

# IceCube Computing in the Cloud

Benedikt Riedel  
UW-Madison

OSG All Hands  
Sept 2, 2020






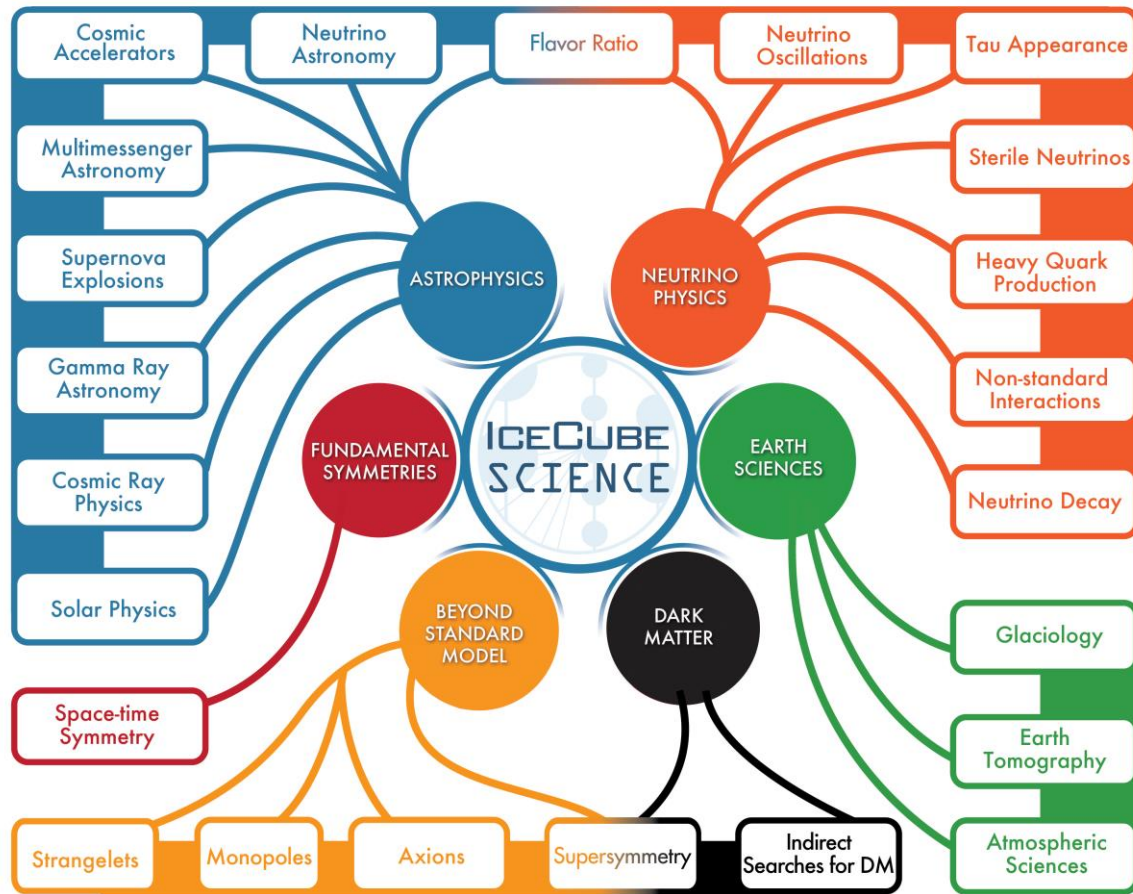
# Outline

- IceCube
- Science
- IceCube Computing
- Cloudburst Experiments

# IceCube




# 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


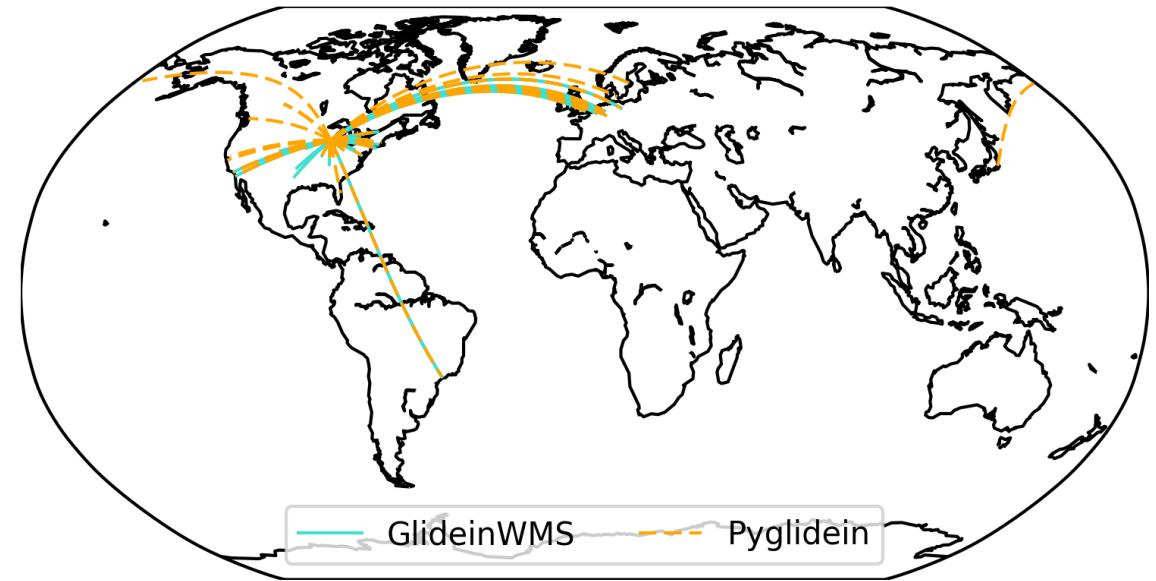
# IceCube Computing – 30,000 Foot View

- 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.


Glidein Locations



# 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









# Why GPUs?



# 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 – 1<sup>st</sup> Experiment

Number of GPUs by Cloud Region



Provisioned PFLOP32s over time (mins)




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
- 1<sup>st</sup> demo: In and output data stored in cloud
- 2<sup>nd</sup> demo: Input came from UW, output stored in cloud





# What did we learn?

- There isn't a cloud, there aren't three vendors, there are 30+ clouds – Each region is it's own cloud
- The cost can break you – IceCube estimated yearly cost is O(\$50-100M)
- Lots of resources are available without significant planning
- Social engineering - One of many customers, need someone on the inside to advocate for you
- University and funding agencies don't quite know how to deal with cloud yet
- Network can break the bank – Networking costs more than compute

Thank you!

Questions?

Acknowledgments:

The Cloud burst would not have been possible without the hard work of Steve Barnet, Vladimir Brik, Frank Würthwein, David Schultz Igor Sfiligoi, and DoIT networking team, among others I gratefully acknowledge the NSF grants OAC-1941481, MPS-1148698, OAC-1841530, OPP-1600823, OAC-190444, and OAC-1826967.

Some of the costs for the networking tests were covered by credits provided by Amazon, Microsoft and Google.