The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes[1]
December 4, 2019, 12-2 pm ET
NCO, 490 L'Enfant Plaza, Ste. 8001
Washington, D.C. 20024

Participants (*In-Person Participants)

Lisa Arafune (CASC)	David Martin (ANL)
Richard Carlson (DOE/SC	Matt Selmici (UW Madison)
Dhruva Chakravorty (TAMU)	Nathan Tallent (PNNL)
Vipin Chaudhary (NSF)	Kevin Thompson (NSF)
Kaushik De (UTA)	Kamie Roberts (NCO)
Sharon Broude Geva (UMich)	Birali Runesha (UChicago)
Joyce Lee (NCO)*	Alan Sill (TTU)
Brian Lyles (ORNL)	Sean Wilkinson (ORNL)
Brian Lin (UW Madison)	Frank Würthwein (OSG, UCSD/SDSC)

<u>**Proceedings**</u>
This meeting was chaired by Richard Carlson (DOE/SC). October 2019 meeting minutes were approved.

<u>**Guest Speaker:**</u> **Frank Würthwein, Open Science Grid (OSG) Executive Director, UCSD/SDSC, *Running a GPU burst for Multi-Messenger Astrophysics with IceCube across all available GPUs in the Cloud***
(New information about cost and performance added to SC19 presentation)

<u>**Largest Cloud Simulation in History**</u> - NVIDIA (Slide 2)
380 Petaflops for 2 hours – only accomplished by working across 3 providers (AWS, Microsoft Azure and Google Cloud Platform).

- Worked with them since August 2019
- Biggest challenge: 3 providers allowing the purchase of their entire inventory
- Will discuss science motivation and the "how", cost, performance

How did we get here?
Annual IceCube GPU use via OSG (Slide 4, diagram)

- Longterm partnership with IceCube, OSG, HTCondor provided technical capabilities; (NSF-funded)
- Biggest challenge: Convince vendors and do everything needed for first time; set up accounts, etc
- Did everything with production software
- Normal production for last year IceCube via OSG (Slide 4 graph) – peak typically at ~3k GPUs/daily

[1] Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.

- Pie chart: distribution of contribution of various clusters.
- Cloud burst: produced ~3% of annual photon production in the last year in 2 hours

The Science Case

IceCube (Slide 6)

- Digital Optical Module contains photo tube that detects light
- Facility has diverse science goals, but talk focuses on high energy astrophysics, the motivator for this exercise
- Building at South Pole: towers on side (cable trays)

High Energy Astrophysics – Science case for IceCube (Slide 7-8)

Shows reason to use something other than light. While normally most of what we learn from universe is via light, universe is opaque for very high energy photons (light not transmitted).

- Use neutrinos to verify sighting of high energy particles
- Then direct telescopes, which measures low energy from source; thus, understand source.
- Top graph: X-axis energy from photo; Y-axis: distance it will fly
- Bottom graph: Measured spectrum. Extragalactic neutrinos (red) – high energy neutrinos
  - High energy neutrinos from outside the solar system (Slide 8)
  - IceCube showed there are high energy neutrinos from outer space.

Understanding the Origin (Slide 9)

Cause of these high energy events?

- Synchrotron radiation source causes protons which interact with photons from synchrotron radiation.
- Source with highest energies produces light and neutrinos. If see neutrinos, look for light.
- As high energy source will have very broad spectrum, looking at visible spectrum (low energy photons) to understand activity
- Direct telescope based on what IceCube sees.
  - Likewise Gravitational waves: Tell telescopes where to look based on what LIGO sees.
  - “Multi-messenger astrophysics”

The v detection challenge (Slide 10, graph)- Neutrino detection

- Left graph: Deposited energy vs. median angular error - dominated by systematics
- Ice properties: model IceCube uses. Change is function of depth. Understanding ice that detector is sitting in is important to beat down systematics
- Different ice models impact where you think neutrino came from – systematic effects are severe

Goal:

- improve electron power reconstruction to increase neutrino flux and make more observations
- Improve ice model to improve pointing accuracy of detector – in order to conduct multi-messenger astrophysics

First evidence of an origin (Slide 11)

Sept 2017- IceCube directed observation of high energy neutrino

IceCube’s Future Plans (Slide 12)

Upgrade to improve facility; if funded, will take data until 2050.

Capabilities establishing today would be used for next 20-40 years.

Details on Cloud Burst

Idea (Slide 14)- IceCube- normally do steady state. Not benchmark burst capabilities.

- Part of objective: how fast can burst. Want burst capabilities.
- What is current bursting capability and IceCube capable of running on exascale systems?
    - IceCube submits photo propagation workflow to HTCondor pool established across global cloud capacity and run GPU bursts
    - Handle input, jobs on GPUs and output as a single globally distributed system

Gobal HTCondor pool (Slide 15)

Relies on Condor for resource orchestration; Dedicated HW setup to avoid disrupting OSG production system

HTCCondor Distributed CI (Slide 16)

Special software deployment

- Each region in world had fixed set of collectors (aggregation piece of condo deployed in that cloud region; i.e., VMs registered with one of 28 cloud local collectors→talk to San Diego collector attached to negotiator->did global scheduling.
- Pieces installed in each cloud region; built HTCCondor system that managed entire workload during burst

Using native Cloud storage (Slide 17)

Pre-staged Input data into native Cloud storage (storage cost negligible compared to GPUs)

- Kept simple: 18/28 regions had storage used.
- To prepare, measured intraregion transfer bandwidth between all regions in cloud and providers (complete measurement of IO capabilities)

Deployed simple wrappers

Conclusion: No difference when comparing job performance that pulled data across regions vs. performance of jobs using native storage

Testing Ahead of Time & Science with 51k GPUs achieved as peak performance (Slide 18-19)

- Tested ahead of the run with 90k jobs running;

  To assemble at this scale, must run globally across all regions
- Actual jobs running simulation for IceCube (black curve). Ramped down once sustained peak at 51,500 GPUs
- Lessons learned: how to deal with Cloud resources. When ramped down, harder to turn Cloud off than it to turn on. Paid more in struggle to ramp down.
- 8 generations of NVIDIA GPUs used. See number of GPUs used (chart). If not use all GPUs available for sale, would not have reached peak.

Heterogenous Resource Pool (Slide 20)

- Different from typical supercomputer - Globally distributed; maximally heterogenous.
- No 1 region or GPU-type dominates – must use all in order to reach scale**.**
- Top Left: only region, not GPUs' source disclosed; roughly equal amounts exist in each region

Science Produced (Slide 21, pie chart)

Pie chart: Events processed per Cloud Region

- Distribution more heterogenous than the distribution of the prime resources that IceCube has for its own prime resources
- If restrict to largest region, only get 10.8% of the total; Can only be done through HTC and distributed HTC.
- This is the only computing paradigm that can be used

Performance and Cost Analysis
<u>Performance vs. GPU type</u> (Slide 23, pie charts)
Time spent per GPU type: (e.g., V100- 42% done in V100 in 19% of wall time)
Events processed per GPU type – amount of science accomplished on each GPU type

<u>IceCube Performance per dollar</u> (Slide 24, chart)
GPU types used: compared to V100, P100, T4 and M60

- K90 and K520 not have good science per dollar performance, with some exceptions- cannot provide vendor-specific information (See Slide 19)
- Second row (relative to V100)
- 4 GPUs and given 2 other GPUs that IceCube uses for on-prem (what IceCube purchased)
  <u>Discussion</u>: IceCube workflow is single precision and not care about ECC memory; ideal software to run on cheap GPUs.
- Last row: actual amount paid on each provider, assuming would have paid spot prices everywhere. 2 or 3 providers charged “spot” and third charged “on demand” prices.
  - V100, P100 and M60: roughly equal in science/dollar
  - T4: most cost-effective way of doing business for Ice Cube in the cloud today
  - Change over time as things evolve
- For high end GPUs: IceCube performs better than TFLOP32 (compare first 2 rows)
- Science/dollar is 1.5 better for T4 than V100 because software only requires 32-bit
- Price differential between vendors (~10-30%) and on demand vs. spot (~x3)
- If buy gaming GPUs in violation of NVIDIA’s sales rules, another factor of 3 (~x3)

IceCube and dHTC
<u>IceCube Input Segmentable</u> (Slide 26)
Of 10.2B events that we produced, ran 175k GPU jobs. Arbitrary – how many events produced in job

- IceCube prepared 2 types of input files: 1) designed for K80 and K520; 2) other run by newer GPUs. Differed only by number of events per file. 1 file/GPU
- Left graph - Distribution of runtime (T4)- normally identical files; spread reflects unpredictable timing of simulations. Identical: Data local and remote, so will run data remotely in future.
- Right: K80 a bit shorter in run time; can choose runtime of jobs by preparing input file appropriately.

<u>Applicability Beyond IceCube</u> (Slide 27)

- All large instruments and midscale instruments can think of, can do this.
- A large fraction of Deep Learning, but not all.
- Anything with bundles of independently schedulable jobs that can be partitioned to just workloads, can do this.

<u>Cost to Support cloud as a “24x7” capability</u> (Slide 28) – can use slide to calculate cost on-prem vs. cloud
Paid much more because of 1) difficulty in paying for GPUs; 2) K80 and K520 - not cost-effective.

- Now, would only run on spot price –because only way could have all of their GPUs; preemption by regular customers ok because negligible cost <2% of jobs were killed.
- Today, if only run on spot $15k per 300 PFLOP32 hour

Discussion: Numbers for CPU vs. GPU comparison? Cost for executing on CPUs?

  - Unaffordable if wish for 300 PFLOP.
  - Physics: so large a differential, IceCube does not run photo propagation on CPU.
  - Code more efficient on GPU than on CPUs.

Burst: cheap from human perspective

- FTE and a half: to build infrastructure and run and deal with budgetary issues.
- For 24/7 offering, up to 32 hours of dHTC computing

Discussion

Regarding Cloud providers offerings: Can't do this without going across all 3 providers, can't reach this scale with less effort.

  - Exascale in cloud exists: scale infinite because vendor will do what is needed
  - Compare to finite on-demand business: No exaFLOP in the cloud can be bought for an hour

How much available in 6 months is unknown (constantly building out)

- Would predict that number of GPUs more important than number of FLOPs (Slide 19). In a year, will likely have same number of GPUs available in cloud for sale, on-demand, but more high end and less low end.
- This experiment provides ability to roughly estimate availability of GPU without having to repeat experiment.

IceCube ready for Exascale (Slide 29)

- Just missing capacity
- Computing for large collaborations: Fundamentally fits into global distribution: Should use paradigm to globally put resources and use exascale systems to extent not interfere with people who need those systems

Discussion:

  - If consider networking where no parallel jobs. If interprocess communication, can't do in cloud (Largest scale is in Microsoft which offers InfiniBand-connected clusters up to scale of 800 V 100
  - Doing for first time: Paid $100K in order to do this in its entirety (includes being charged on-demand by a provider). Now know how to buy 300PLOP hour for $15K

CY2020 Speakers

Individual topics instead of multi-month series; lay out 6-8 topics and assign to someone to identify potential speakers. Other speaker nominations? Can email to Joyce Lee

Roundtable

DOE: Rich Carlson (No updates)

CASC: Alan Sill

Survey: https://go.osu.edu/CASCsurvey

Next Meeting: February 5, 2020 (12 noon ET)