Computing Report

Benedikt Riedel
UW-Madison

IceCube Management and Operations
NSF Site Visit March 16, 2020

Outline

- Overview
- Computing Infrastructure
- Production and Physics Software
- Data Processing
- Non-M&O Funding
- Summary

Deliverables

- Data Warehouse and Storage Infrastructure for experimental, simulation, and analysis data, including data retrieval from Pole
- Computing cluster for timely offline data analysis and simulation production, including GPU computing
- Data Center Infrastructure, i.e. infrastructure to maintain data warehouse and cluster
- Provide infrastructure and support to utilize collaboration computing resources
- Offline/analysis software support and maintenance, including distributing workloads across a global computing grid

IceCube Computing

- 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.

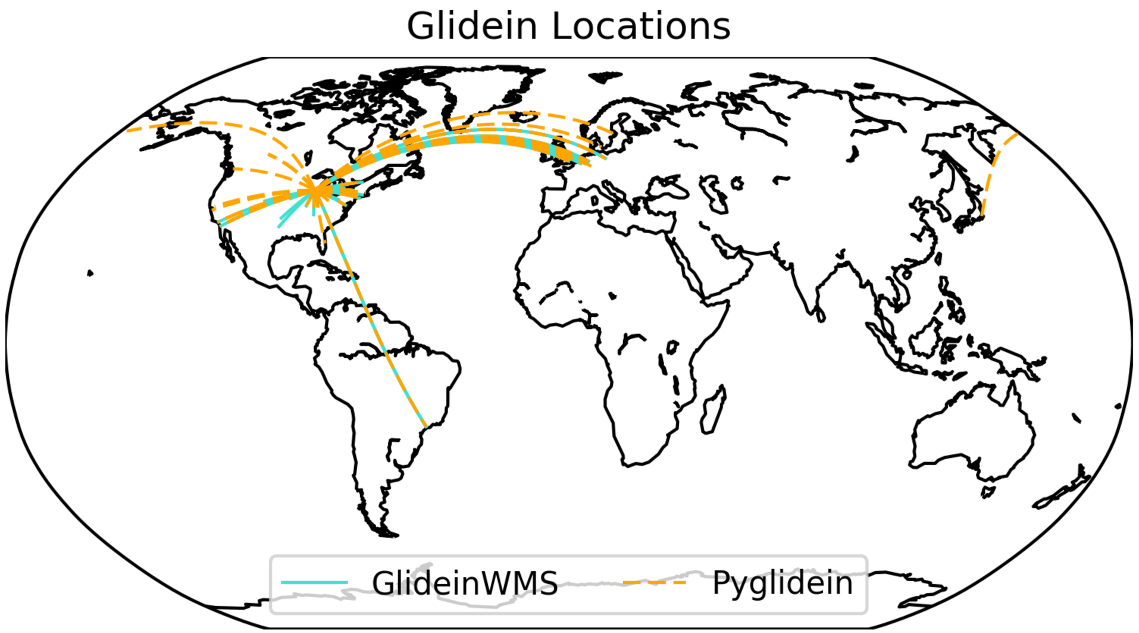
Glidein Locations

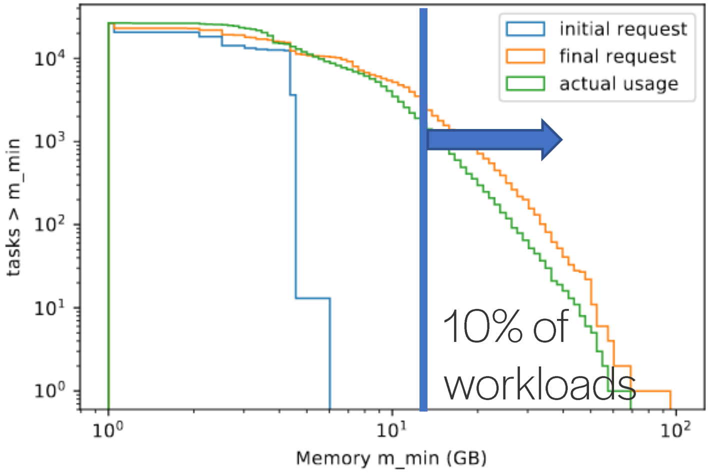

GPU Cloudburst Experiments

- Original Goal: Create an ExaFLOP compute pool in the cloud (80,000 NVIDIA V100) and address SCAP 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 Cloud Simulation 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
- 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
- More experiments planned - Half the funds are left
- Predominately funded through non-M&O awards
  - Required a total of 2 FTE for ~6 months
  - E-CAS for IceCube personnel
  - CESER for UCSD personnel
  - EAGER for cloud credits

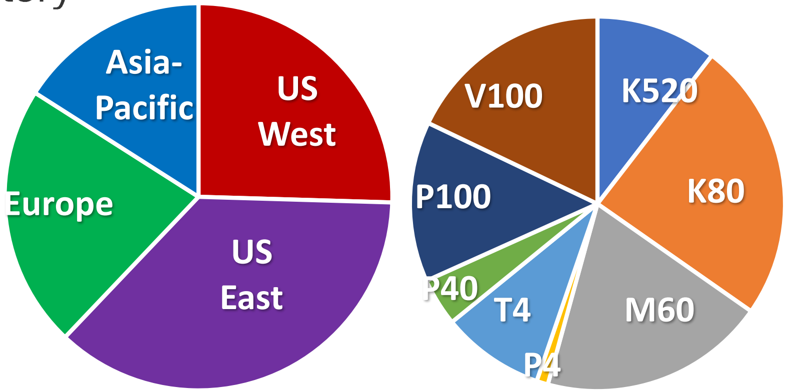

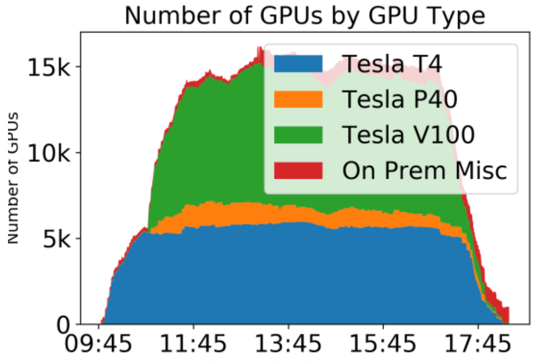

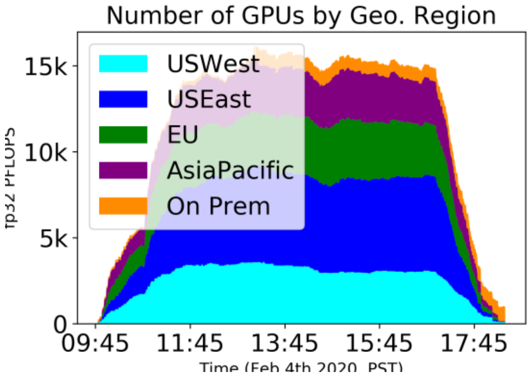

Challenges

- IceCube is in transition – Discovery to Precision Science
- Diminishing returns of more hardware
  - More hardware will not solve resource shortfall
  - More hardware will be used inefficiently
- Significant "technical debt"* to address
  - Evolving landscape for scientific computing, code, etc.
  - Data movement in simulation and data processing
  - CPU and GPU efficiency
  - Adjusting for resource pool – Larger resources for shorter time
- Machine Learning/Artificial Intelligence on the rise:
  - Specialized hardware required – Some only available in the cloud
  - More interactive computing – Iterative rather than batch computing

* Code is not up to standards, not maintained, etc.

Computing Infrastructure

Compute Infrastructure – UW

- WIPAC and UW resources are the backbone of computing infrastructure for IceCube
- WIPAC hosts the central data warehouse for IceCube detector and simulation data, and central data analysis facility
  - 10 PB of storage available (1 PB more storage already purchased, not deployed)
  - ~6000 CPU cores (90+% usage), ~300 GPUs (90+% usage) dedicated to IceCube
  - Interactive analysis and support infrastructure
- Resources are split between UW Physics Department, and OneNeck facility in Madison
  - UW Physics Department - Compute cluster and ancillary storage
  - OneNeck - Core services, storage infrastructure, etc.
  - OneNeck replaced 222

Computing Infrastructure – UW

- Storage infrastructure needs re-design
  - More storage for users
  - R&D area to study feasibility of different storage technologies: Ceph, dCache, etc.
- GPU capabilities stagnating – Will decrease soon
  - Demand outpacing supply of GPUs compute – Machine learning, new reconstructions, etc. require GPUs, e.g. student estimated needing **half the GPU resources of the entire collaboration for a single analysis**
  - Largest expansion of GPU/accelerator resources through applying to outside resources, e.g. NSF's XSEDE program – Not guaranteed (recently saw a decrease in XSEDE allocation)
- CPU resources are showing their age – Last significant refresh 7 years ago

Computing Infrastructure – Collaboration

- Computing pledge system has had
  - Positive effects:
    - Investment from international partners
    - World-wide distributed resource pool
    - Additional GPU resources from collaboration
  - Negative effects:
    - **In-kind contributions require significant management overhead**
    - **Cost per resource significant**
    - Inefficiently resource utilization – More on this later
- Established **long-term archive** at NERSC for IceCube raw data
- Working on substituting resources: People rather than CPU/GPU

Computing Infrastructure – National

- Significant invest in GPU resources on national-scale HPC resources
- USA
  - **Open Science Grid** (OSG) infrastructure and resources are essential
  - NSF high performance compute program – XSEDE (Comet, Bridges), leadership-class
  - DOE resources an option – Question how to access and motivate access
- EU
  - Significant number of possible resources targets, e.g. LHC facilities, supercomputers, etc. - Some come with significant restrictions
  - Lots of issues boil down to “Who will pay for it?” or “Who paid for it?” – Not used to sharing
  - PRACE (EU-version of XSEDE) support requested – No results yet
- Race to Exa-Scale
  - Heterogenous and “new” architectures become the norm
  - Unclear how to address

Physics Software

Physics Software – Crossroads

- Physics Software is at a crossroads
    - Lack of Support – Bus Factor* ≤ 1
        - Some essential packages have **limited to no official support**
        - Some software has grown **too complex to be maintained by in-kind contributions** (read: grad student, post doc, faculty)
        - **Improvements take years to implement**
    - "Herding cats" - Death by fractional FTE (0.1 + 0.15 + 0.1 + … = 0)
        - **In-kind model is not sustainable** – 100+ packages of "physicist" code to maintain
        - Lack of sense of responsibility
            - Unsupported software with unknown impact being used
            - "IceCube Software: where asking the absolute bare minimum is asking too much."
        - Participation is lacking – Spending significant resources trying to get people involved
        - Major releases take 1.5 – 2 years of development
    - "Bit rot"
        - No support to explore improvements
        - Large fraction of "donated" resources – No control over base software stack, not frozen in
        - Hardware not used efficient – More hardware just a waste
- **To reduce scope or not reduce scope? – Reduce scope is current answer**
- **What has led us here?**
    - Number of software interested students and postdocs leaving, graduating, becoming faculty, etc.
    - Software complexity increasing – Hard to make meaningful contributions without significant investment

* What happens to a project or package if maintainer gets hit by a bus

Physics Software – Technical Debt/Complexity

- Dynamic Stack CORSIKA
  - Example of lack of support, and complexity of software
  - Initial simple settings show factor of 2 reduction in CPU across all energy ranges
  - > 1.5 years of development – Still not production
- Photon Propagation/Ray Tracing
  - Current GPU interface hard to maintain or vendor lock-in – Newer interfaces available from exa-scale efforts
  - Newest generation GPUs have “ray-tracing” capabilities
    - Could be beneficial
    - Compatible resources available – No support within collaboration

Physics Software – Upgrade

- Upgrade physics software has **many open issues**
  - Lack of support for new modules in software – Today **can not simulate IceCube + Upgrade in full fidelity**
  - Upgrade integration will need to address some fundamental assumptions in the software – Shape, dimension, properties of module
  - New reconstruction algorithms needed for Upgrade
  - “Analysis software” is more resource intense than IceCube – Needs to be addressed given resource pool
- In-kind contributions can **only cover small fraction of effort**
  - Experts needed for photon propagation, reconstruction, etc.
  - No clear timeline given in-kind-only support – **May have detector without software to analyze data**
  - Quality of software will suffer – Physics results will suffer as well – Reproducibility, long-term viability, etc.

Production Software

Production Software – Keeping afloat-ish

- Non-production users using production systems
  - Good in moving people to using distributed resources
  - Requires more support than currently allocated
- Essential pieces of software have no to bare minimum support – Inefficient resources utilization (50+%) due to lack of support
- Maintenance and operations the focus
  - Making choices what “fires” to attend to
  - Not moving forward in terms of efficiency, new technologies, etc.
- Upgrade production will exasperate issues
- Addressing SCAP recommendations very slowly

Production Software – Cloud

- Cloud have impactful, but limited uses:
  - Off-loading compute-intensive, short-term tasks – Real-time reconstruction
  - Specialized machine learning hardware – Either commercially or to test capability
- Scaling exercises
  - Is the infrastructure ready for X% scaling – SCAP 2018 recommended 10x scaling in resources
  - Understanding shortcomings in infrastructure – Networking to Europe, etc.
- Too expensive for most services

Production Software – Current Projects

- Moving cloud real-time scan into production
- Single-Sign On and new user management – SCAP recommendation
- Transition to G Suite (Google) – Email, documents, calendars, etc.
- Maintenance and operations of current infrastructure

Data Processing

Data Processing – Status

- Transition to new workflow management system on-going - > 1 year development, goal is Run Start 2020 (April/May 2020)
- Near real-time data processing continues
  - Bugs in processing have led to repeated reprocessing campaigns – Limited in compute effort, significant human effort
  - Significant delays possible
- Pass 3 postponed indefinitely

Non M&O Funding

Non M&O Funding

- Awarded
  - CESER Award – Funds Long-Term Archive development
  - Exploring Cloud for the Acceleration of Science Phase 1 – Cloud funds
  - EAGER for Exa-Scale Demo – Awarded through UCSD, only cloud credits
- Not awarded
  - Convergence Accelerator – Data portal combining neutrino events and gamma-ray catalog
  - Humans Advancing Research in the Cloud
  - Mid-Scale R1 – GPU cluster hosted at UCSD
- Submitted
  - 3 AI Institutes @ UW, 1 AI Institute @ MSU
  - CC* with UW HEP group
- Planned
  - 3 proposals planned for later this year
    - "Event Management Service" – EvMS
    - Data Portal
    - Upgrade Software development
  - MRI with UW IT

Summary

- IceCube Computing is at a crossroads – Reduction in services and/or service quality imminent
- Unable to keep up with recent developments – Significant technical debt
- Computing landscape is changing at an accelerating pace – Falling behind quickly
- In-kind model is not sustainable in maintaining software essential to IceCube
- External funding has only slowed the accumulation technical debt

Thank you!

Questions?

BACKUP

Data Processing

- Pole Filtered Data arrives via satellite - Arrives at UW-Madison and is processed further
- Raw data is written to archival disk at pole, retrieved once a year
- Raw data is archived at National Energy Research Scientific Computing Center (NERSC)
- Filtered data is archived at Deutsches Elektronen-Synchrotron (DESY)

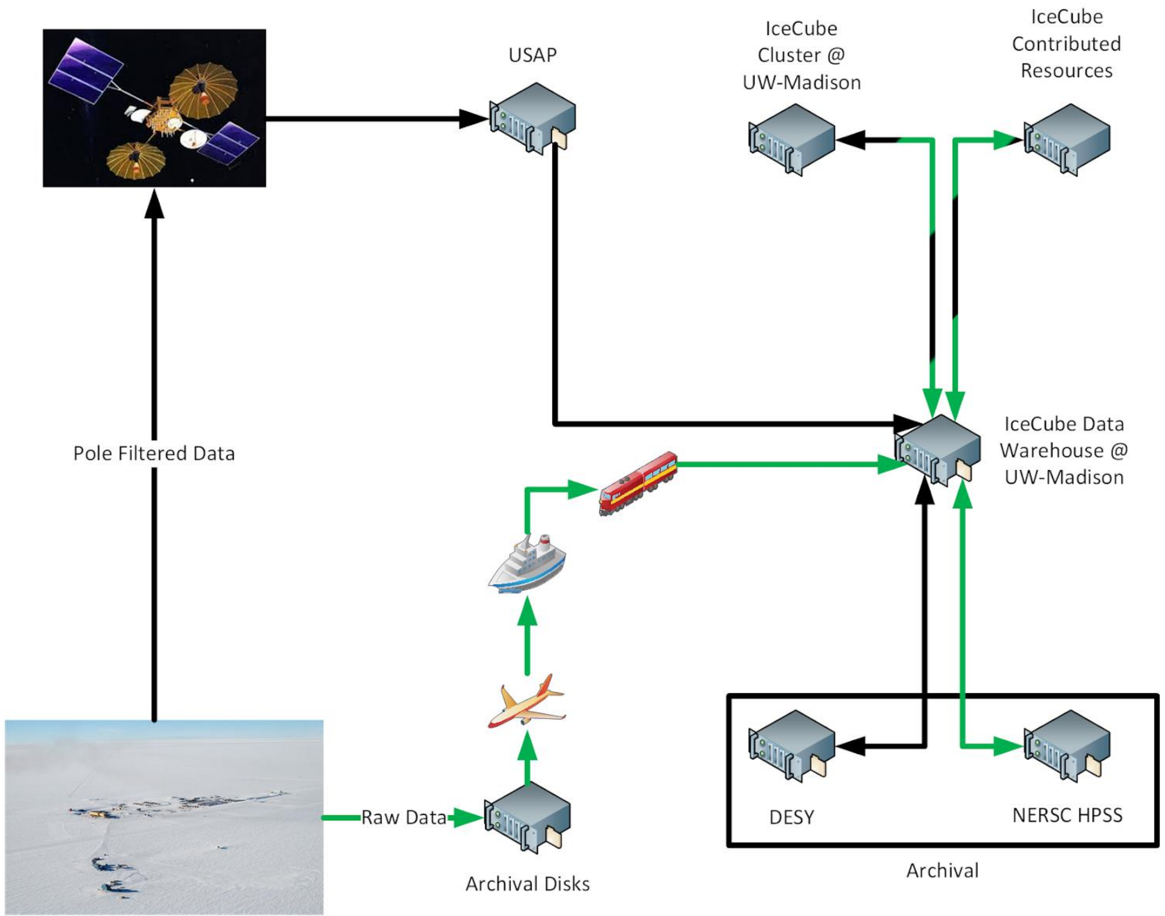

Simulation Chain

- 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

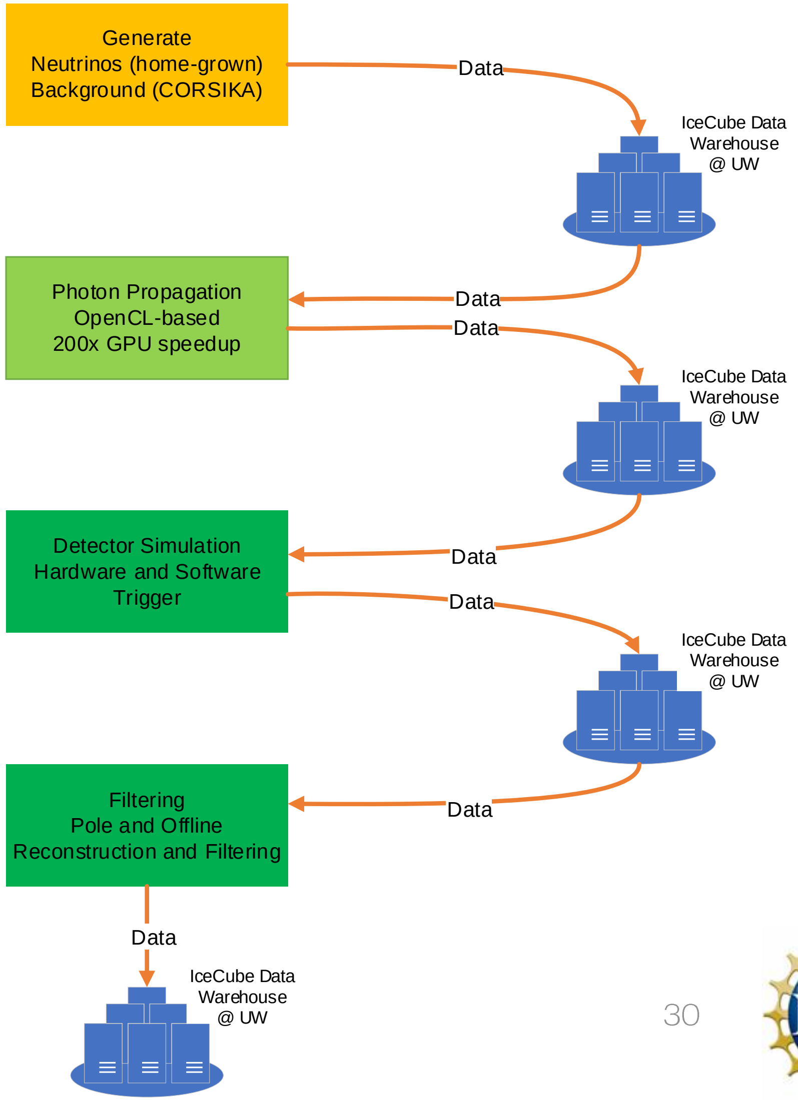

Generate
Neutrinos (home-grown)
Background (CORSIKA)

Photon Propagation
OpenCL-based
200x GPU speedup

Detector Simulation
Hardware and Software
Trigger

Filtering
Pole and Offline
Reconstruction and Filtering