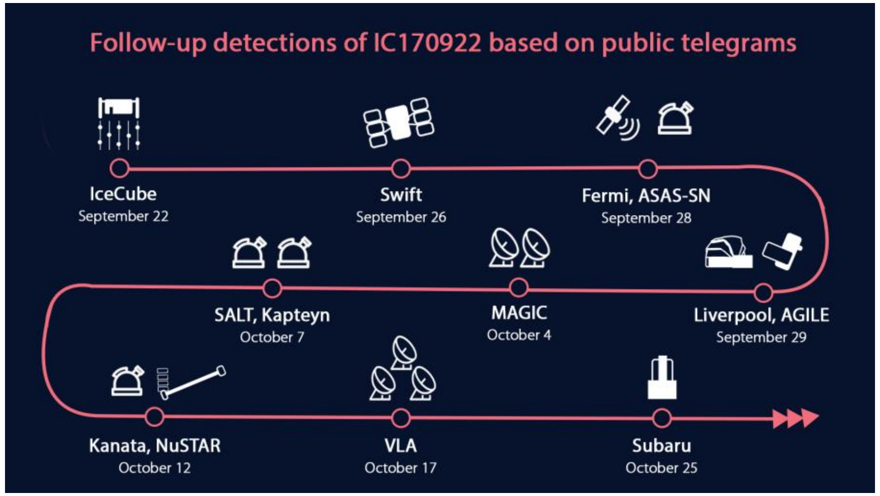

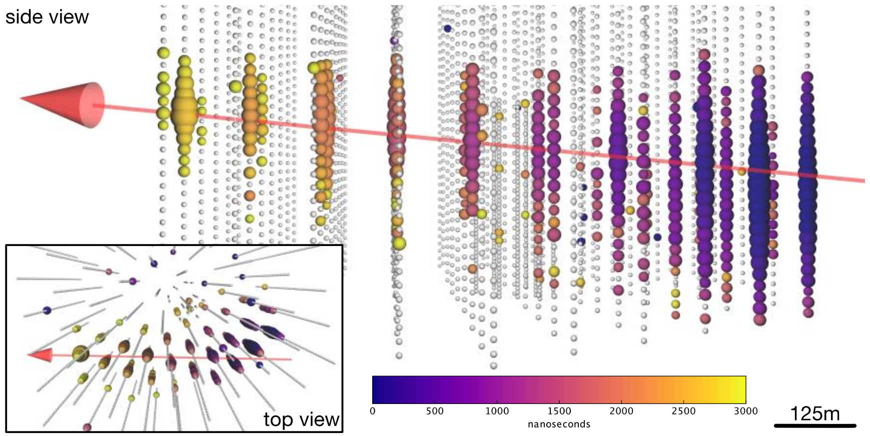

- Multi-Messenger Astrophysics (MMA)
  - Observing astrophysical phenomena with more than one "messenger" (gravitational waves, neutrinos, EM)
  - One of NSF's Big 10 Ideas
- IceCube detected an event that came from Blazar TXS 0506+056
  - Follow up observations by several observatories/telescopes showed signal
  - Back catalog showed access for this source
- **Fast response to alerts requires significant cyberinfrastructure**
- NS+NS merger would be ideal IceCube+LIGO observation
- Core collapse supernova would be ideal for IceCube+DUNE observation

Multi-Messenger Astrophysics

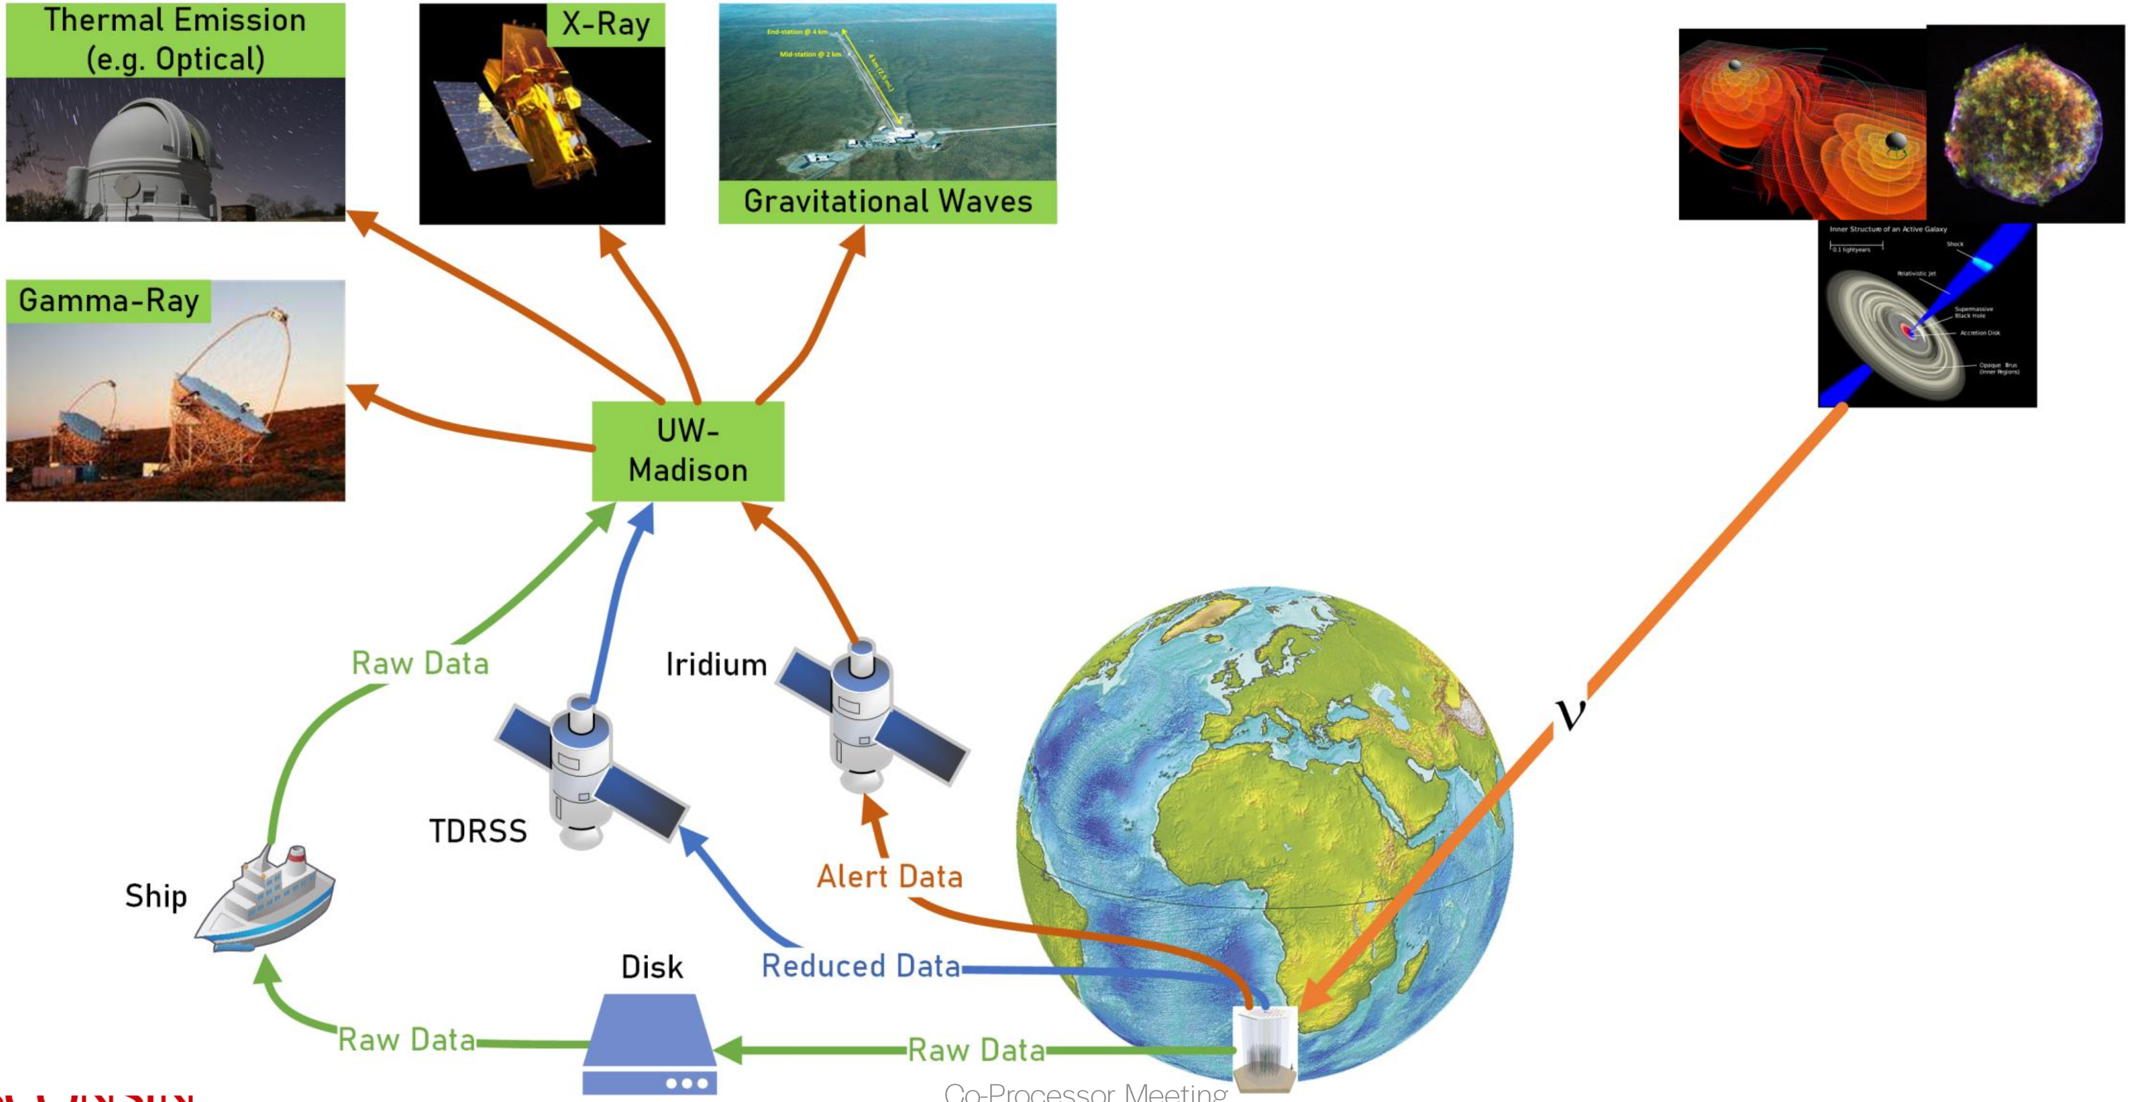

Multi-Messenger Astrophysics – Reconstruction

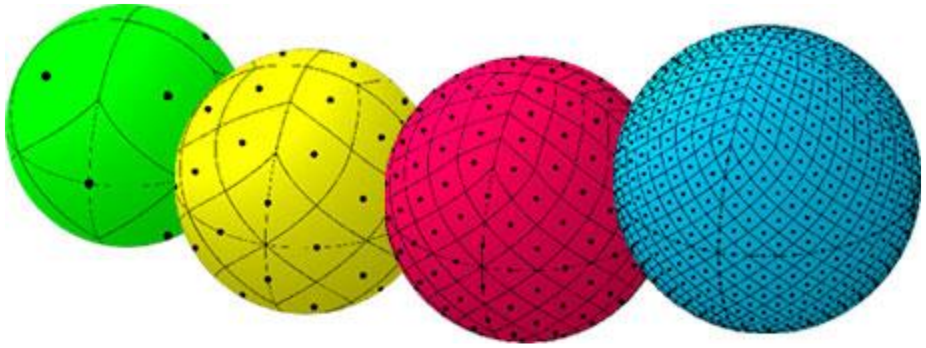

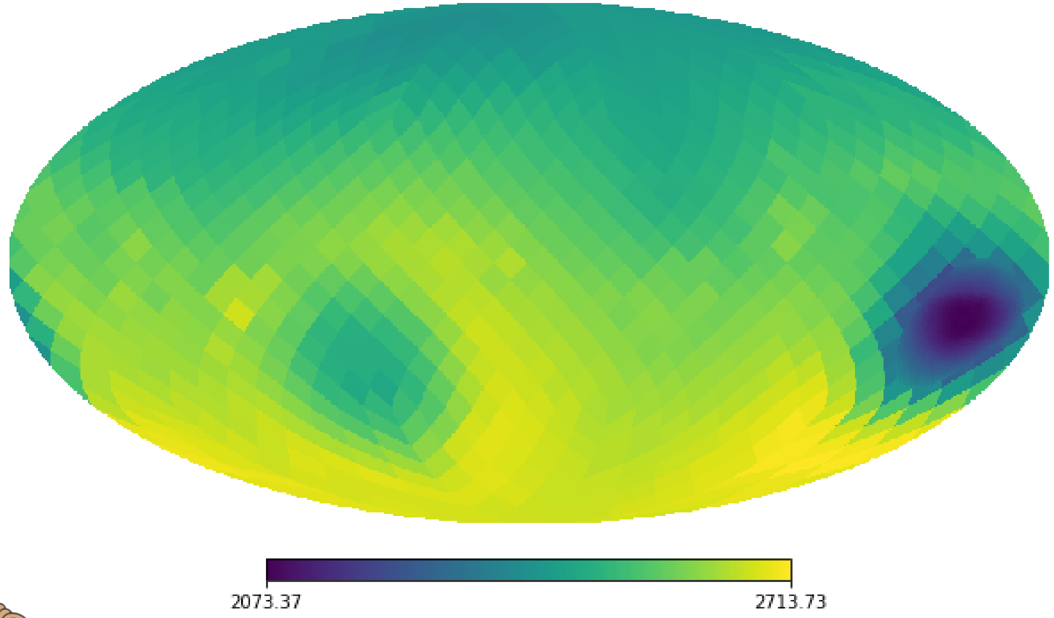

- Most accurate (and with well understood errors) directional reconstruction comes by scanning across the sky
  - Split sky into constant surface area pieces
  - Test each directional hypothesis against likelihood
  - Create directional likelihood map
  - Gives most probable direction and error
- Each hypothesis calculation is independent – Easy to split up workload across O(1000[000]) or cores

But… ML in MMA?

- What is the issue? – Errors
  - Errors from ML are poorly understood – Derived from simulation rather than statistically rigorous and “understandable” method
    - Is there a bias? Are we using the right systematics?
    - We know the simulation doesn’t describe the detector 100%, so can ML?
      - Poor training data in, Poor results out
  - In some events multiple sources are within in 90% bounds

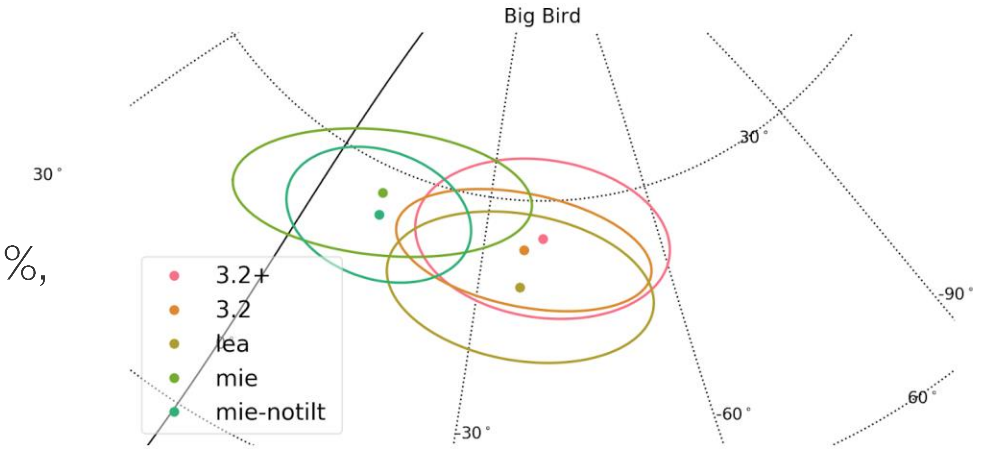

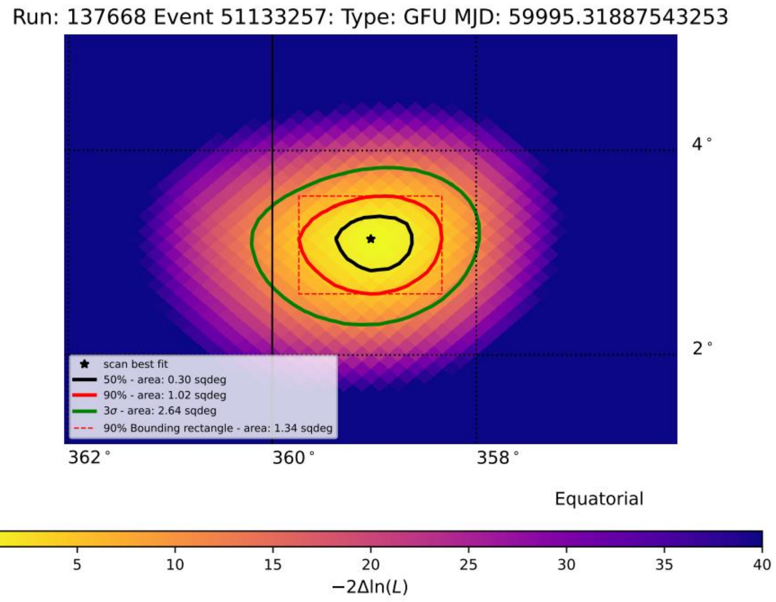

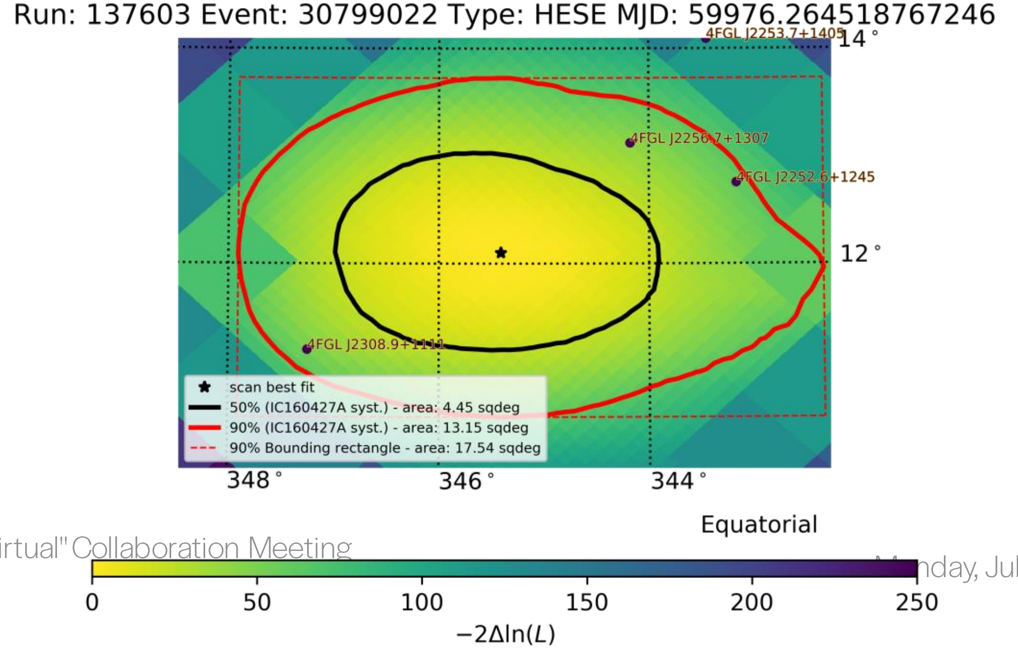

Bottlenecks

- Most MMA sources are transient
  - Fast follow-up is essential
  - 2 non-transient sources: Galactic Plane and NGC 1068
- IceCube pipeline is ~30 seconds from data acquisition to preliminary result for IceCube-initiated events
  - Partner-initiated takes longer due to network limitations to South Pole
- Limited hardware/power overhead at the pole – Accelerators and ML/AI can **save power and speed up computation**
- Biggest computing sink is waveform unfolding/deconvolution - 30-50% of CPU time (closer to 50% for MMA events)

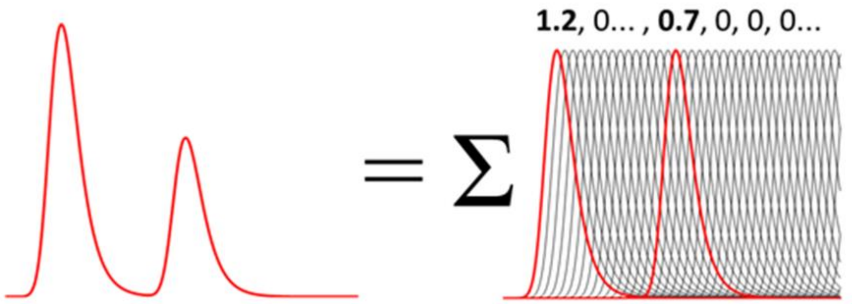

$$\begin{pmatrix} y(t_1) \\ \vdots \\ y(t_m) \end{pmatrix} = \begin{pmatrix} f(t_1 + z_1 dt) & \cdots & f(t_1 + z_n dt) \\ \vdots & \ddots & \vdots \\ f(t_m + z_1 dt) & \cdots & f(t_m + z_n dt) \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

ML unfolding

- Use an MLP to find pulses in PMT voltage waveforms
  - Scan over waveforms for more compact size
- Ability to find SPEs with a comparable performance to standard template fitting method
- Able to distinguish separate pulses 12 ns apart
  - Will be able to improve
- Eventually plan to implement on GPU and FPGA

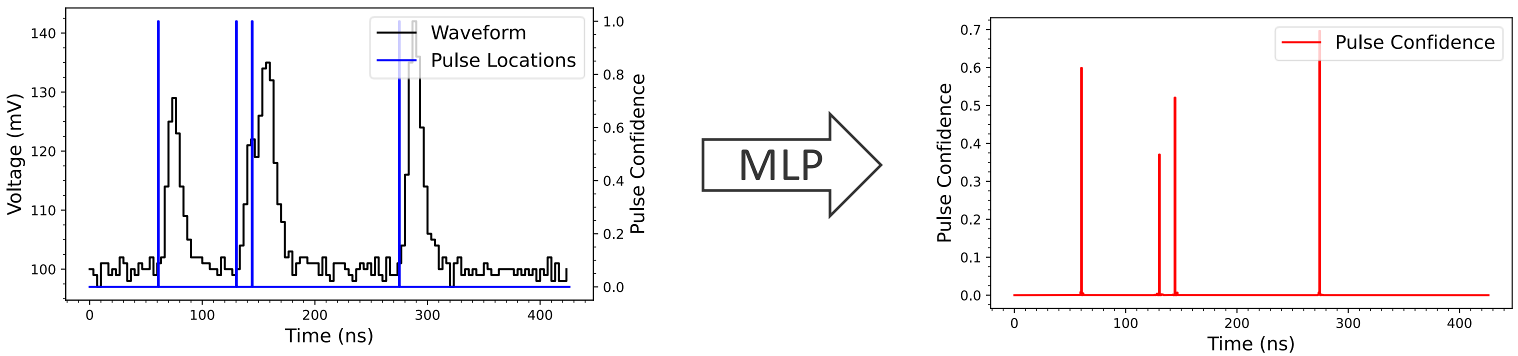

Performance with SPEs

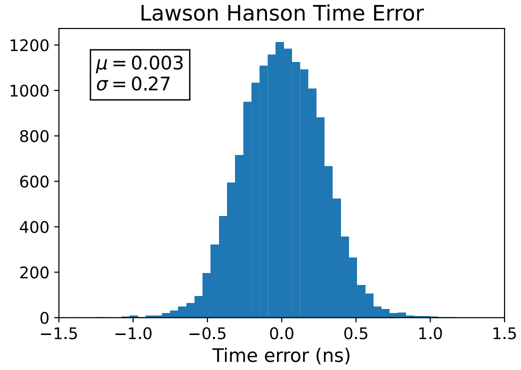

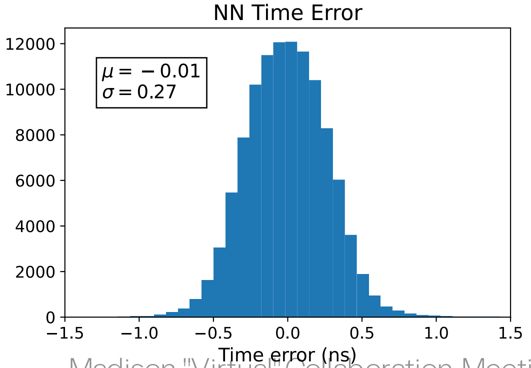

Two Pulse Performance

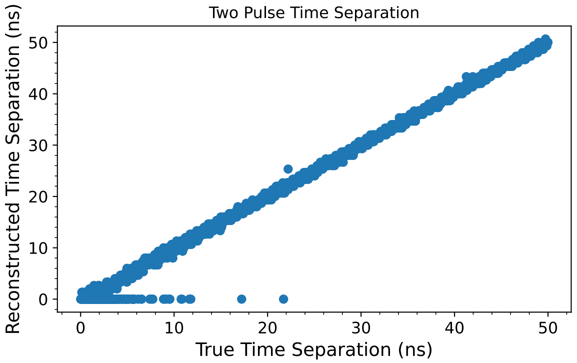

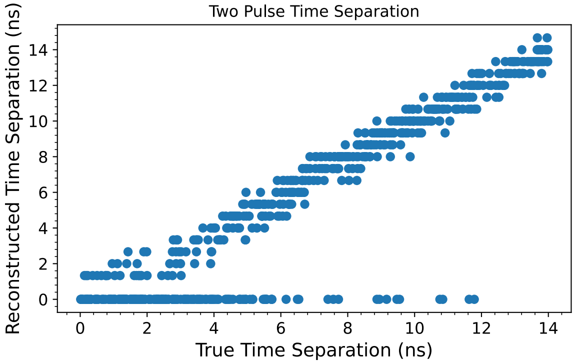

GPU unfolding

- Currently use the Lawson Hanson non-negative least squares algorithm to fit SPE template functions to voltage waveforms
- Our current algorithm runs on CPU, could perhaps speed up the algorithm by running on GPU
- Plan on adapting current C++ code to be compatable with CUDA

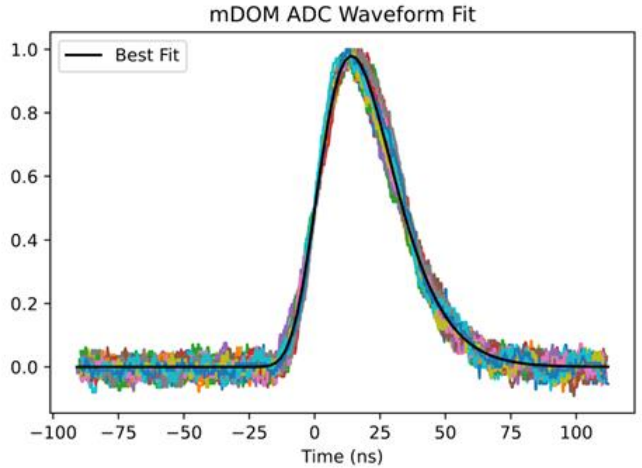

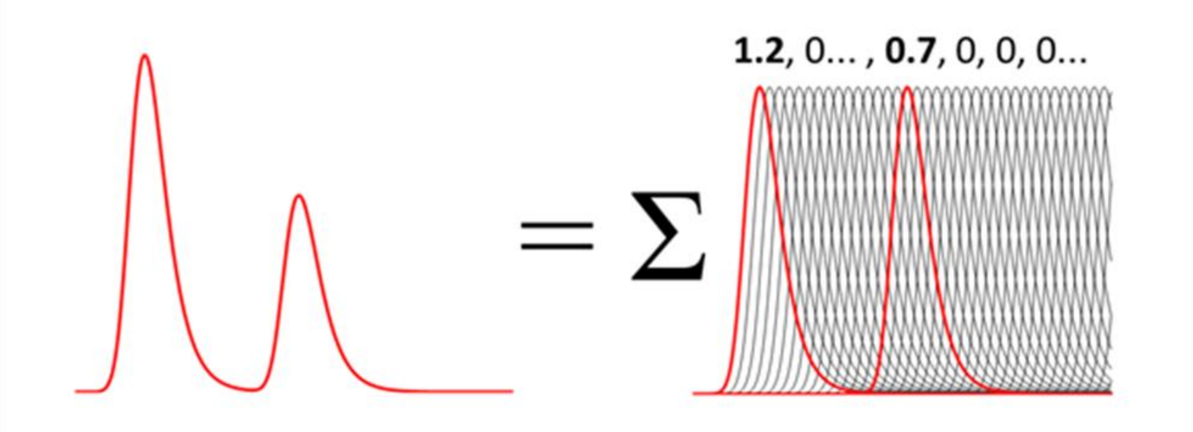

$$\begin{pmatrix} y(t_1) \\ \vdots \\ y(t_m) \end{pmatrix} = \begin{pmatrix} f(t_1 + z_1 dt) & \cdots & f(t_1 + z_n dt) \\ \vdots & \ddots & \vdots \\ f(t_m + z_1 dt) & \cdots & f(t_m + z_n dt) \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

Thank you!

Questions?

ML in MMA?

- Current pipeline takes ~30 seconds from data taken to alert
  - Biggest hurdle is extraction charge information from electronic readout
- Work on a single shot ML pipeline
  - From trigger to final selection in one step
  - Can we use this with even less information?
  - Can we accelerate the waveform unfolding?

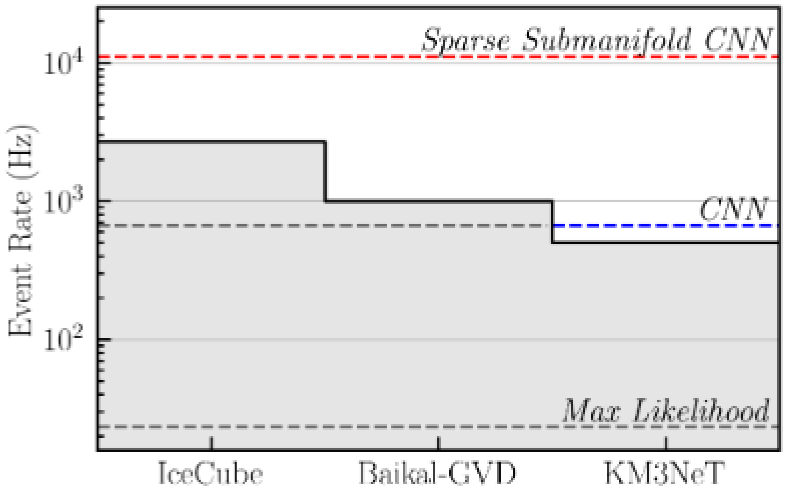

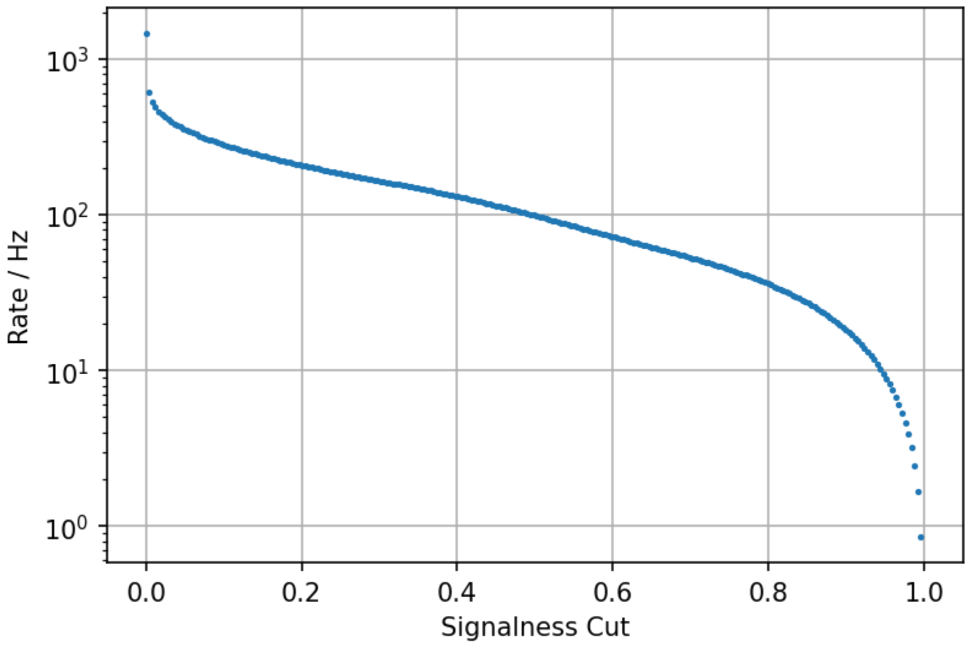

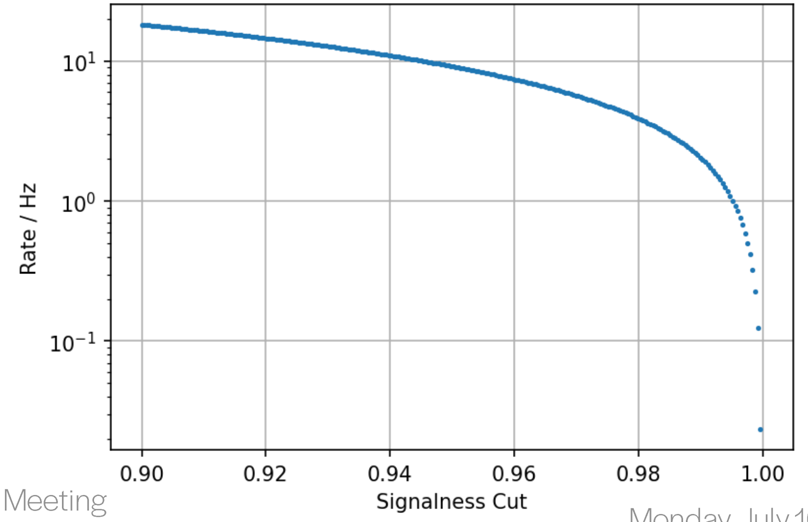

GNNs

- Neutrino Telescopes
  - Collection of individual detectors – Time-evolving 3D point cloud
  - GNNs a natural fit – Nodes (Individual detectors) with a connection (geometry of telescope)
- GraphNet
  - GitHub, IceCube Paper
  - Effort among neutrino telescopes to create a GNN framework and pre-trained models for reconstruction, particle ID, etc.

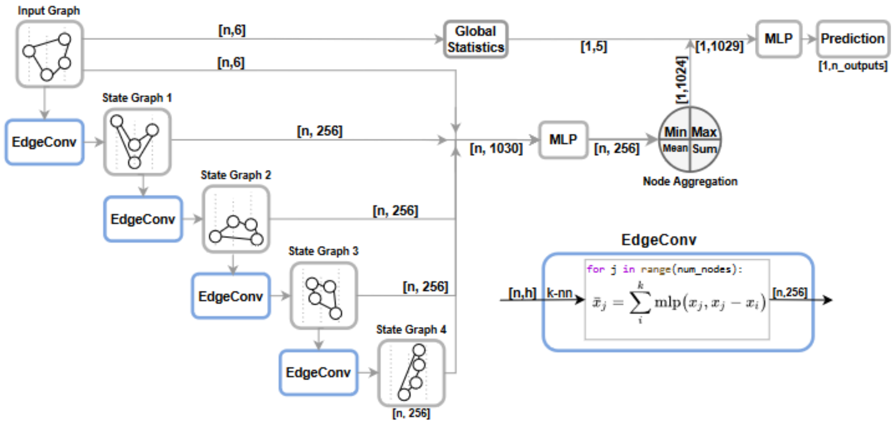

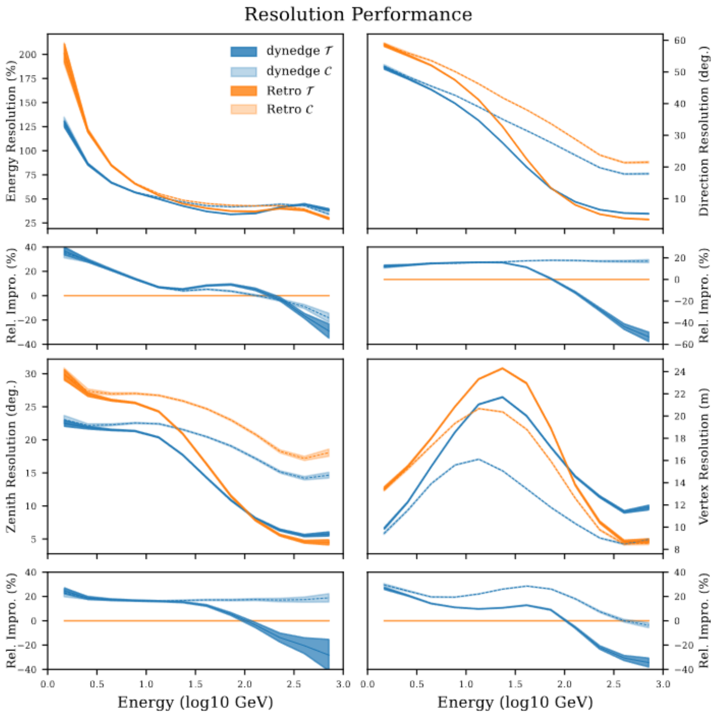

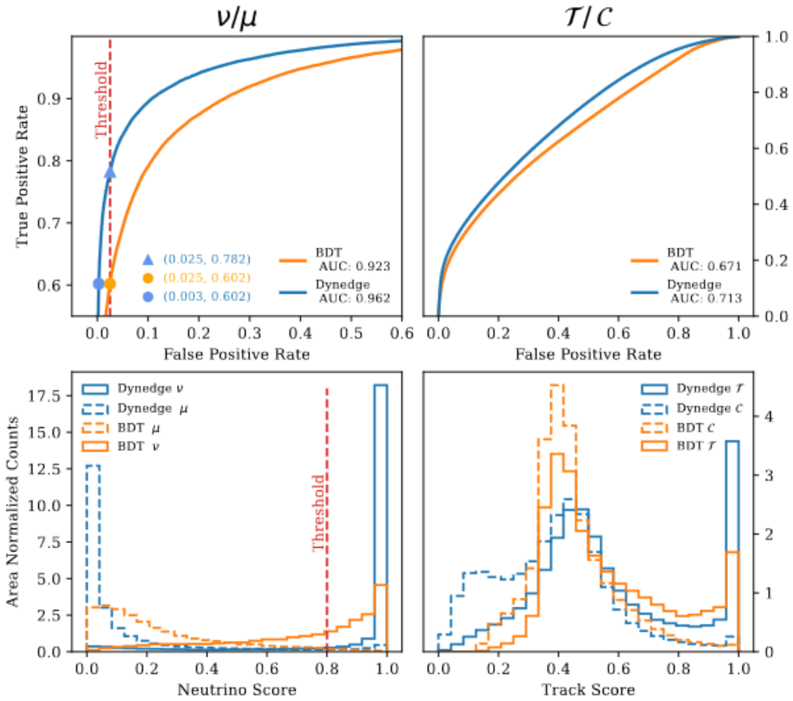

Maximum Likelihood – Short Comings

$$\vec{\theta} = (x, y, z, \varphi, \theta, E, t)$$

$$\mathcal{L}(\vec{x}|\vec{\theta}) = \prod_i p(x_i|\vec{\theta})$$

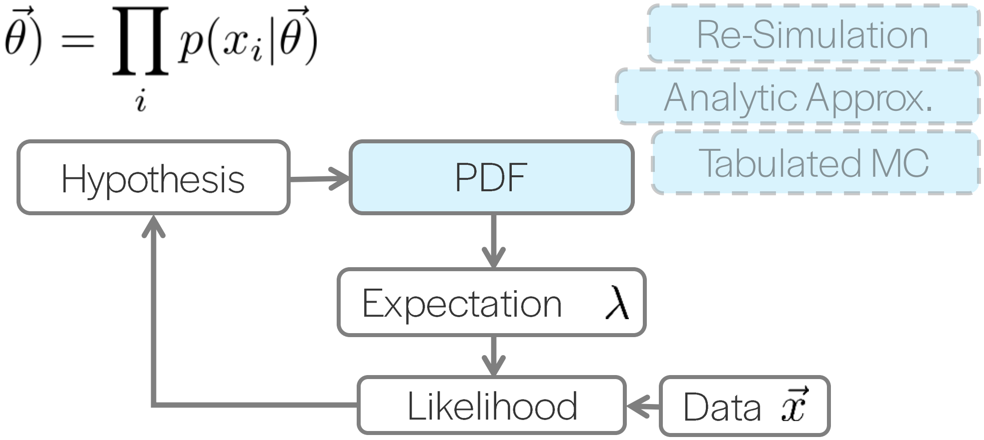

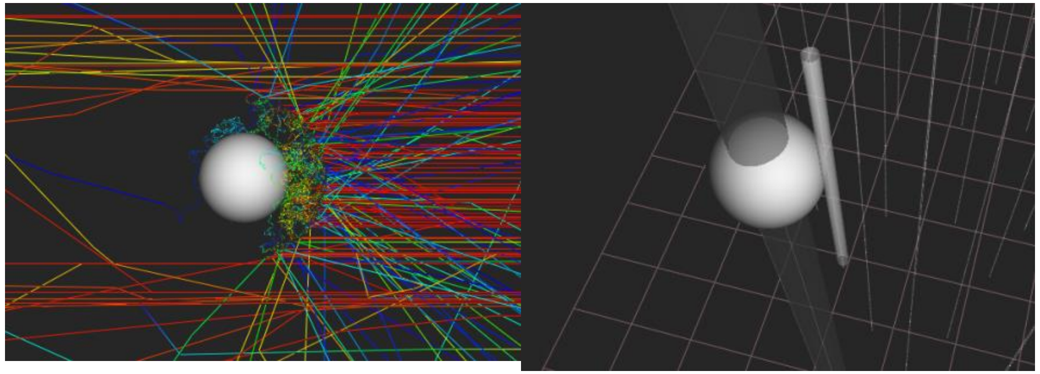

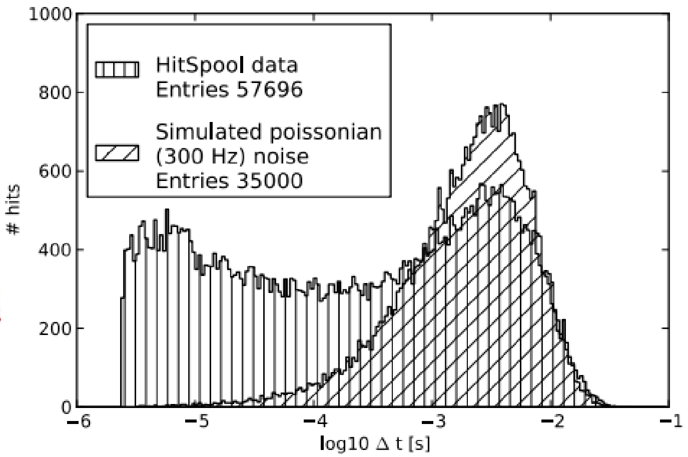

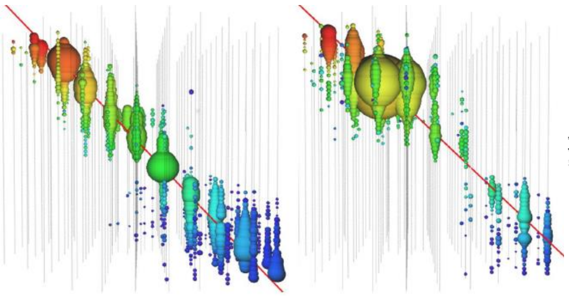

- Largest issue for reconstructions is the estimated light yield
  - At high energies – Ice and catastrophic energy losses
  - At low energies – Ice and non-poissian noise
  - Ice
    - Low energies – Area close to individual detectors
    - High energies – Bulk properties

Applications – Splitting

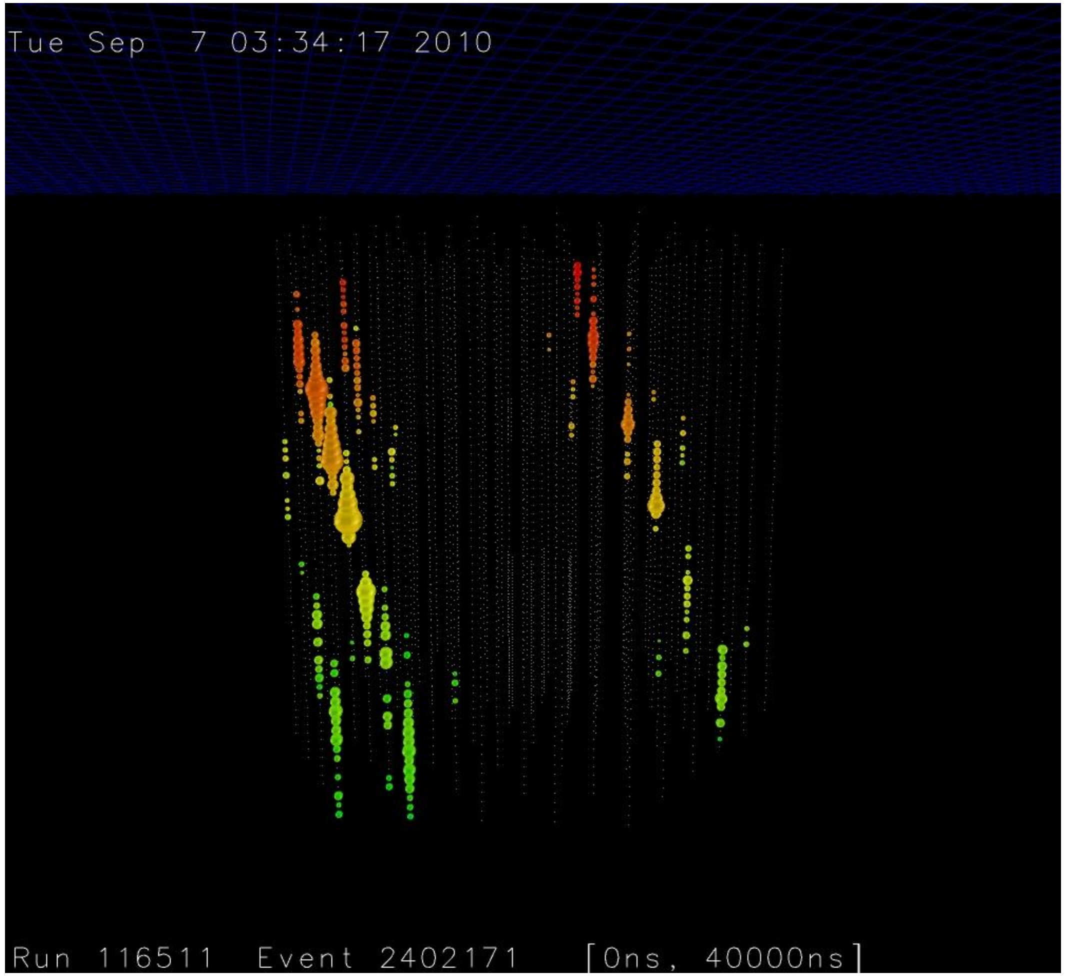

- Coincident events – Two or more events in the same trigger window
- Semantic Segmentation Nets – Question how to differentiate objects of the same class

Applications – Splitting

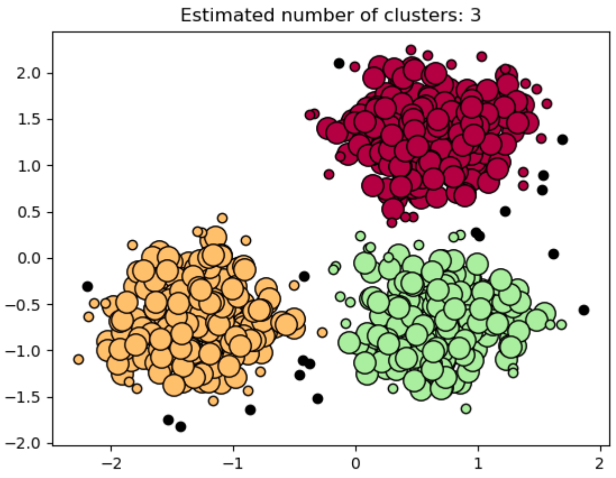

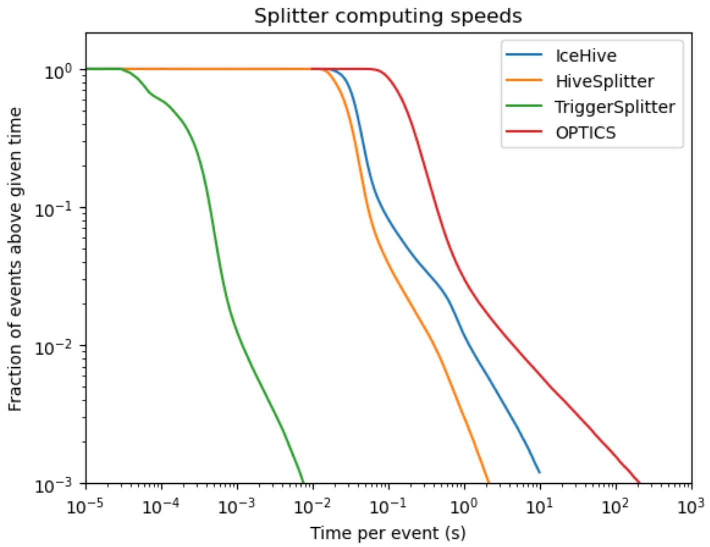

- Unsupervised clustering algorithms have proven either too slow or not as performant as existing algorithms
- There are advantage in noise identification with unsupervised learning – Needs to be revisited in new detector designs
- Looking at GNN-based clustering

D-Egg

mDOM

Applications – Simulation

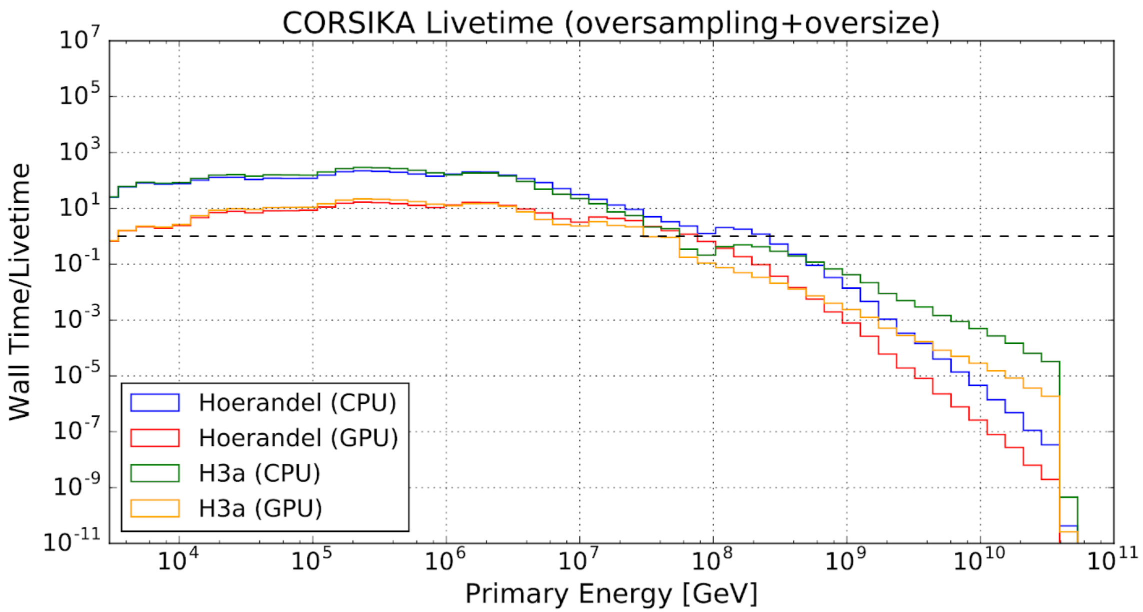

- Physics is the issue here
  - Lots of background
  - Lots of possible physics
- Enormous energy range and flux – **3 kHz of muons across 9+ orders of magnitude in energy with running at 90+% uptime for 10+ years**
- Desire is **100+ years** of simulation
- Up to **500x the compute time compared to livetime**
  - 2 seconds of detector runtime:
    CORSIKA Showers: 1M events
    CORSIKA Output: 92K events
    Photon Prop/Ray-tracing: 18K events
    Detector Sim output: 10K events
    Final output: 5897 events
  - Anywhere from **34-98% of CPU/GPU time is spend on first principal particle simulation**

Applications – Simulation

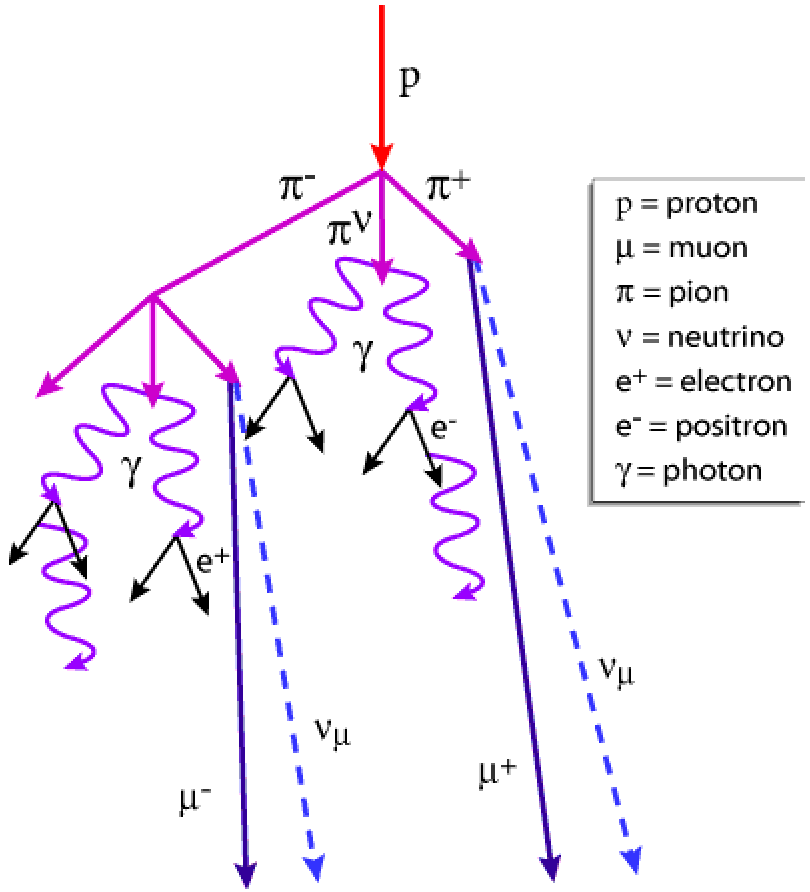

- First order simulations are good for deep dives
  - Not good if you need to generate lots of it
  - Edge cases important
- Existing parametrizations are lacking – Missing "muon bundles" (groups of related particles)
- Potential CPU saving solutions use flawed assumptions – Local network connectivity needed

IceCube Events

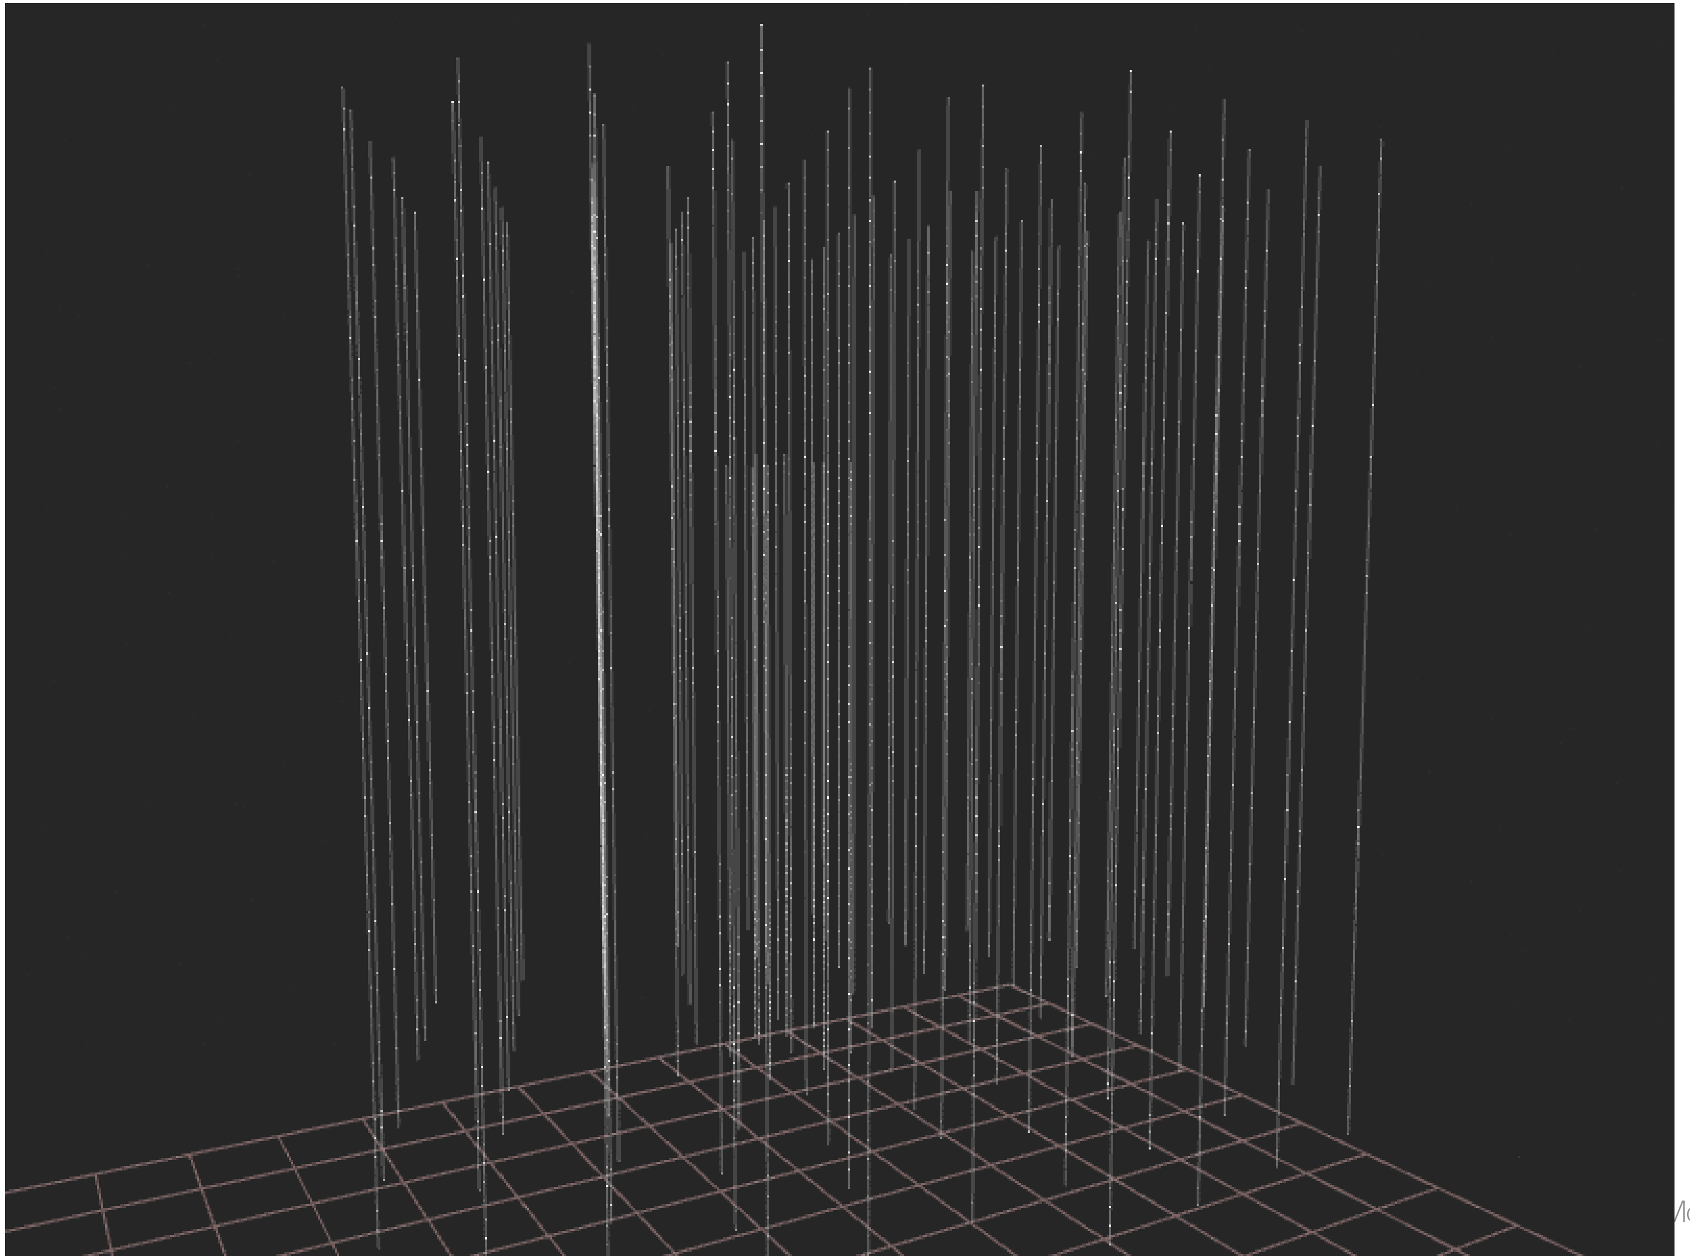

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Ice Model

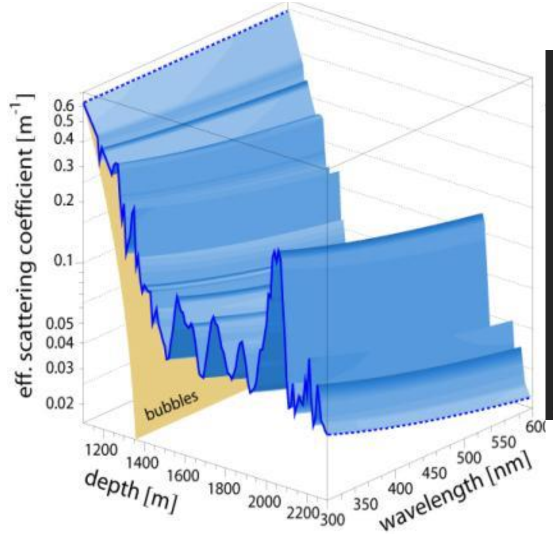

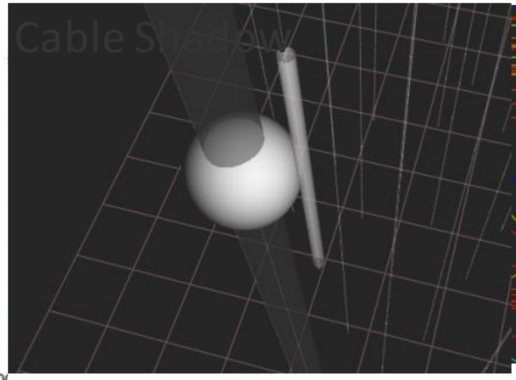

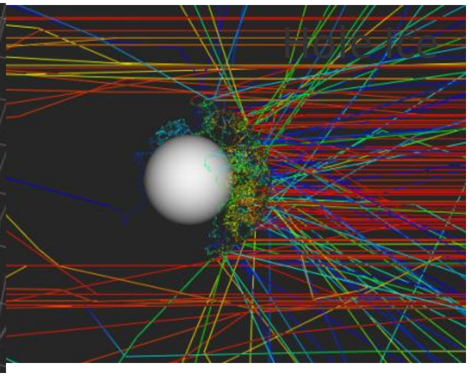

- Natural medium – Hard to calibrate properly
- Dropped a detector into a grey box
  - The ice is very clear, but...
  - Is it uniform?
  - How has construction changed the ice?
- Ray-Tracing with GPUs
  - Too complex for a parametrized approach
  - Needs brute force approach
- Drastic changes in reconstructed position with different ice model

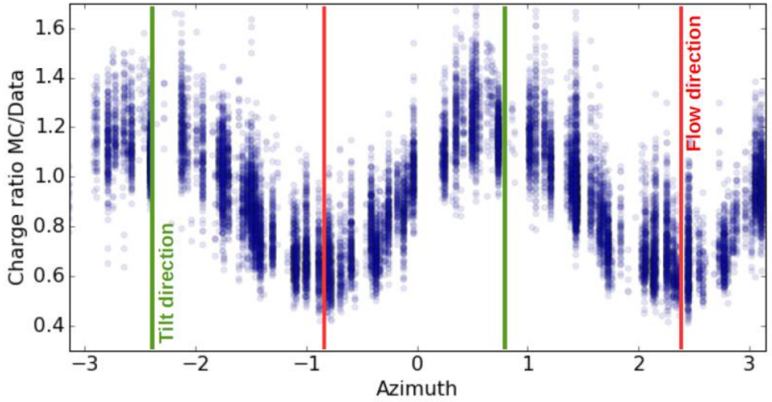

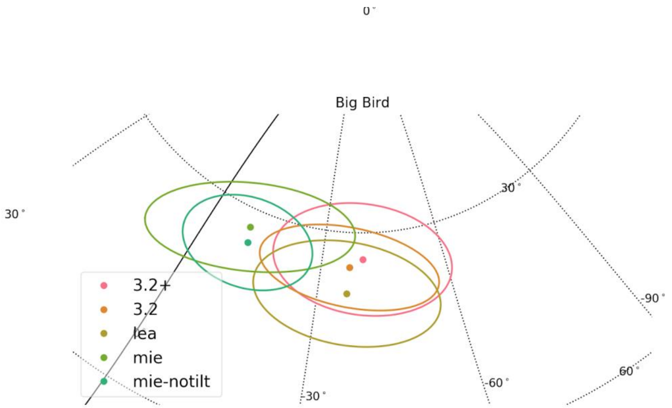

Cascade Events – Results with cascades and ML coming soon, stuck in journal

Ice Model

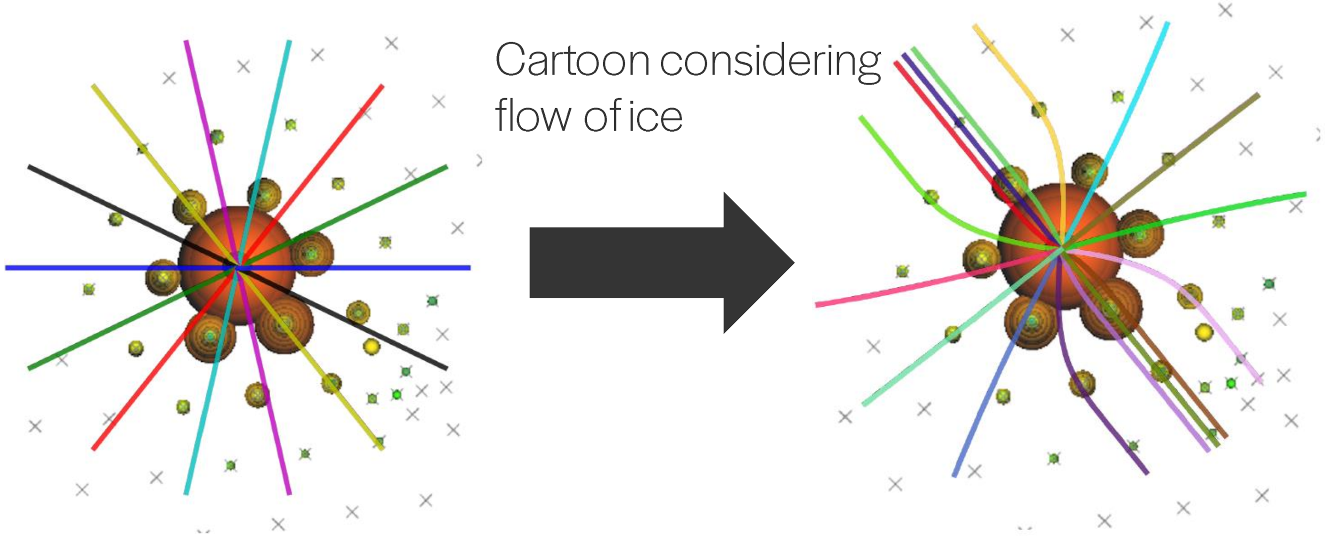

- New calibration data is available every year – More data being data, new ice cores being taken, etc.
- Want to have most up-to-date information in detector
- “Fitting” the model is compute intensive – Up to 400 GPU years
- Newer analyses want the most up-to-date simulation
- Ice model is the largest systematic effect in detector – Adding new systematic effects
- The deeper we dig the more complicated it gets
  - Birefringence – Ice Crystal Size
  - Layering of ice over millenia and moving across a changing bedrock

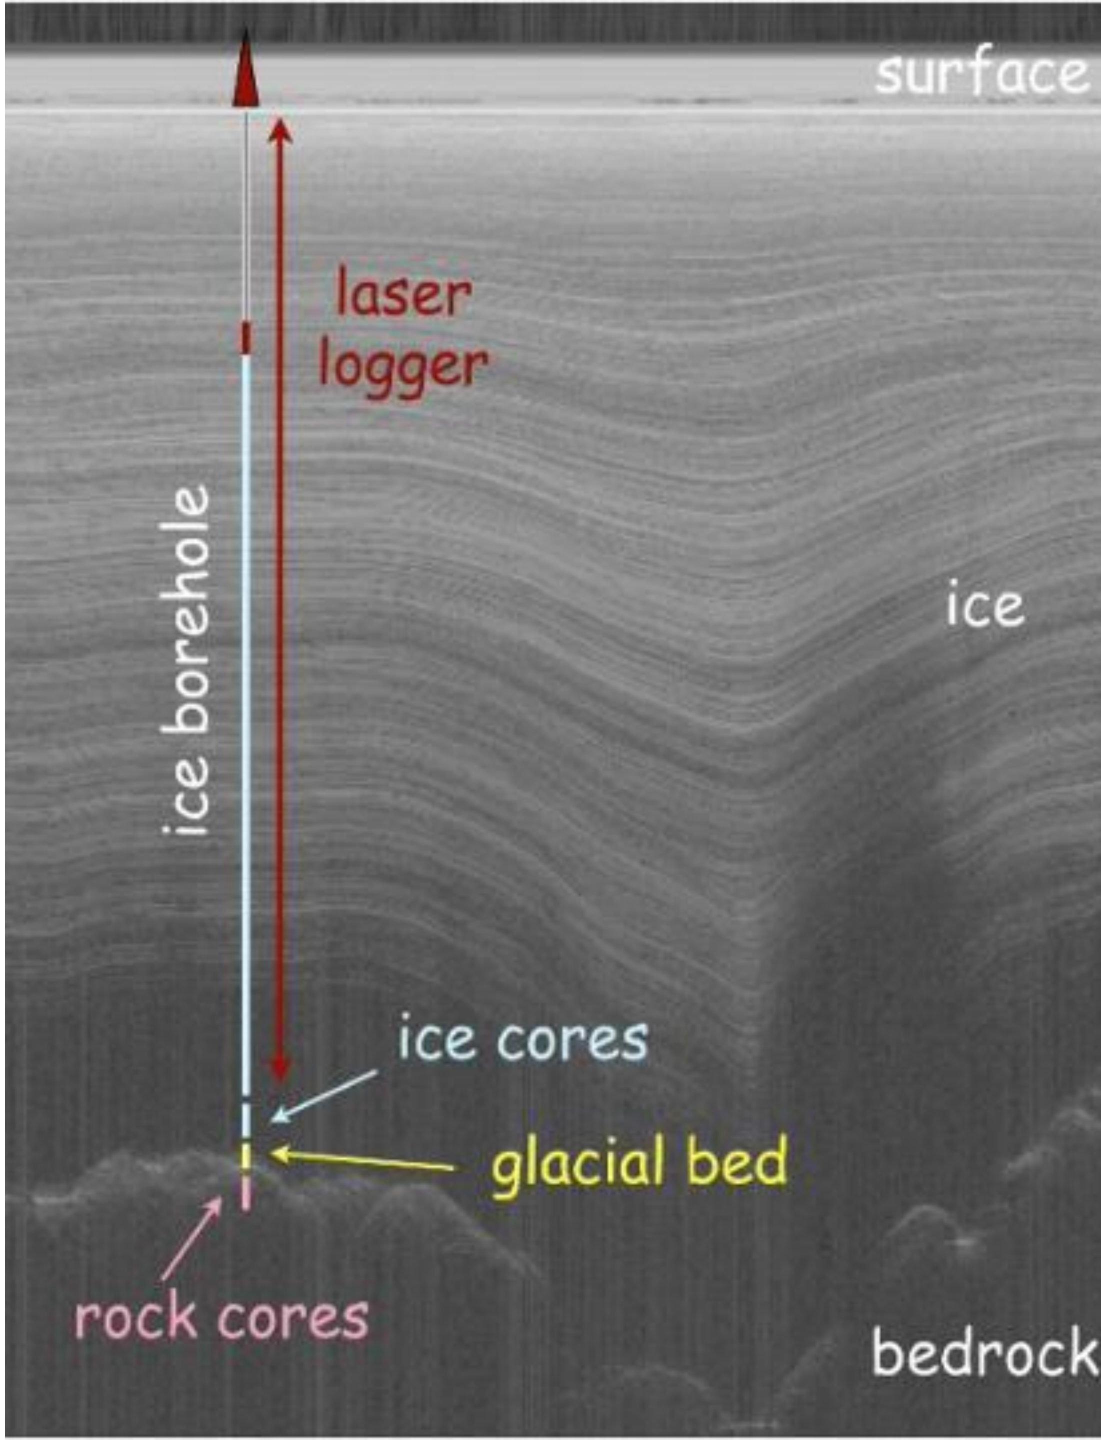

Classification

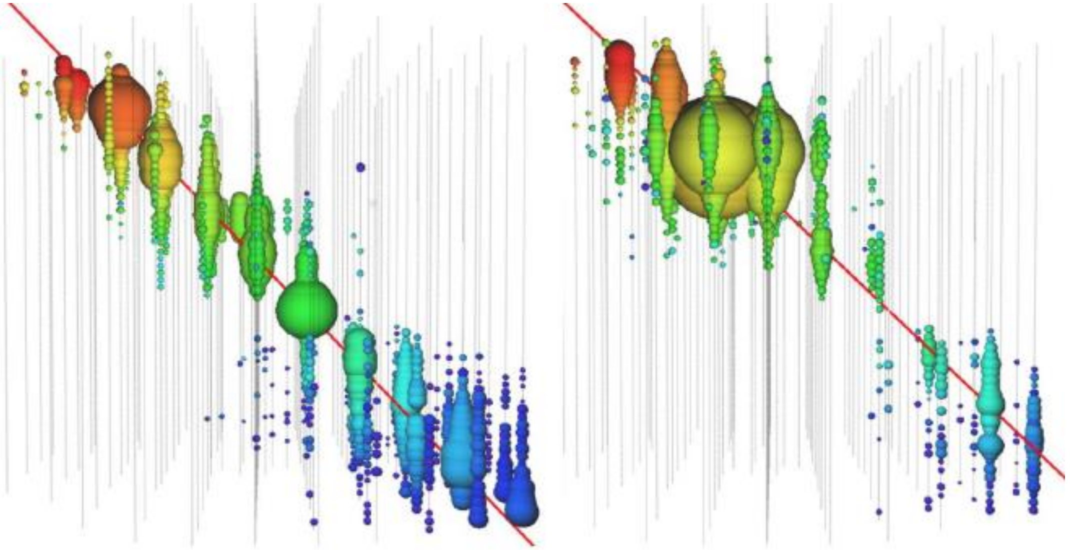

- CNN being added to “standard” processing to identify:
  - Through-going Tracks– Pass through entire detector
  - Stopping Track – Enters the detector and stops in it
  - Starting Track – Starts in the detector
  - Starting Cascade – Starts in detector and doesn’t exit
  - Skimming – Enters and exits at the corners of the detector
- Graph Neural Nets
  - IceCube is a collection of points in space with a relationship
  - Graphs are a “natural” fit
  - First signs of impact – Differentiation in events performs better than CNN
  - https://arxiv.org/abs/1809.06166

Convolutional Neural Networks

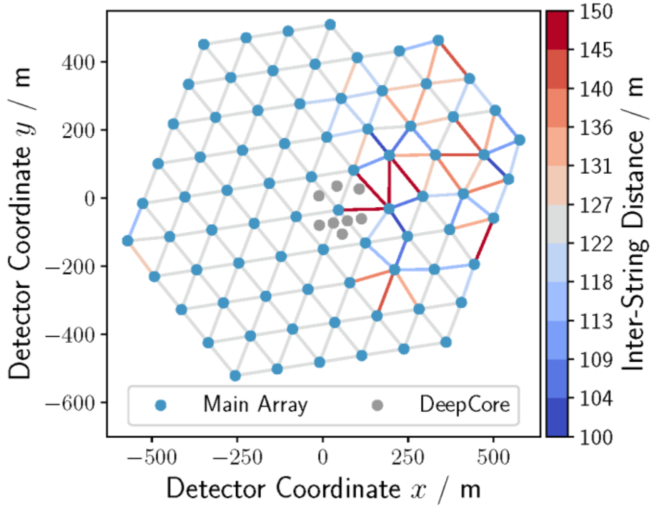

- Convolutional Neural Networks – Square Peg Round Hole Problem
- IceCube is not a CCD – Irregular geometry
- Time evolution of events - Contains information about detector, events, etc.

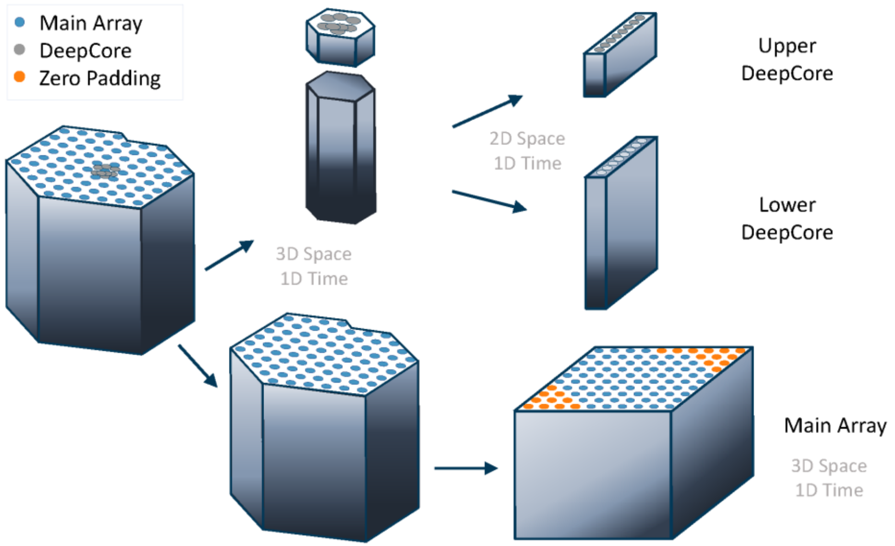

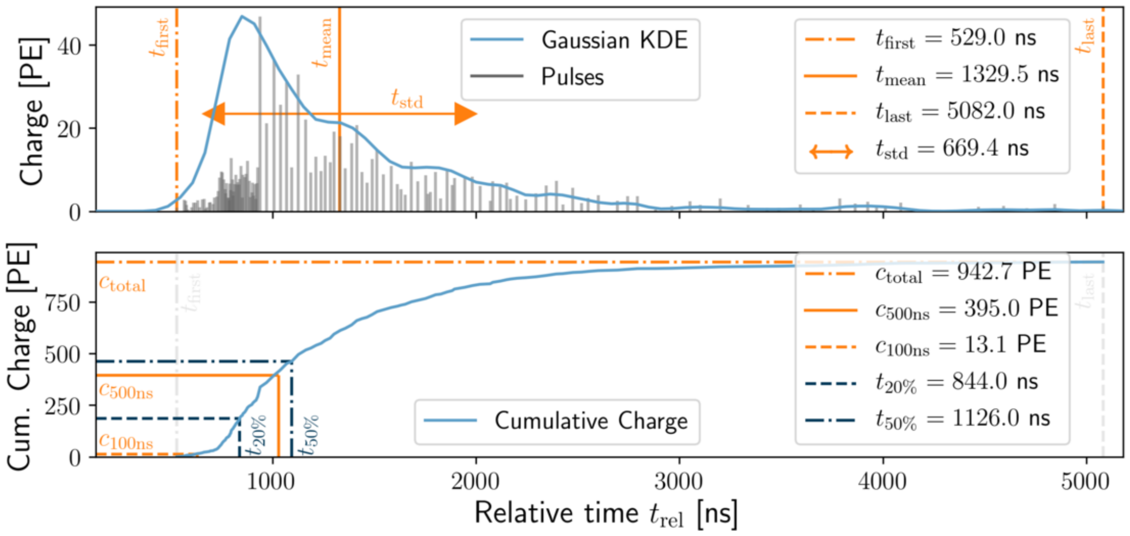

CNN Performance

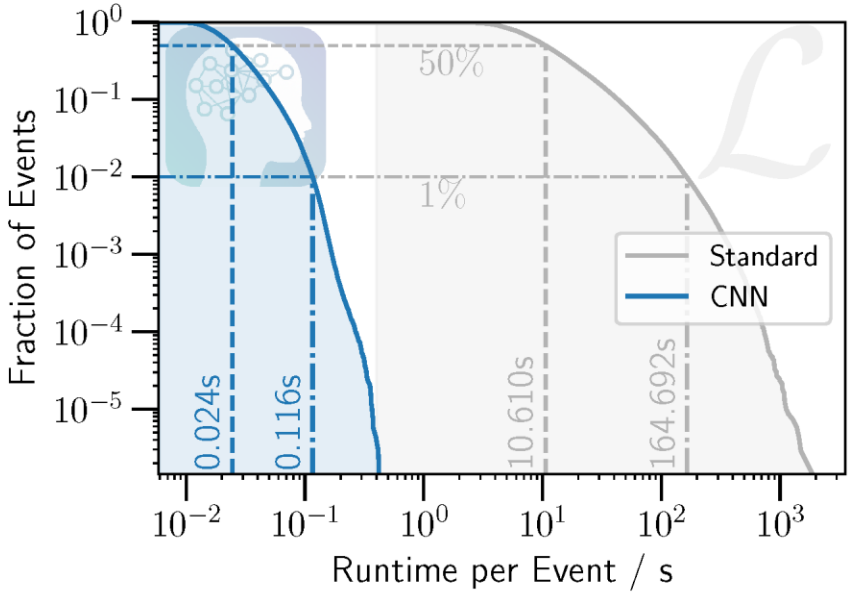

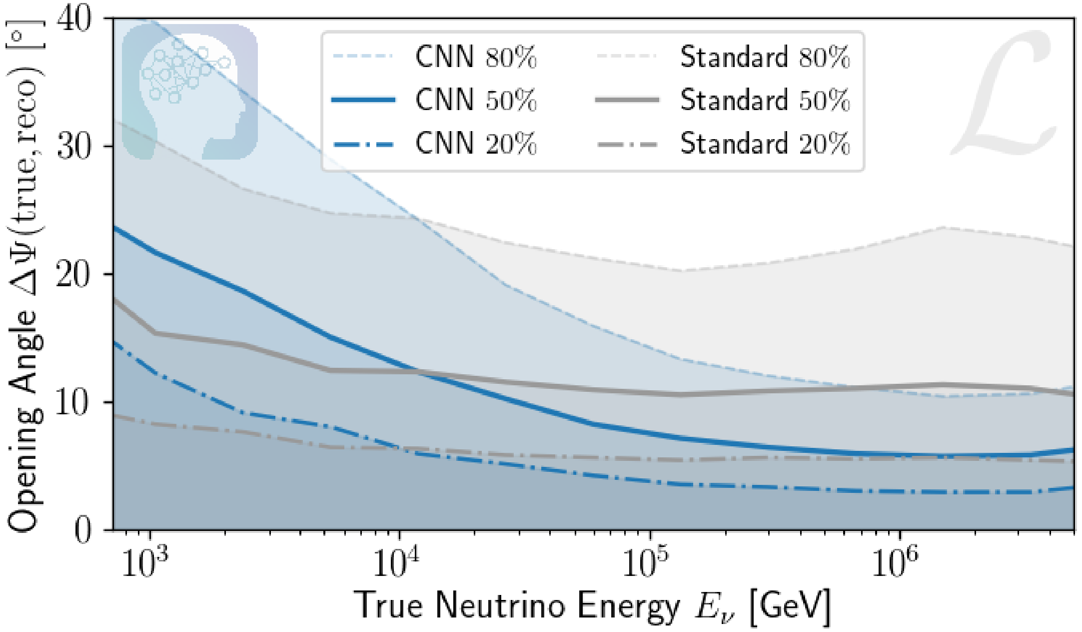

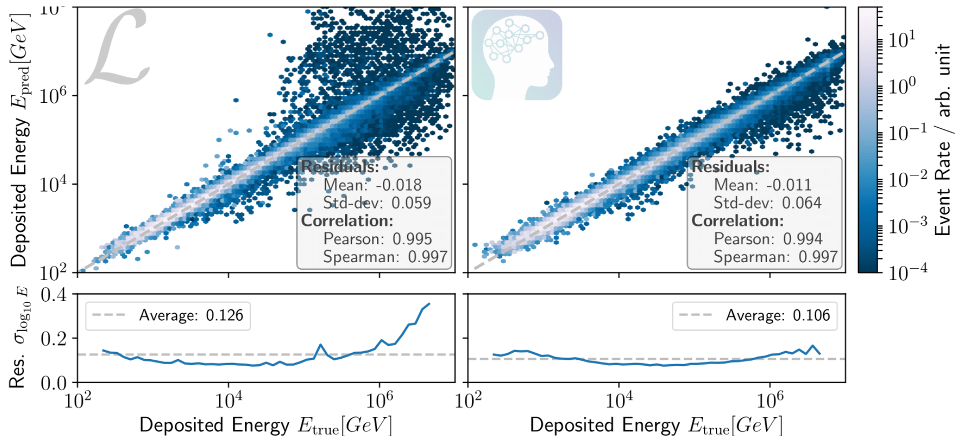

- Faster
- Better resolution

Today

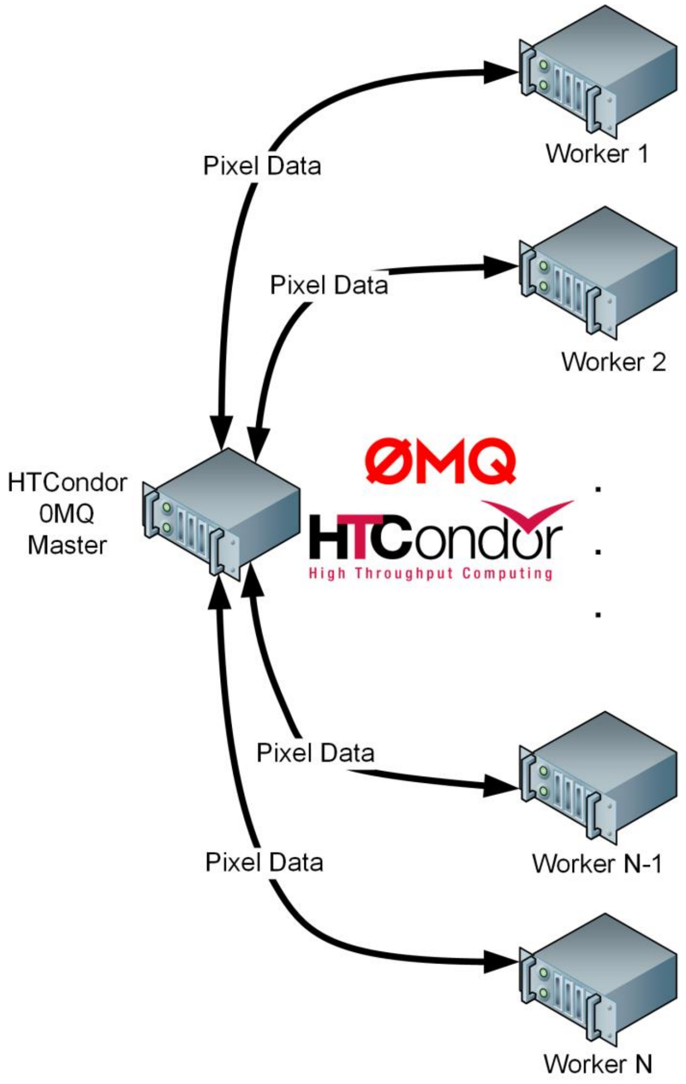

- Manager-worker setup
- Worker resource requirements are the same
- Manager makes decision about next scan
- Data communication via ZeroMQ
  - Easy to use and setup
  - Scaling an issue
    - Issue with communicating over external network – Firewall issues
    - Manager can't keep up > 2000 cores
- Using HTCondor for scheduling workers

CSSI – Event Workflow Management Service

- Introduce Manager-Worker paradigm to dHTC workloads – Map-Reduce for the Hadoop Afficinados
- Plain HTCondor not good for Manager-Worker
  - Schedule overhead dominates in execution time with small tasks
  - Use the right tool for the right job – HTCondor to aggregate resources and schedule workers
- Message Queues (MQ) have all the tooling needed to extend HTCondor
  - Handle many small data packages (messages) efficiently
  - Multi-user, multi-workflow separation
  - Persistence (if needed)
  - Offload to storage (only Apache Pulsar)
  - Monitoring
  - Off-shelf – No first principal derivation

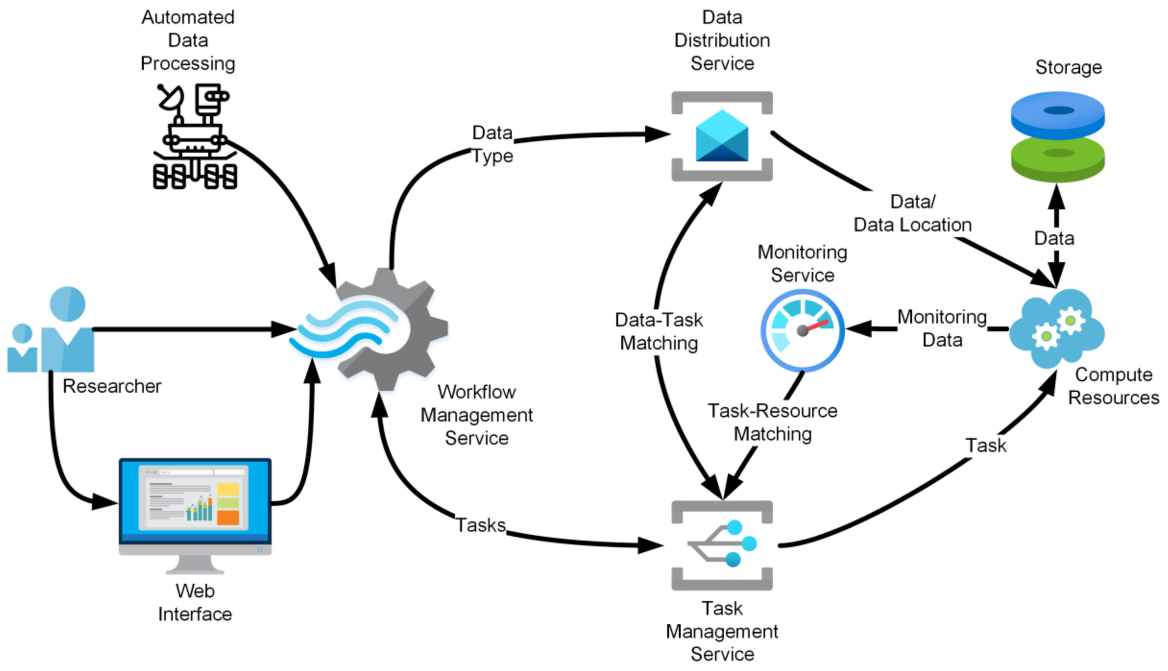

Tomorrow – Grid

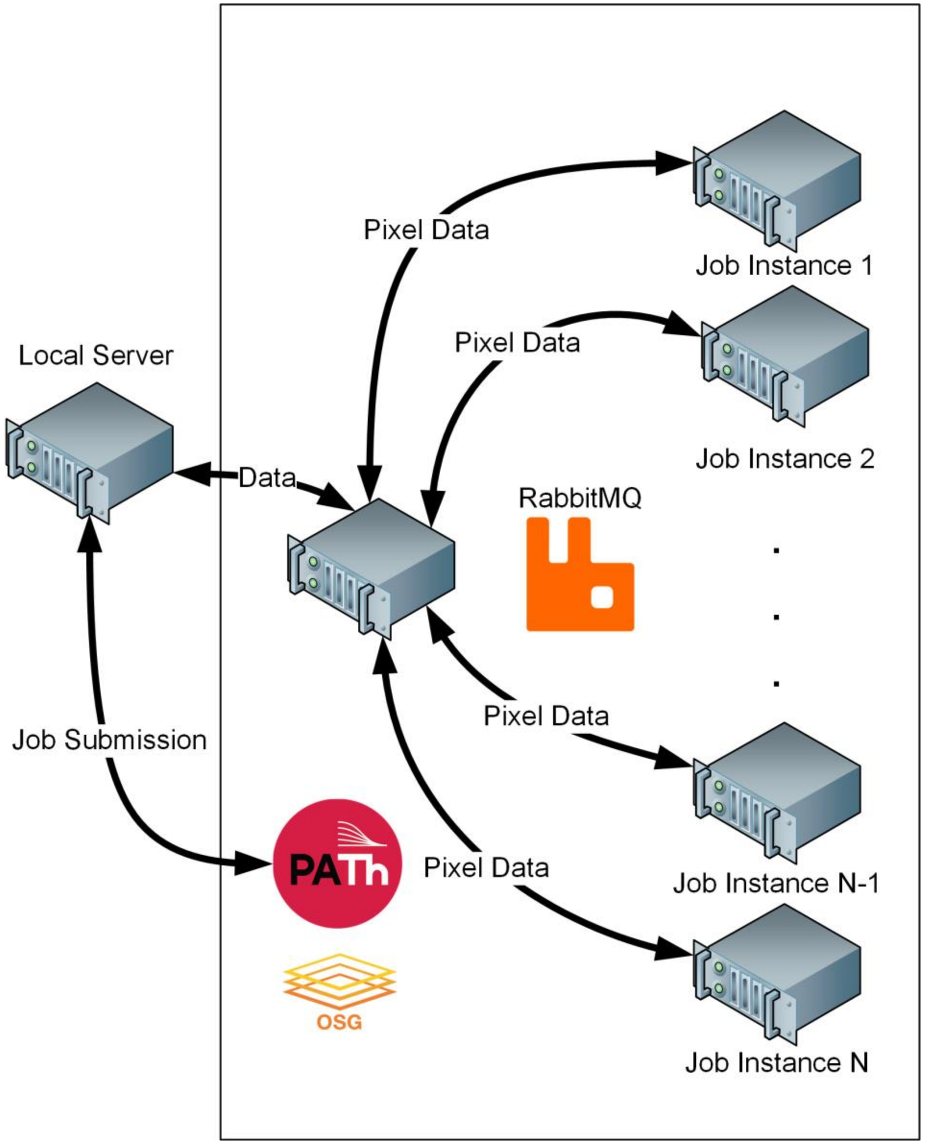

- Replacing 0MQ with RabbitMQ – Why?
  - Build for scaling
  - No network issues
- Why use the Grid?
  - More resources available – Faster completion
  - Free!
- Need an RabbitMQ instance running continuously – PATh facility to the rescue!

Tomorrow – Cloud

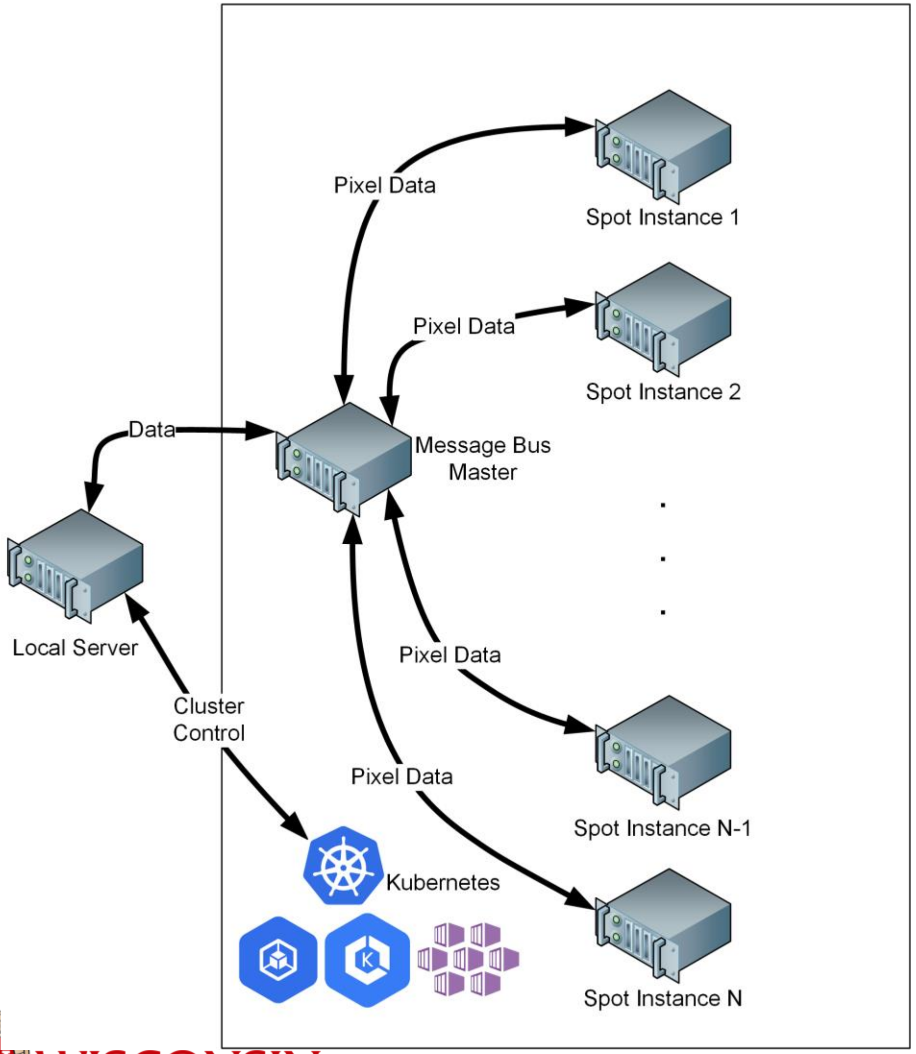

- Substituting Grid resources with Cloud resources
- Much larger pool of resources
- Readily available resources
- More control over resources
- Autoscaling cluster size (Kubernetes as a Service) makes things easier
- $$$ - Question about cost (~$100-1500 per alert)

SkyDriver – Reconstruction-a-a-S

- Researchers submit an event to a service be reconstructed through a scan
- Researcher does not need to worry about the details of the scan

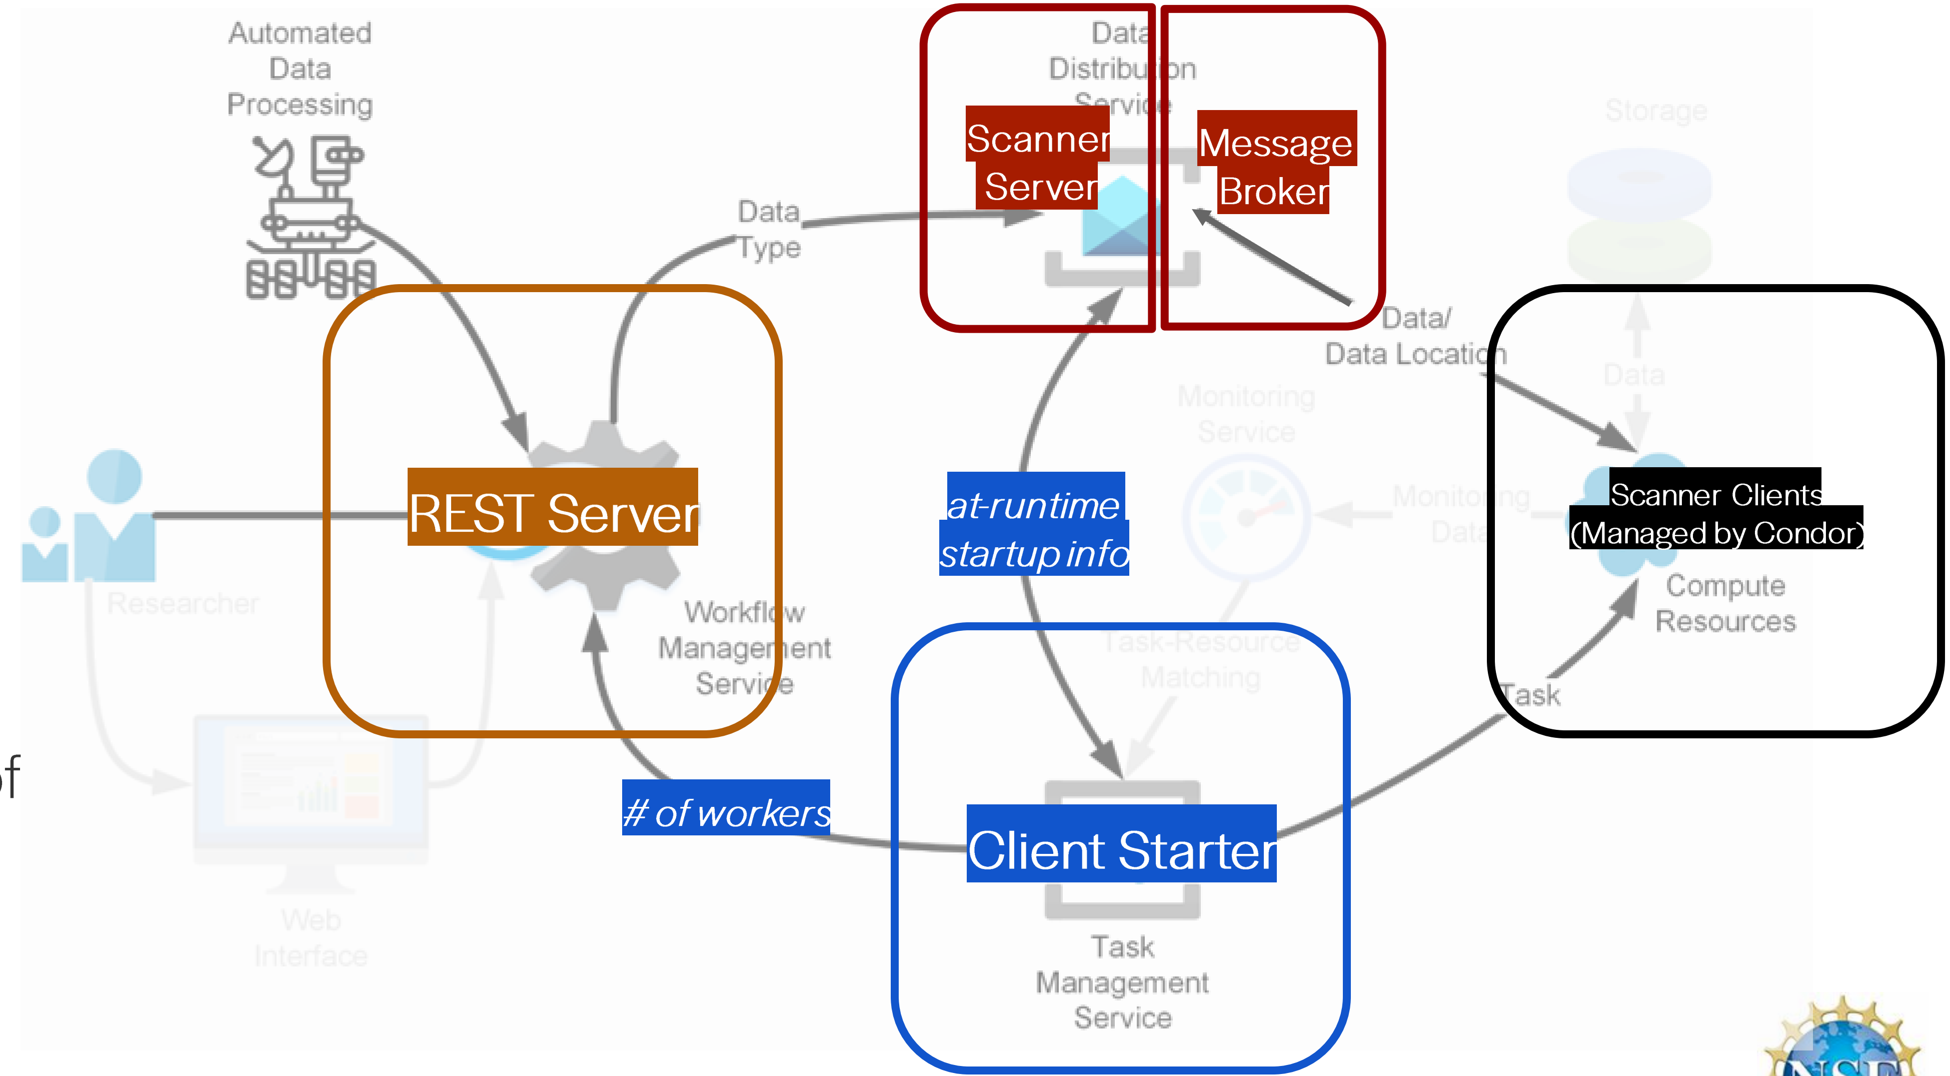

Applications – Reconstruction

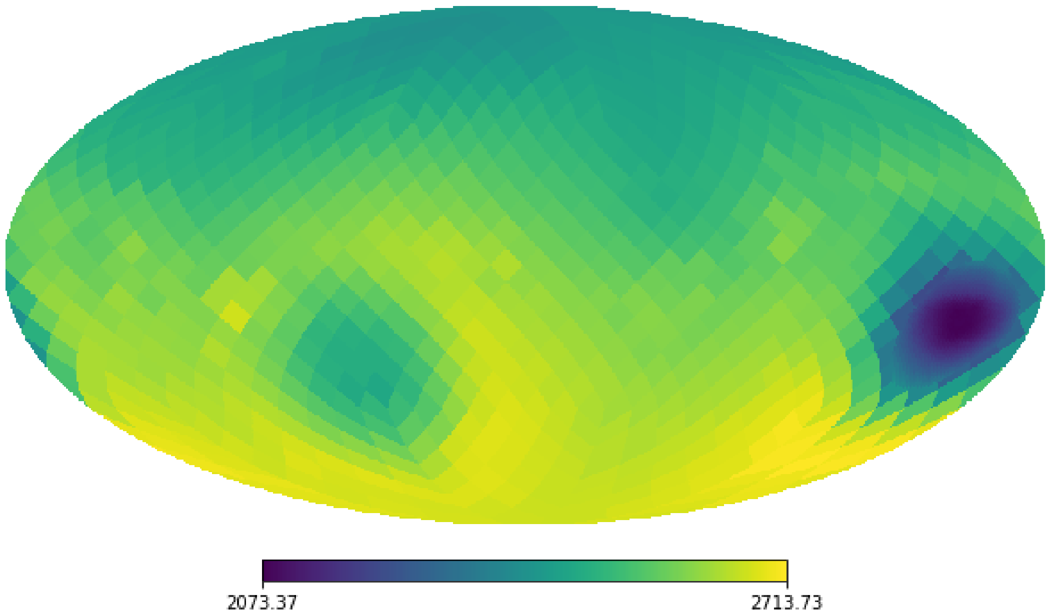

- Determine
  - Direction
  - Energy
  - Position in detector
  - Type
- State of the Art:
  - Maximum likelihood technique – Cannot support newest detector knowledge
  - CNN (>TeV) and GNN (O(10) GeV) reco
- Different type of events require different methods
  - Energy a factor – At low energies events "stops" in detector
  - Lots of physics is "known"

GPU Cloudburst Experiments

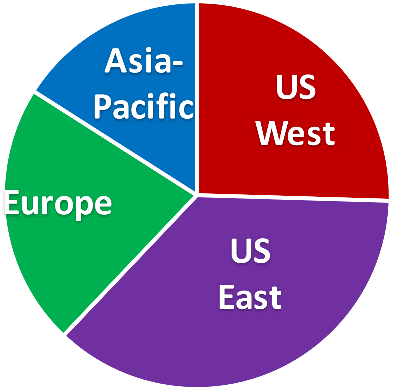

- Original Goal: Create an ExaFLOP compute pool in the cloud (80,000 NVIDIA V100) and address review panel recommendations
- Cloud provider(s) do not have those resources available – We were promised they do
  - Pre-allocated resources
  - Single cloud provider does not have those resources

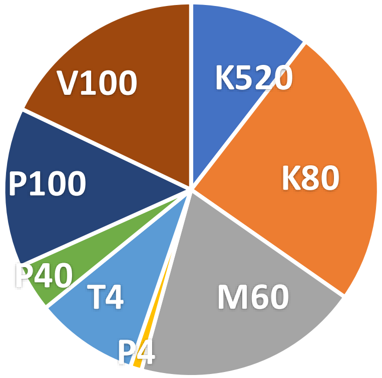

- First Experiment – On Nov 16 2019 we bought all GPU capacity that was for sale in Amazon Web Services, Microsoft Azure, and Google Cloud Platform worldwide - Creating The Largest GPU Cloud Pool in History
  - 51k NVIDIA GPUs in the Cloud
  - 380 Petaflops for 2 hours (90% of DOE's Summit, No. 1 in Top 500)
  - Distributed across, US, EU, and Asia-Pacific
  - Cost: $50-150k (under NDA)
- Second Experiment – More realistic test
  - Most cost-efficient GPUs for 8 hours
  - Achieve 1 ExaFLOP-hour of compute
  - Distributed across, US, EU, and Asia-Pacific
  - Cost: ~$60k

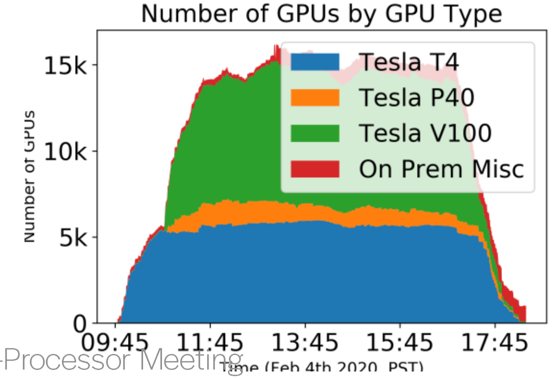

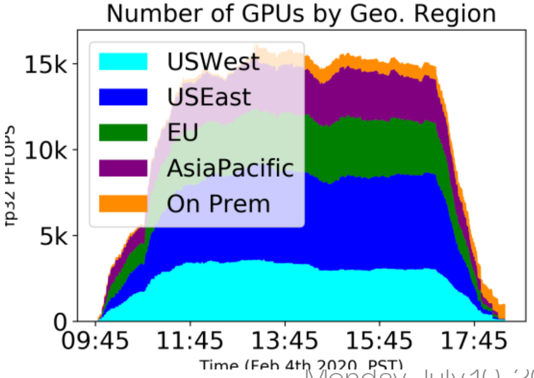

GPU Cloudburst – 1st Experiment

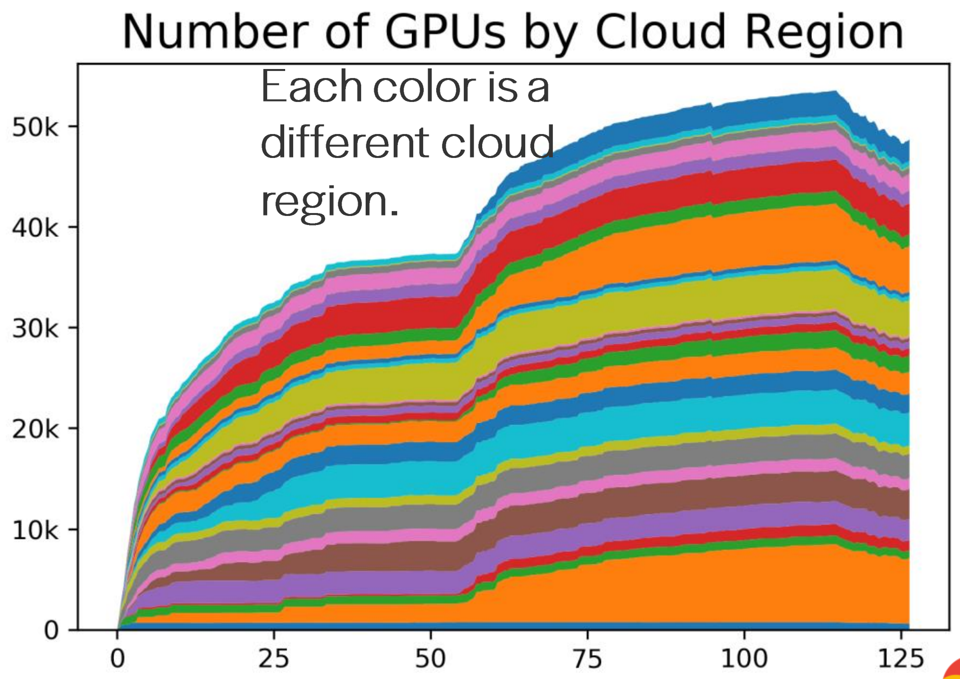

aws

Azure

Google Cloud

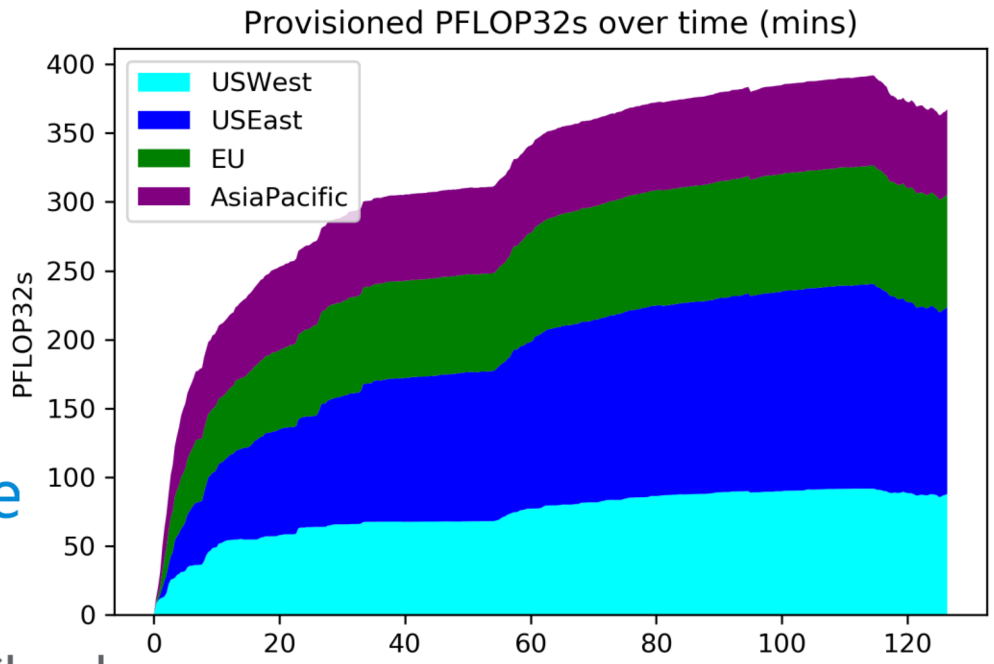

Peaked at 51,500 GPUs

Total of 28 Regions in use.

GPU Cloudburst Technology

- Multi-collector HTCondor setup – Already well-established
- Collector in each cloud region to reduce load on start-up – No idea where resources would be
- Workload is computing heavy compared to typical IceCube load – Reduce potential networking cost
- 1st Demo: In and output data stored in cloud
- 2nd Demo: Input came from UW, output stored in cloud, dedicated network links – Saturated UW SciDMZ network (100G)
- 3rd Demo: Week-long cloud usage, output and input at UW

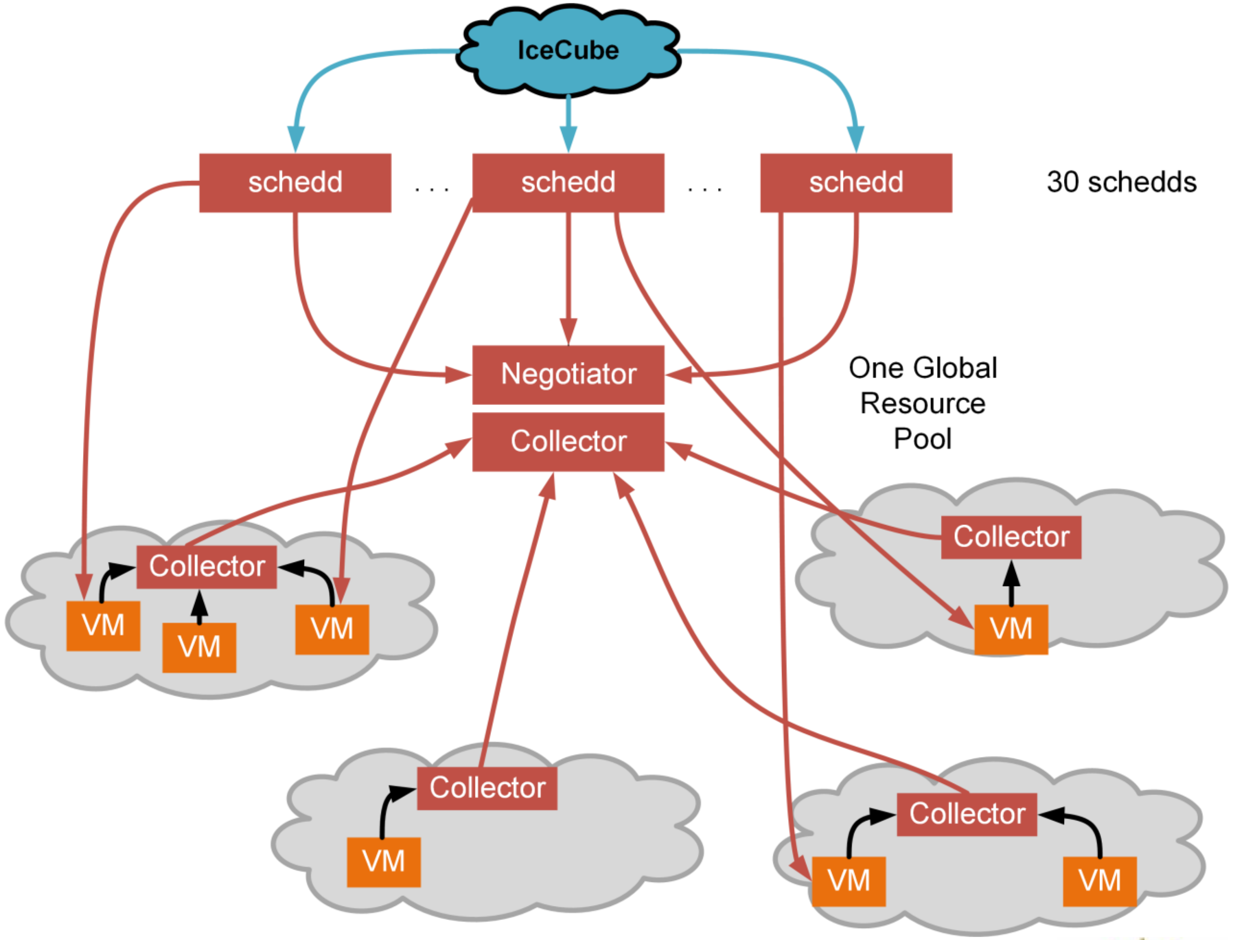

https://arxiv.org/abs/2002.06667
https://arxiv.org/abs/2004.09492
https://arxiv.org/abs/2104.06913
https://arxiv.org/abs/2107.03963

IceCube Computing – 30,000 Foot View

- Classical Particle Physics Computing
  - Ingeniously parallelizable – Grid Computing!
  - "Events" - Time period of interest
  - Number of channels varies between events
  - Ideally would compute on a per event-basis
- Several caveats
  - No direct and continuous network link to experiment
  - Extreme conditions at experiment (-40 C is warm, desert)
  - Simulations require "specialized" hardware (GPUs)
  - In-house developed and specialized software required
  - Large energy range cause scheduling difficulties – Predict resource needs, run time, etc.

IceCube Computing – 10,000 Foot View

- Global heterogeneous resources pool
- Mostly shared and opportunistic resources
- Atypical resources requirements and software stack
  - Accelerators (GPUs)
  - Broad physics reach with high uptime - Lots to simulate
  - "Analysis" software is produced in-house
    - "Standard" packages, e.g. GEANT4, don't support everything or don't exist
    - Niche dependencies, e.g. CORSIKA (air showers)
- Significant changes of requirements over the course of experiment - Accelerators, Multi-messenger Astrophysics, alerting, etc.

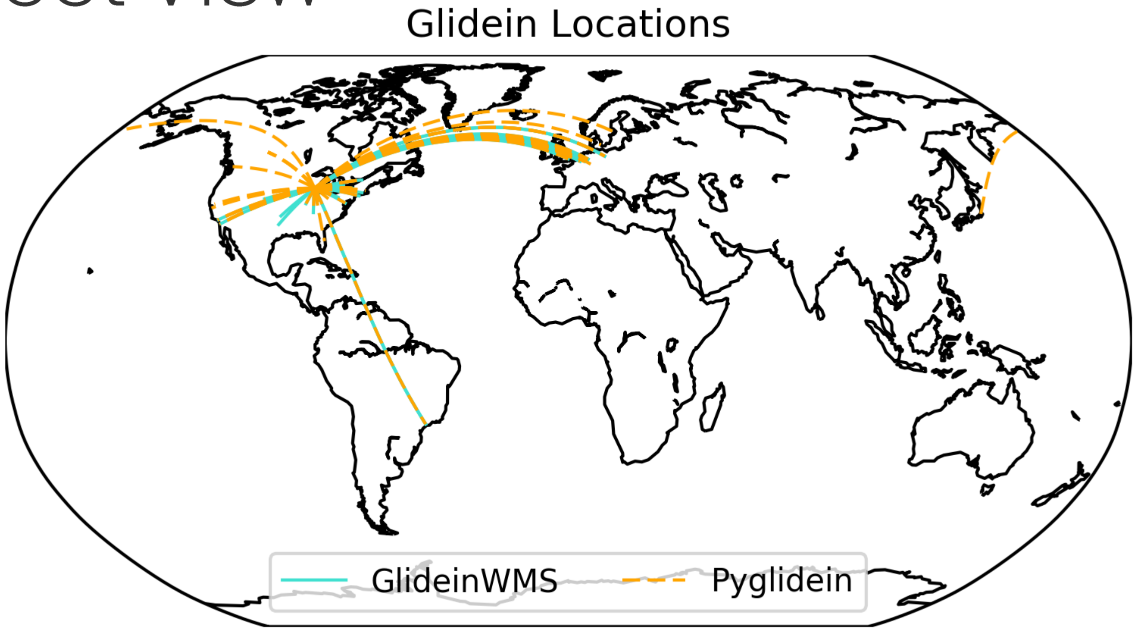

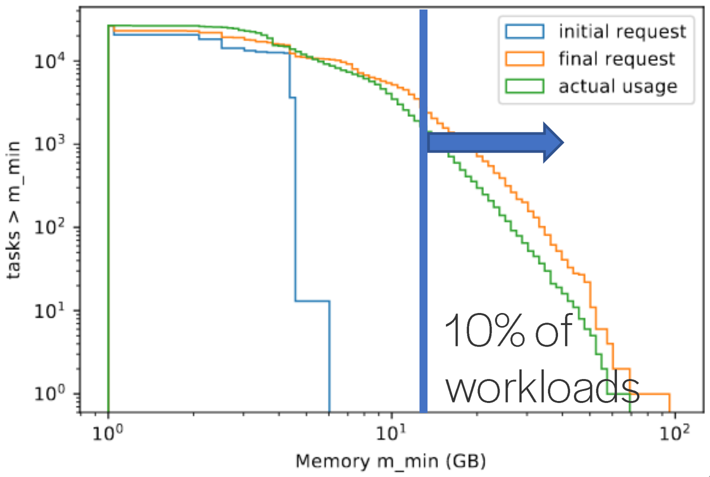

IceCube Grid – Simulation Workflow

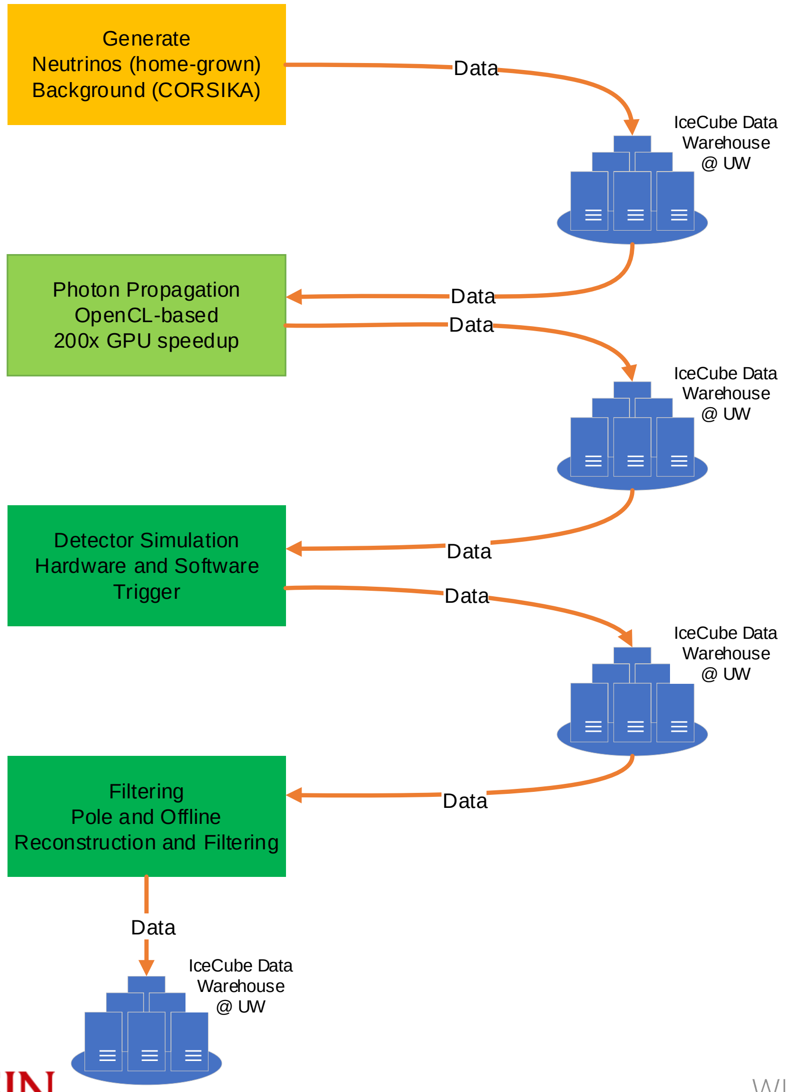

- Fairly straightforward particle physics-like workflow
- Big constraint is lack of dedicated resources
  - No data aware scheduling
  - Lots of data movement – Lots of time wasted to move data
- Different steps can have drastically different requirements

IceCube Grid – PyGlidein - I

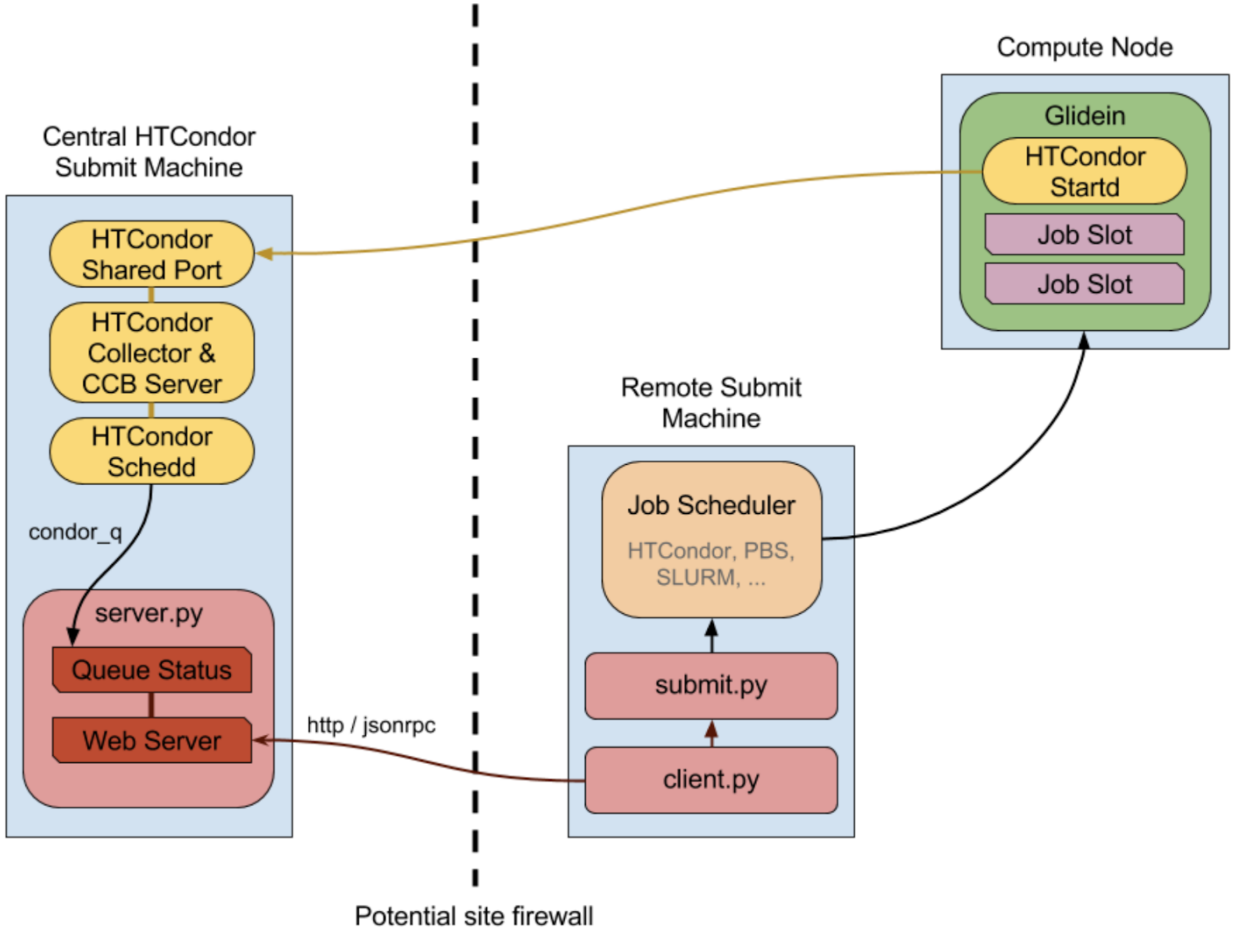

- Separate job submission from workflow management
- Lightweight design as possible
- Only difference between sites is a config file
- Why separate system?
  - Performance issues – Maximum ~3500 jobs
  - Experts needed for deployment, operation, and monitoring
  - Individual users could not use distributed resources

IceCube Grid – PyGlidein - II

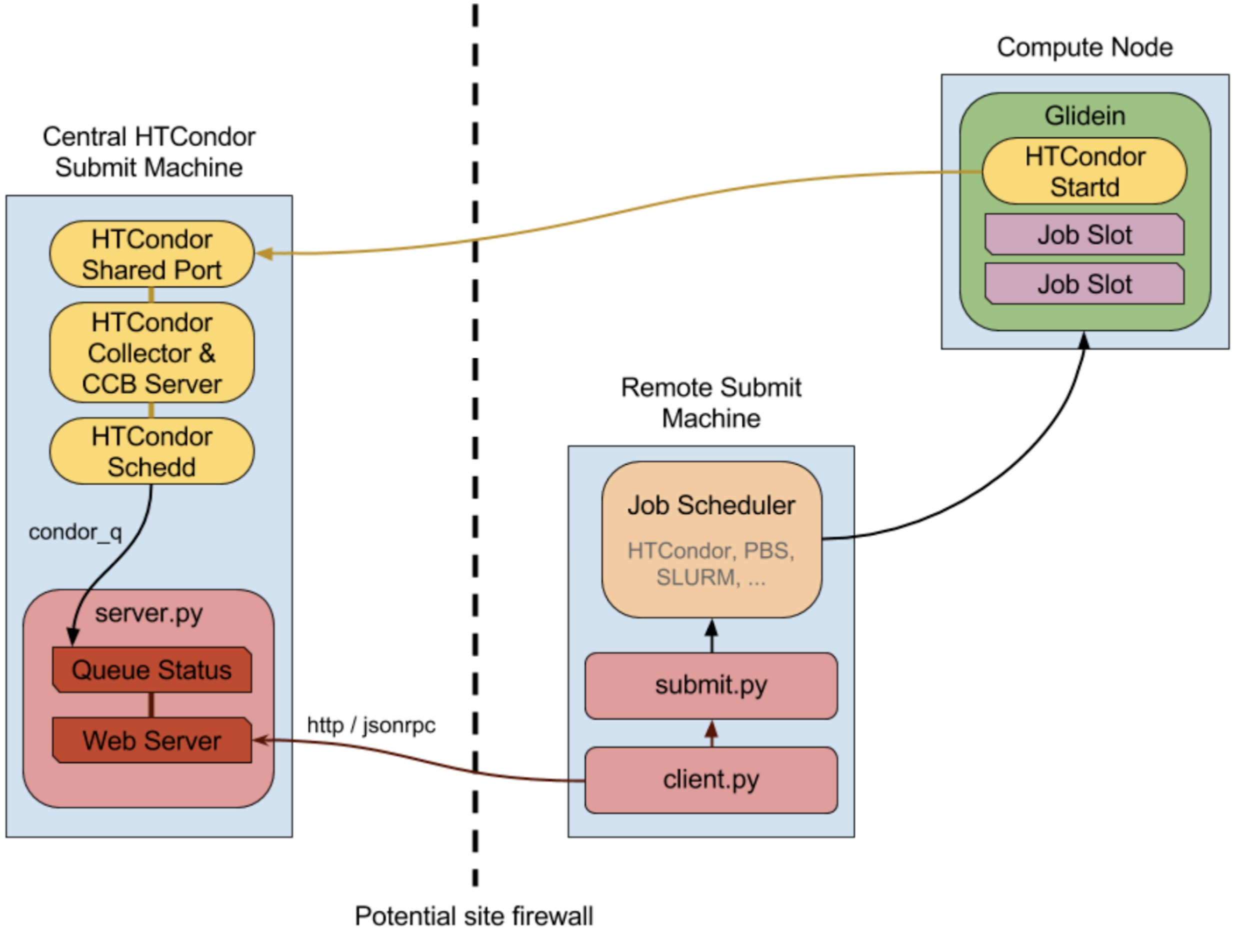

- Exposed the job resource requirements - CPU, Memory, Disk, GPU – via HTTP
- Remote client queries for job requirements and submit **HTCondor `startd`** jobs accordingly within local resource constraints
- When **HTCondor `startd`** executes connects back to central pool
- Multiple jobs are submitted per single job in pool – Assuming other jobs will be able to use slots, otherwise dies within set amount of time

IceCube Grid – PyGlidein - III

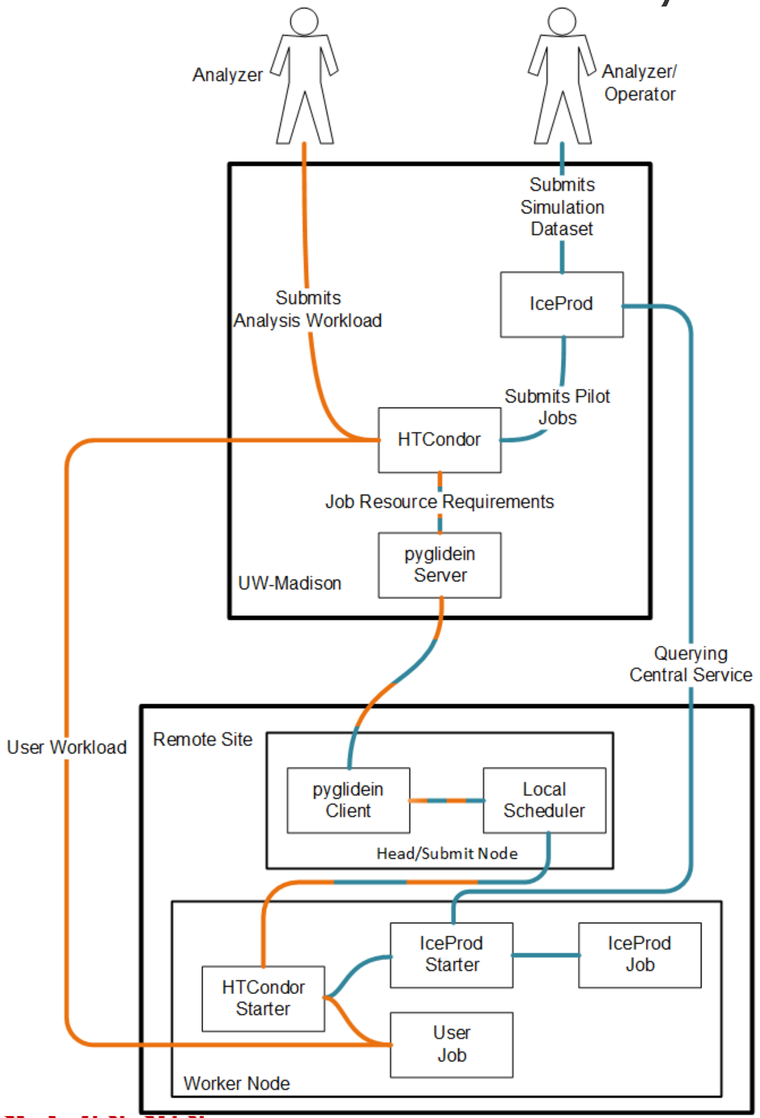

- User perspective
  - HTCondor + Data Management
  - "Just an HTCondor pool"
- Operator perspective
  - Little overhead to add cluster to pool
  - Fairly easy to monitor, e.g. `condor_status`
  - No need for a CE - Use `SSH` or `cron` for submission
  - Local container support
- Future improvements
  - Code needs clean-up – Organic growth to support multiple schedulers
  - User container support