IceCube and the Open Science Pool

Benedikt Riedel
UW-Madison

OSG AHM
17 March 2022

IceCube

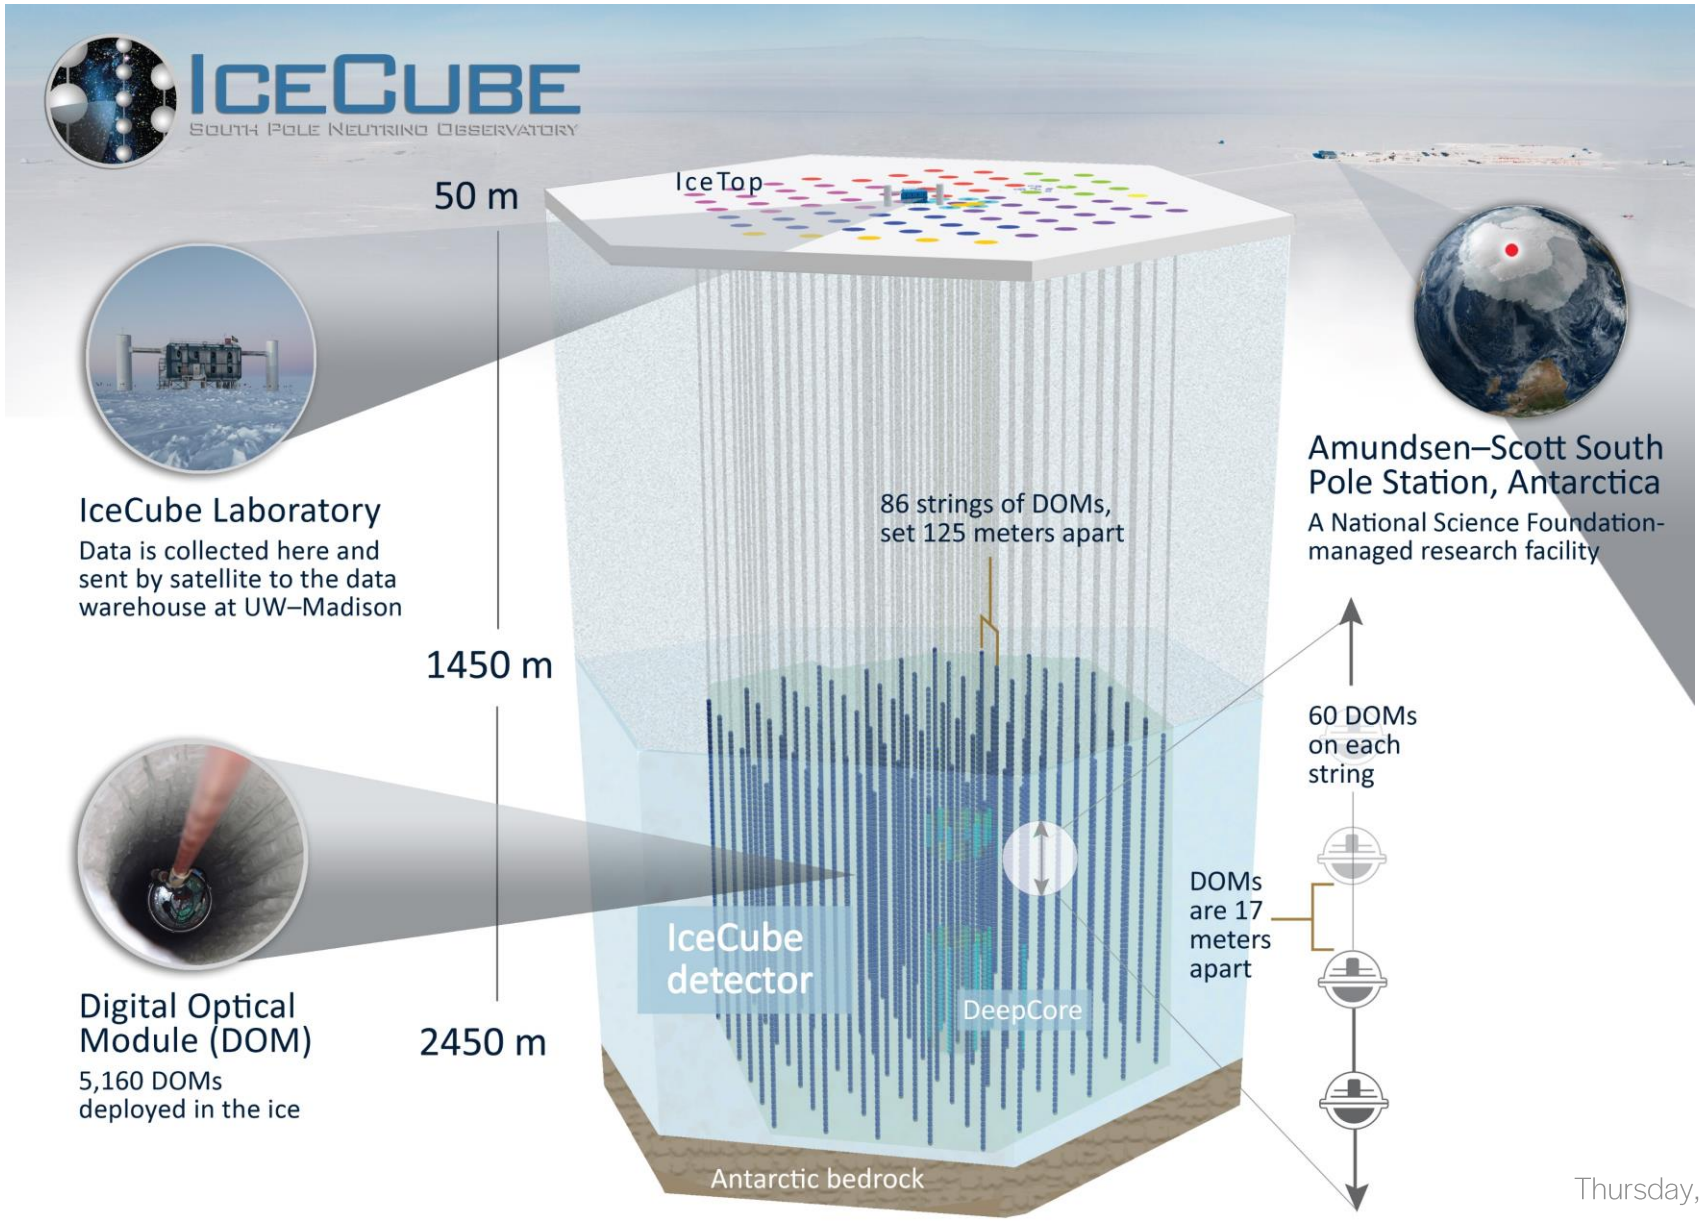
IceCube
South Pole Neutrino Observatory
50 m
IceTop
IceCube Laboratory
Data is collected here and sent by satellite to the data warehouse at UW–Madison
86 strings of DOMs, set 125 meters apart
Amundsen–Scott South Pole Station, Antarctica
A National Science Foundation-managed research facility
1450 m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

IceCube Science

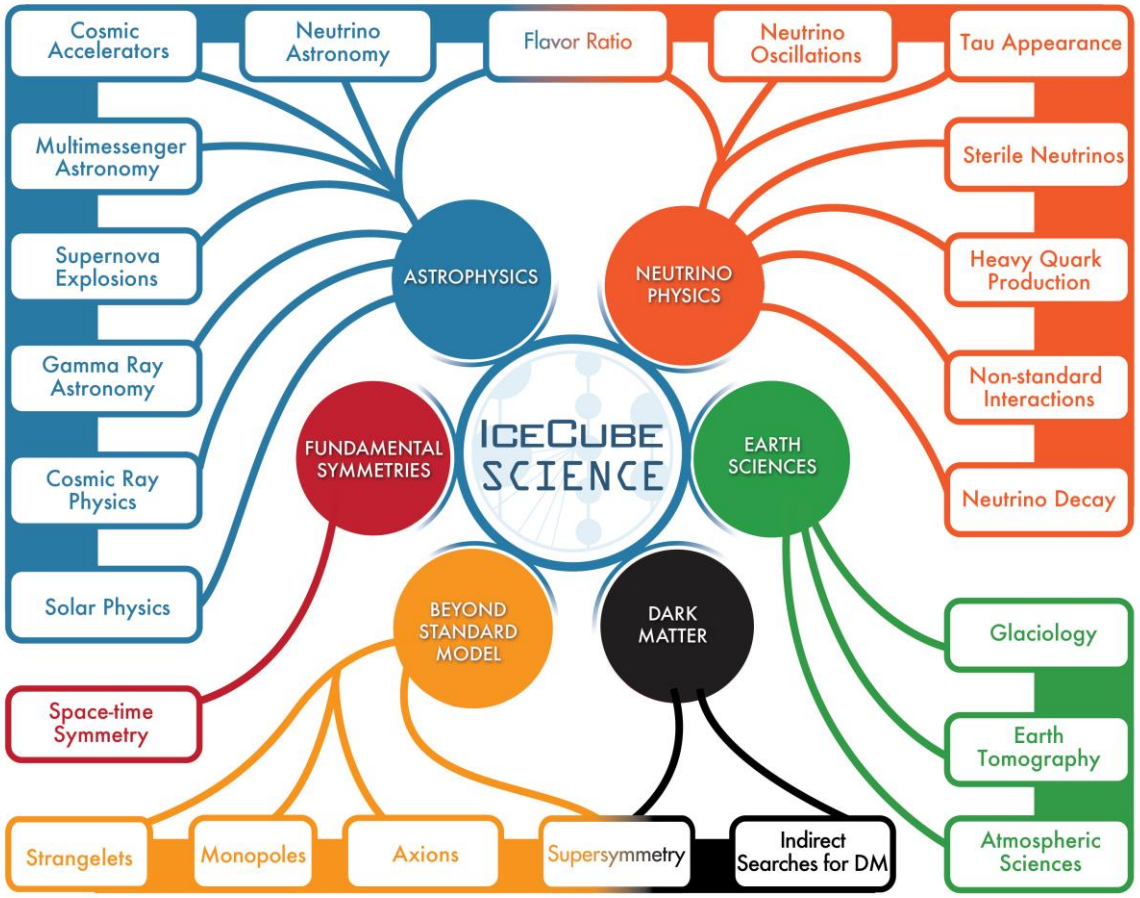

- Novel instrument in multiple fields
- Broad science abilities, e.g. astrophysics, particle physics, and earth sciences
- Lots of data that needs to be processed in different ways
- Lots of simulation that needs to be generated

IceCube Science – Why neutrinos?

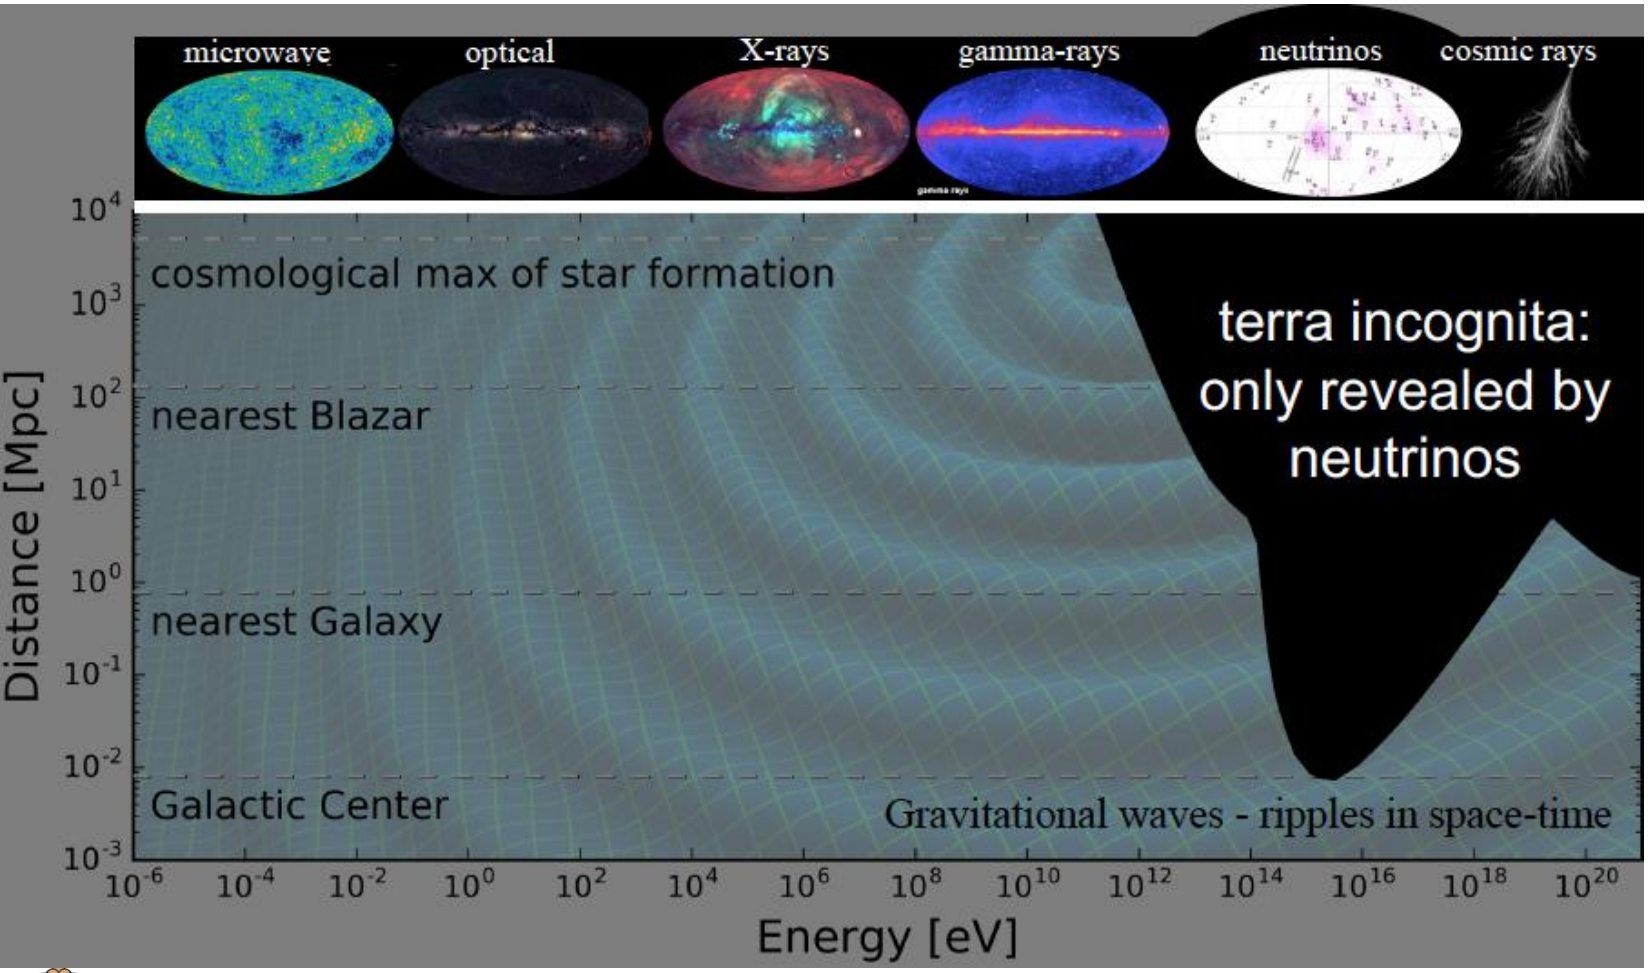

- 20% of universe is dark to "traditional" astronomy, i.e. using electromagnetic waves/light
- Need a new set of "messengers" – Gravitational Waves and Neutrinos

IceCube Science – How does it work?

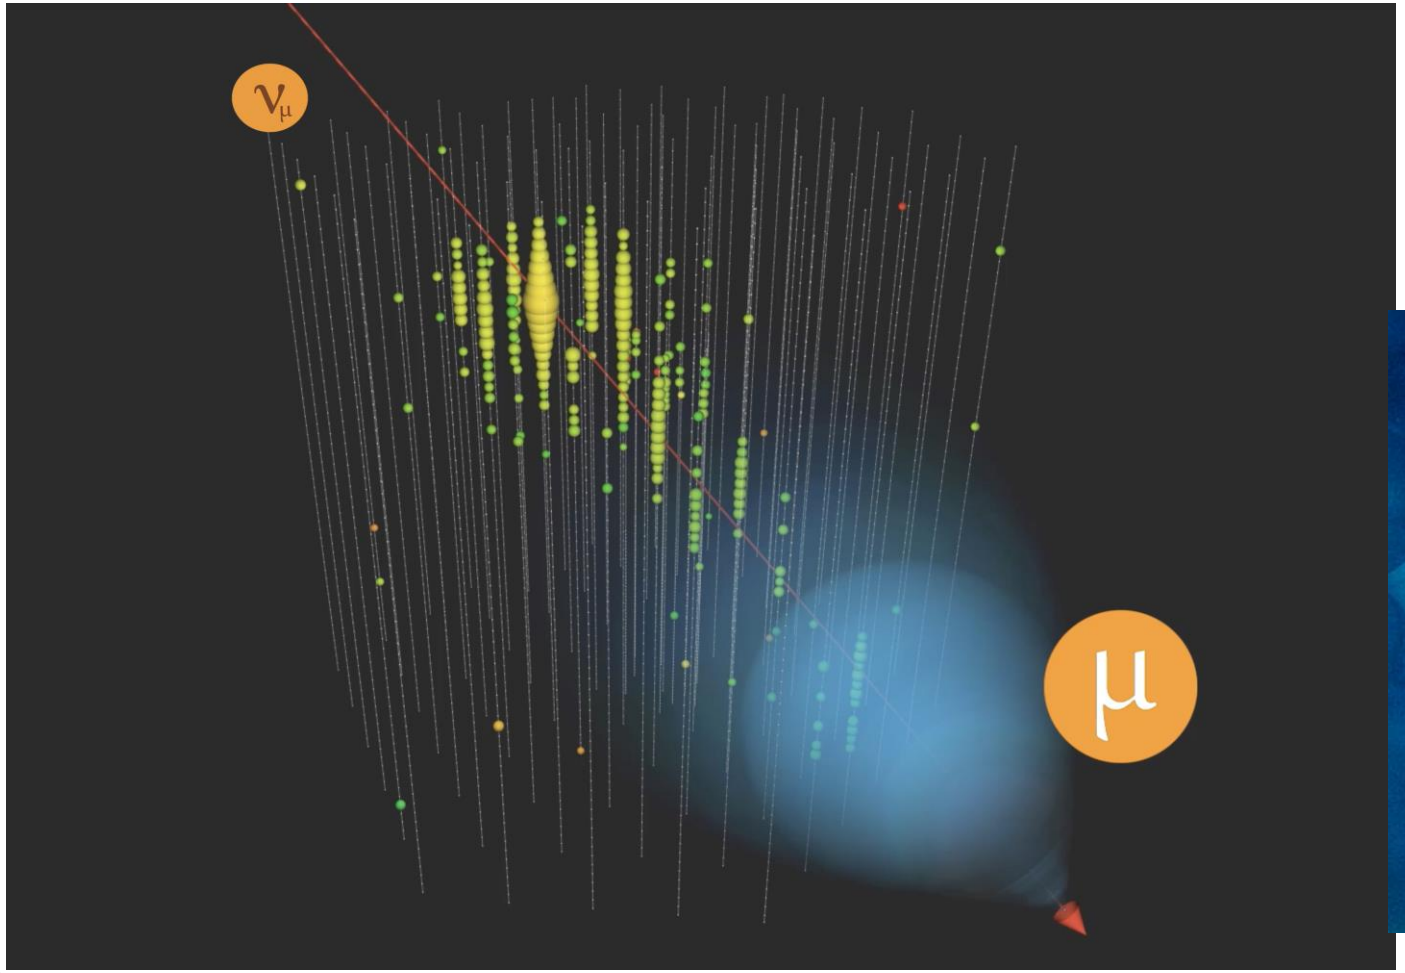

- Cherenkov light - Sonic boom with light
- Cherenkov light appears when a charged particle travels through matter faster than light can

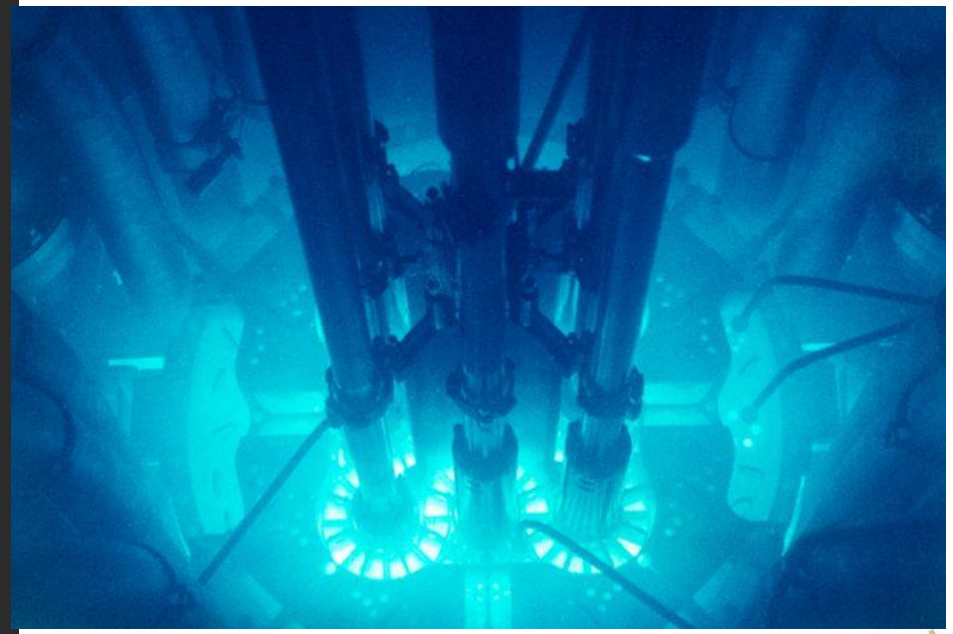

IceCube Computing – 10,000 Foot View

- 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.

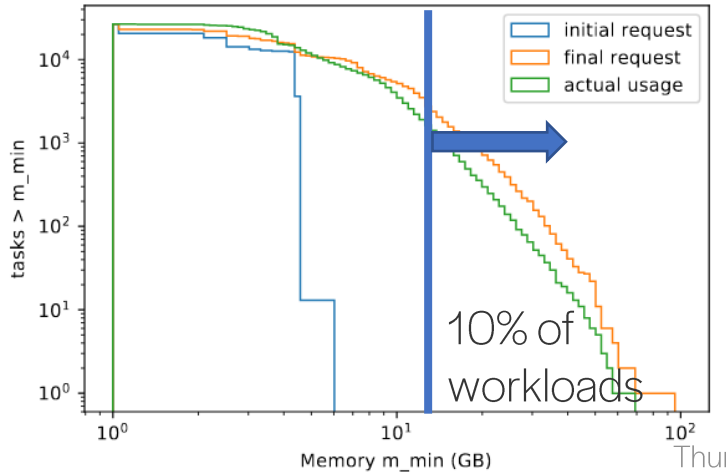

IceCube Grid – PyGlidein 2.0

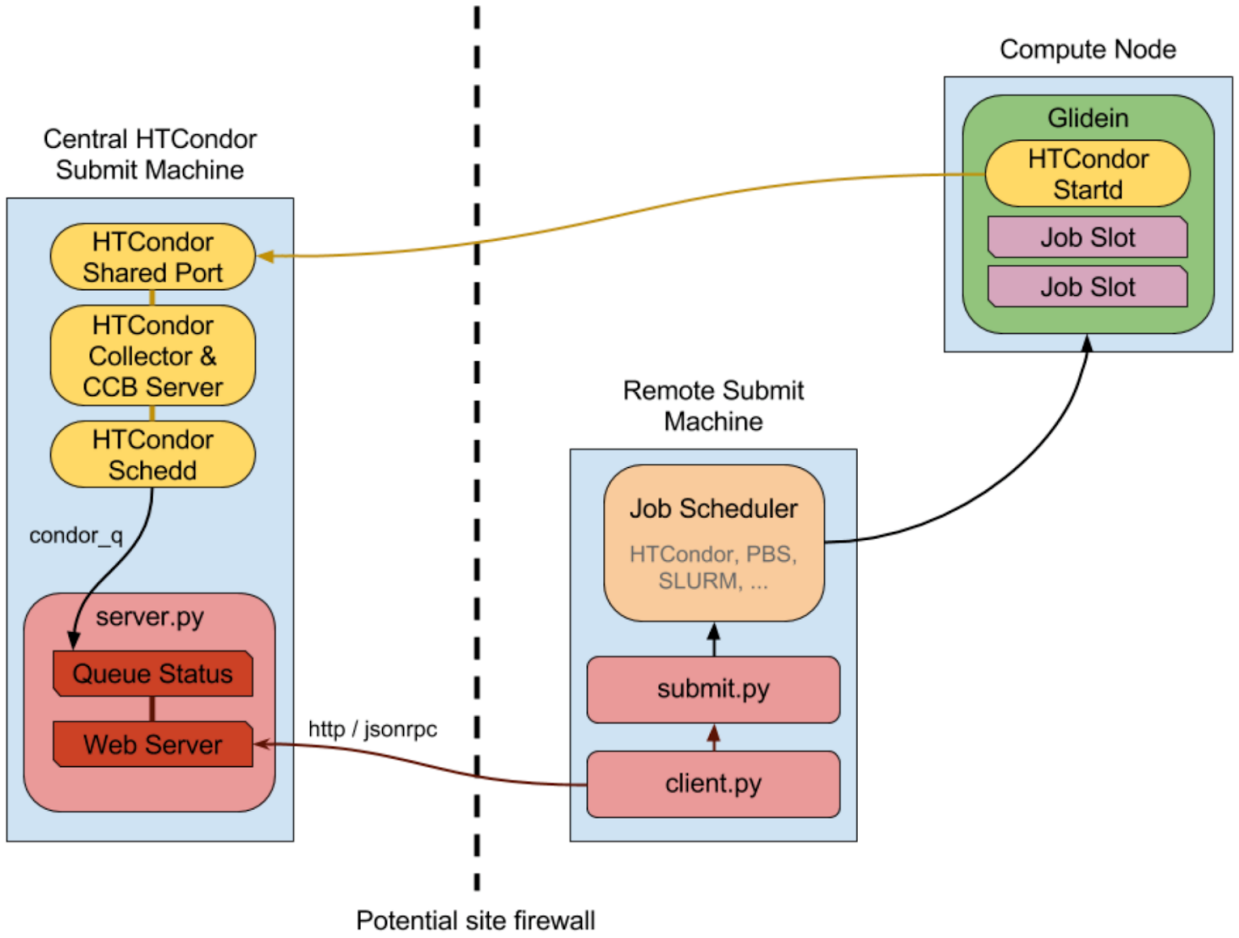

- Separate job submission from workflow management – Thin layer on top of HTCondor
- Switching to OSG-based pilot – Reduced maintenance burden
- Why not go glidewms all the way?
  - Setting up CEs still the biggest hurdle
  - CE much more heavy weight than what we have now

K8S GPU Provisioning

- Outgrowth of 2019-20 CloudBurst experiments
- GCP provided credits to test GPUs through GKE (Google K8S Engine)
- GKE supports GPU-sharing with limited availability
  - Independent of NVIDIA MIG
  - Similar to NVIDIA vGPU for VMs
- GKE & PRP IceCube GPU provisioning moved from gwms+CE to directly K8S

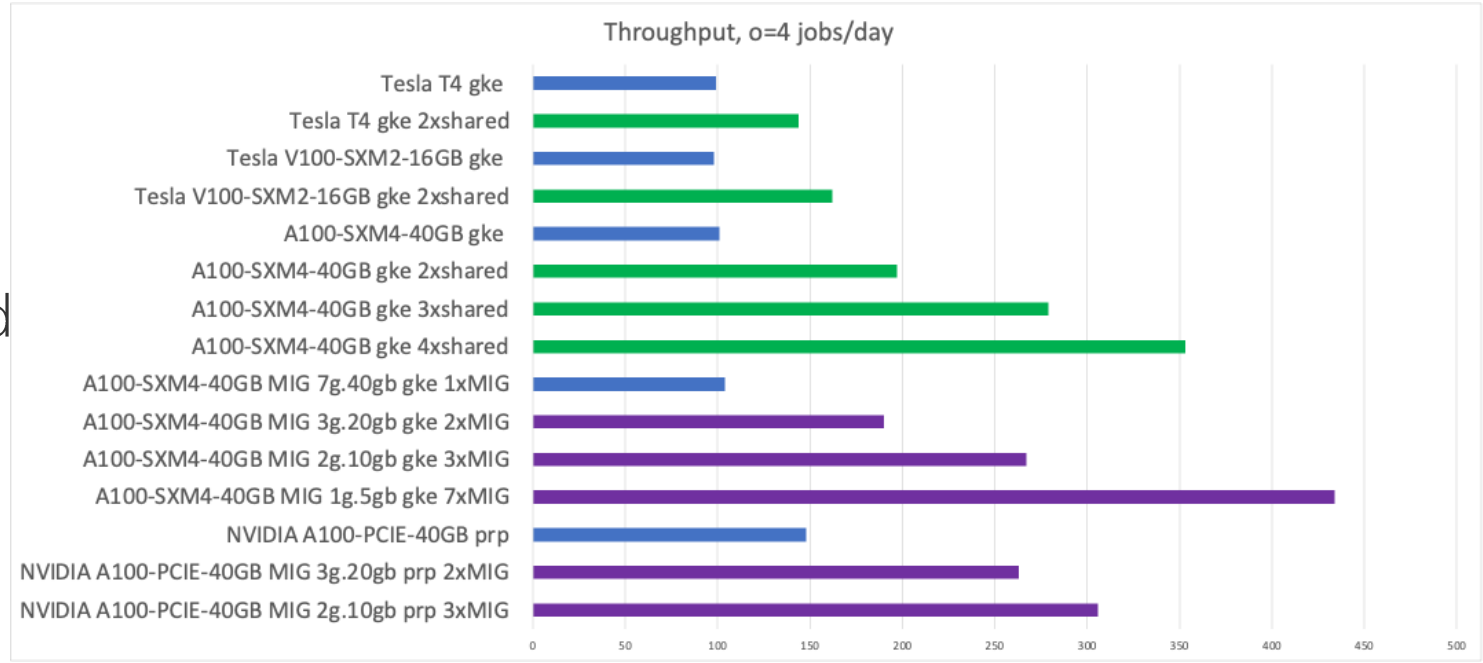

CSSI – Event Workflow Management Service

- 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

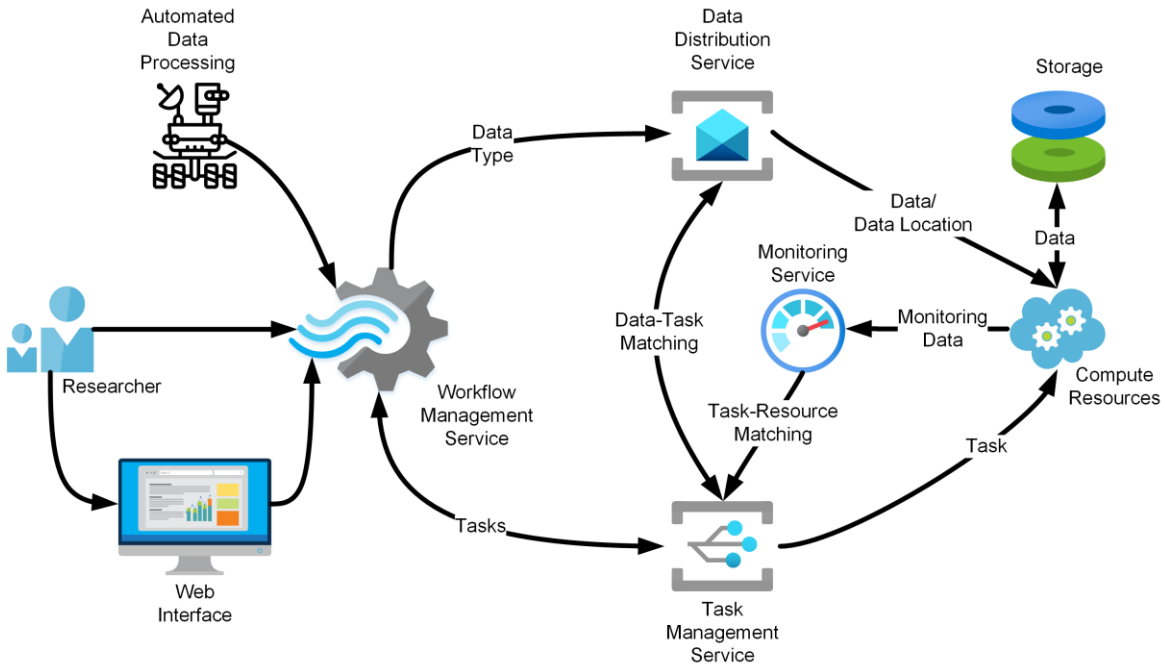

Event-based Processing In IceCube

- IceCube already uses event-based processing in two applications
  - MMA reconstruction
  - Dynamic Stack CORISKA – Feeding the photon propagation algorithm efficiently
- Great speed-ups
  - 100x for MMA reco
  - 16x for Air Shower Monte Carlo
- Both use 0MQ – 0MQ has it's issues:
  - Restricted to special environments – NPX or single nodes
  - Scaling - ~2000 clients maximum
  - Fire and forget model

IceCube Science – Multi-Messenger Astrophysics

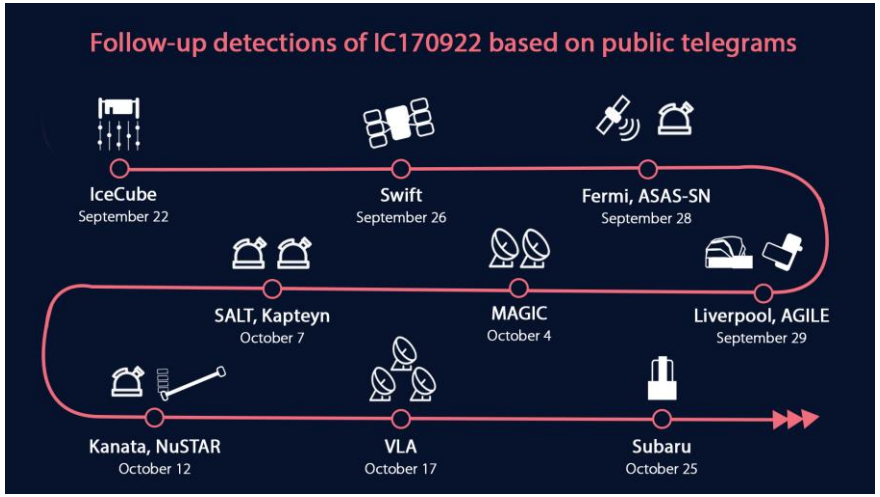

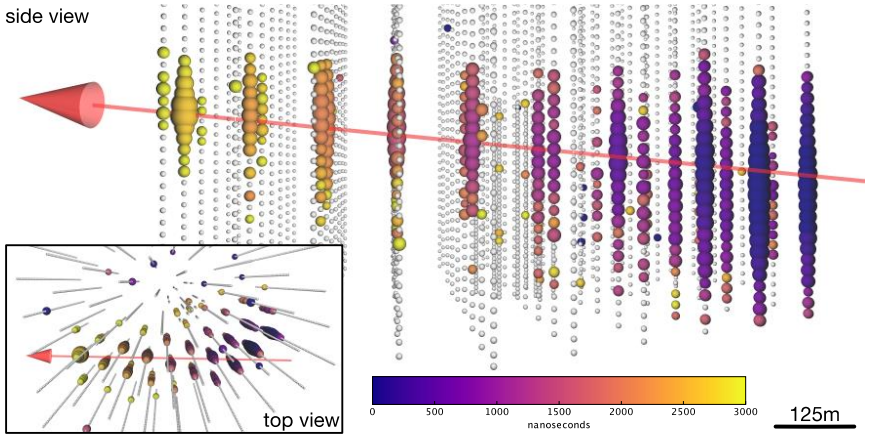

- Multi-Messenger Astrophysics (MMA)
  - Observing astrophysical phenomena with more than one "messenger" (gravitational waves, neutrinos, EM)
  - One of NSF's Big 10 Ideas
- IceCube detected an event that came from Blazar TXS 0506+056
  - Follow up observations by several observatories/telescopes showed 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 – Reconstruction

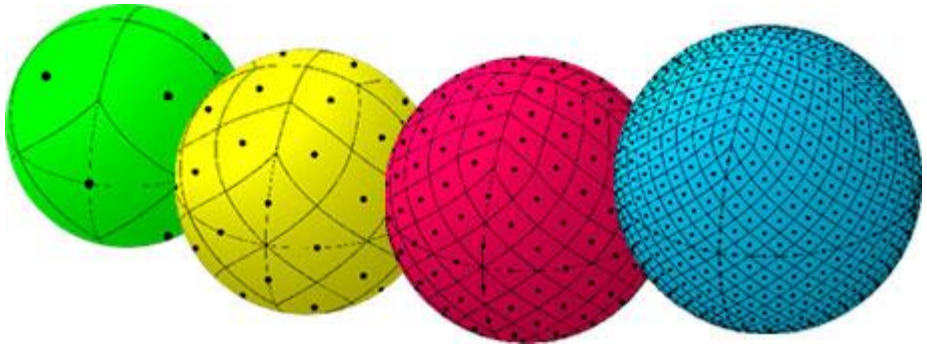

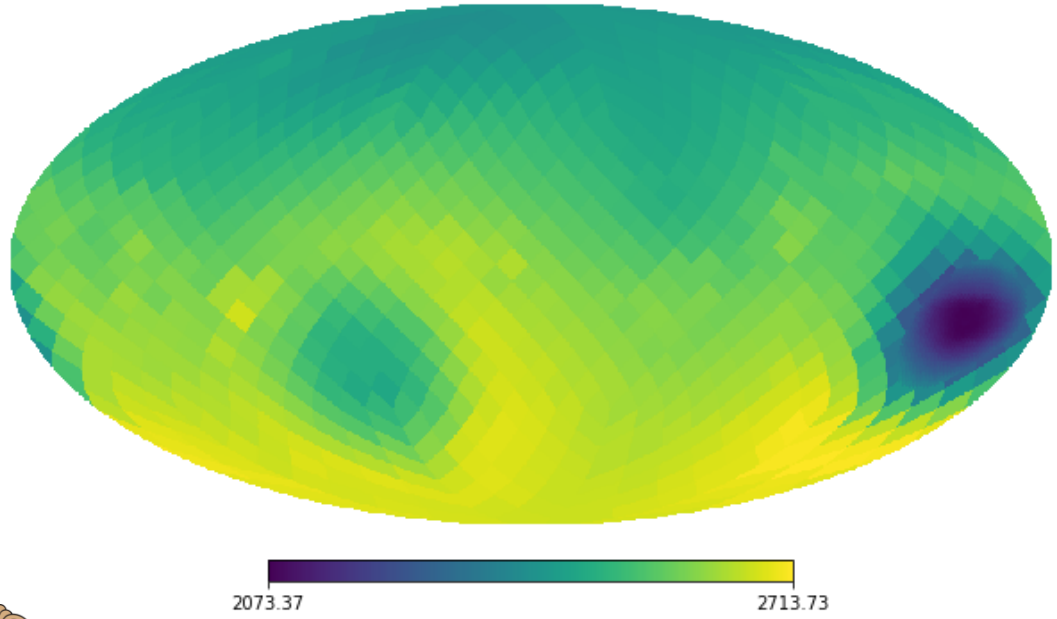

- Most accurate directional reconstruction comes by scanning across the sky
  - Split sky into constant surface area pieces
  - Test each directional hypothesis against likelihood
  - Create directional likelihood map
  - Gives most probable direction and error
- Each hypothesis calculation is independent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

Thank you!

Questions?