Noise in the Clouds

Paper by Densi et. al.

James Osborne
Eddie Brees

The Point

- Network Performance in Modern Computing
  - Central to Cloud and High Performance Computing (HPC) ecosystems.
- Noise - variations or inconsistencies within use (Latency, Bandwidth)
- Purpose of study
  - Objectively compare network performance between on site and cloud-based HPC systems.
  - Inform future scalability and cost models.

Cloud Providers

Amazon AWS,
Google GCP,
Microsoft Azure,
Oracle Cloud.

On Site Systems

Daint
Alps

Network Performance

- Network Performance needs a standard way to be tested
- The “Ping Pong” test
  - Measures Round Trip Time (RTT) by sending data packet from one node to another
  - RTT / 2 = latency
  - Tested various concurrent ping pong connections for bandwidth
- Cloud-based systems (AWS and GCP), often peaked at bandwidth well below their claimed 100 Gb/s
  - Theorized that may be due to presence of several 25 Gb/s Network Interface Cards (NICs), rather than an singular 100 Gb/s NIC

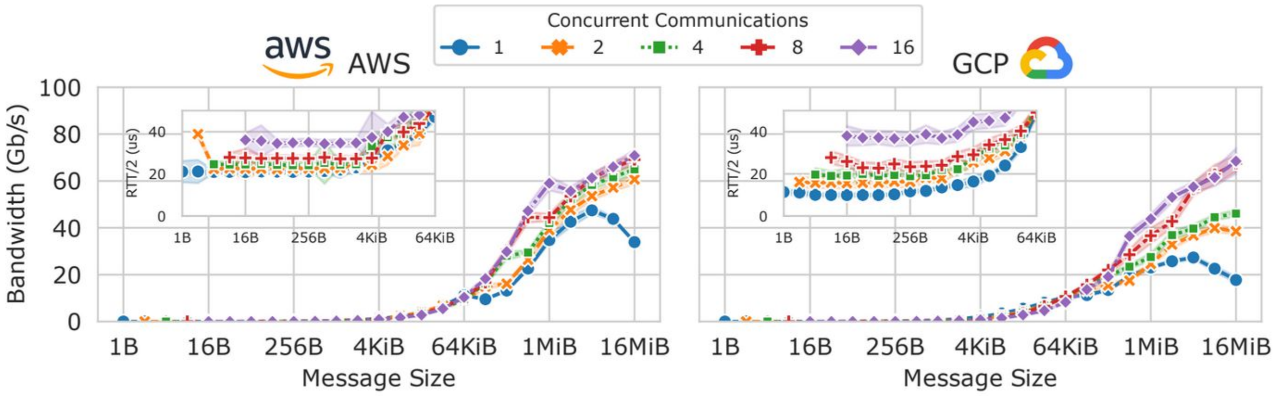

Latency and Bandwidth Findings

- Same testing was done for other systems
- Latency
  - Azure , Oracle, and on-premise systems hit 1-2 microseconds for HPC instances
  - AWS and GCP were Much higher - around 20 and 10 microseconds respectively.
- Bandwidth
  - Various results, with Azure, Oracle, Alps, and DEEP-EST peaking at around 90 Gb/s for 16MiB messages.

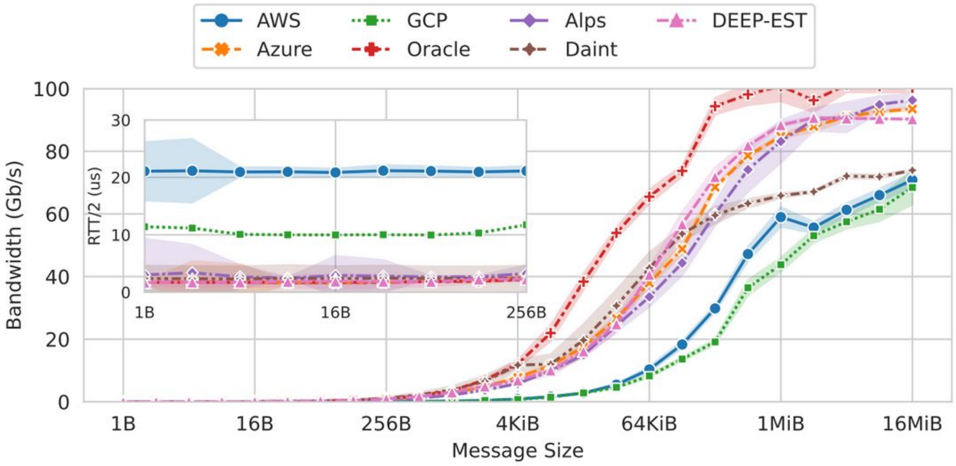

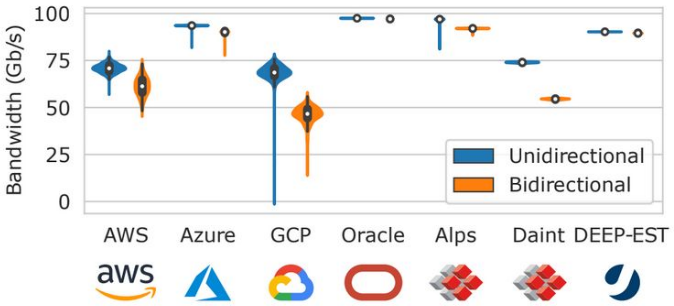

Traffic Burst

- A method called ‘traffic bursting’ was used to see how bursty data transfers might impact the network
  - Used a 1MiB (Mebibyte, 2^20 bytes) "ping-pong" between two nodes, varying the inter-message interval from 0 to 1 second.
- Most providers showed no change in behavior
  - GCP is the exception, and showed large variation, especially with larger inter-message intervals (below)

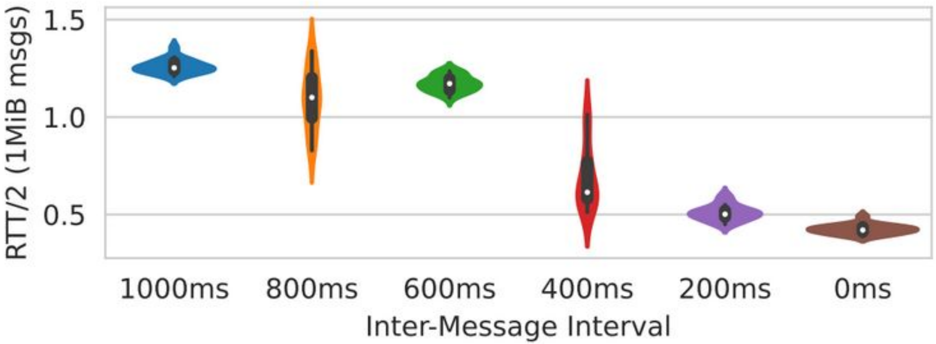

HPC vs. Regular Cloud Instances

- HPC variant is typically optimized for intensity, while regular variants cater to standard cloud applications
- Some observations
  - AWS & GCP: HPC vs Regular instances have similar latencies.
  - Azure & Oracle: HPC instances significantly outpace their regular counterparts in latency, by 10-20 times.

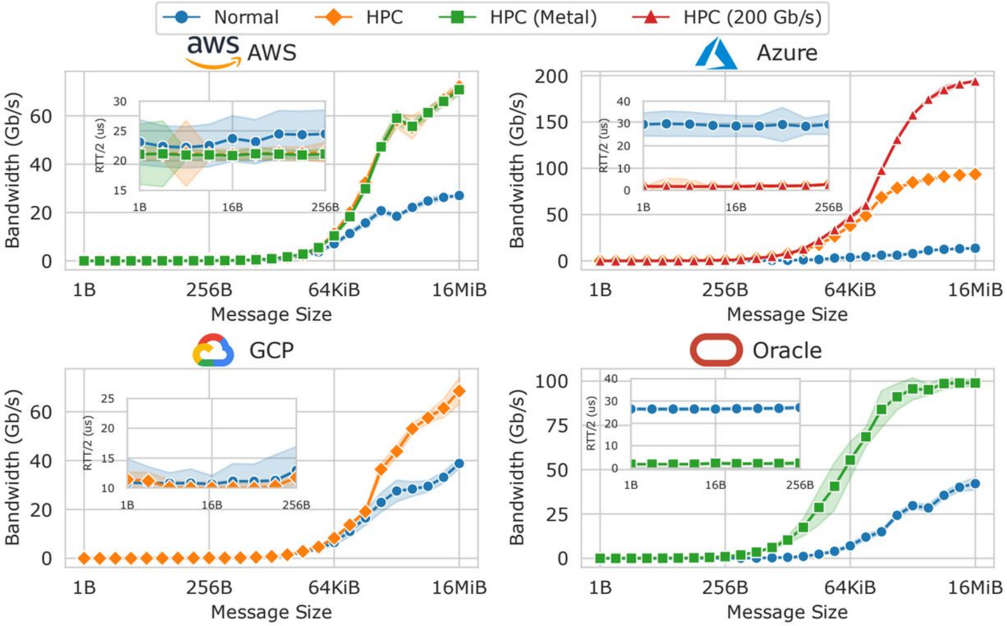

OS Noise

- The interference from the running of the OS, ala the daemons, scheduler, etc.
- Selfish detour test
  - Runs tight loop and records all the iterations larger than 9·tmin, where tmin is the minimum measured time required to complete one iteration.
  - Iterates a set number of times

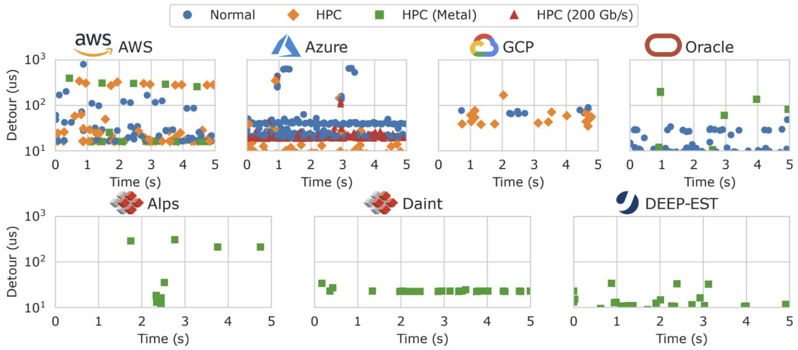

Fig. 8. OS noise for the different instance types described in Table I for the different providers.

Latency Noise

- The noise introduced by the latency in communications
- Ping pong test
  - 1 bit, 1 hour

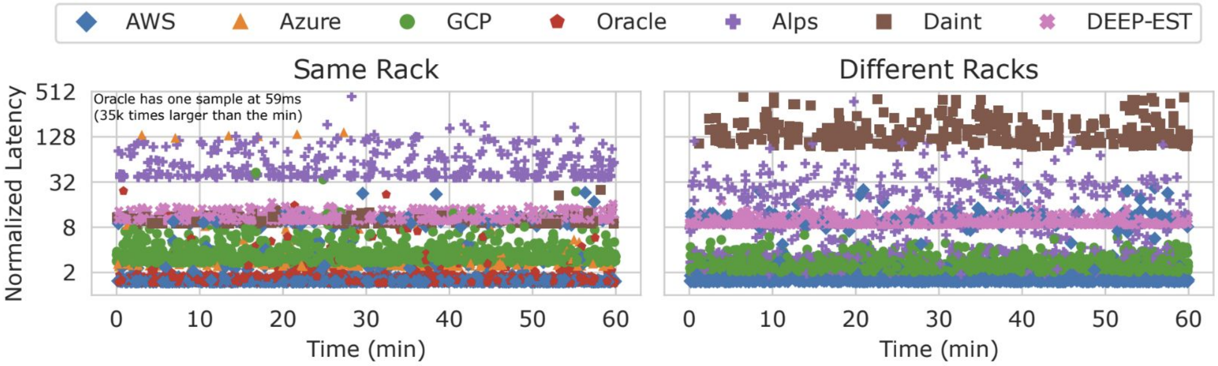

Fig. 9. Latency noise for different node distances for 100 Gb/s HPC instances. Base latency is reported in Table II.

Bandwidth noise

- The noise introduced by bandwidth changes and issues
- Large message ping pong for 1 hour

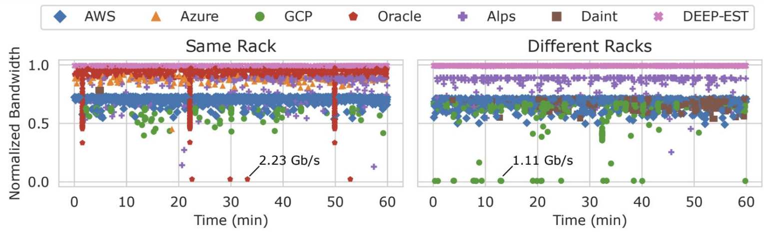

Fig. 11. Bandwidth noise for different node distances for 100 Gb/s HPC instances. Base bandwidth is reported in Table II.

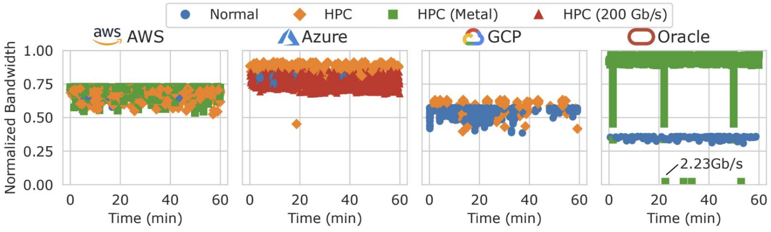

Fig. 12. Bandwidth noise for different instance types. Base bandwidth is reported in Table II.

Simulations

- Ran simulations to test for the effects of noise on application performance

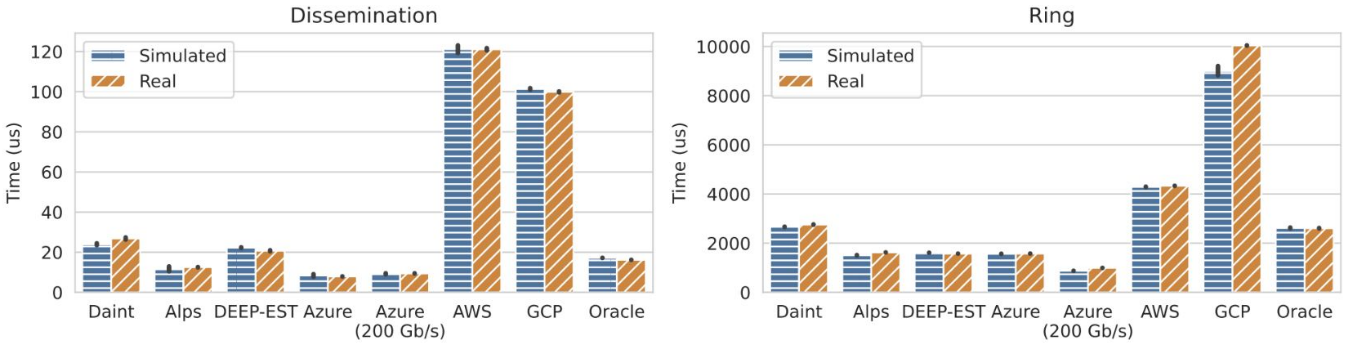

Fig. 13. Comparison between measured and simulated times (on HPC instances) for 16B dissemination and 16MiB ring collectives on 16 nodes. Vertical lines at the top of the boxes represent the 95% confidence interval.

More on dissemination

- Only considered HPC systems from cloud providers
- Effects on dissemination algorithms

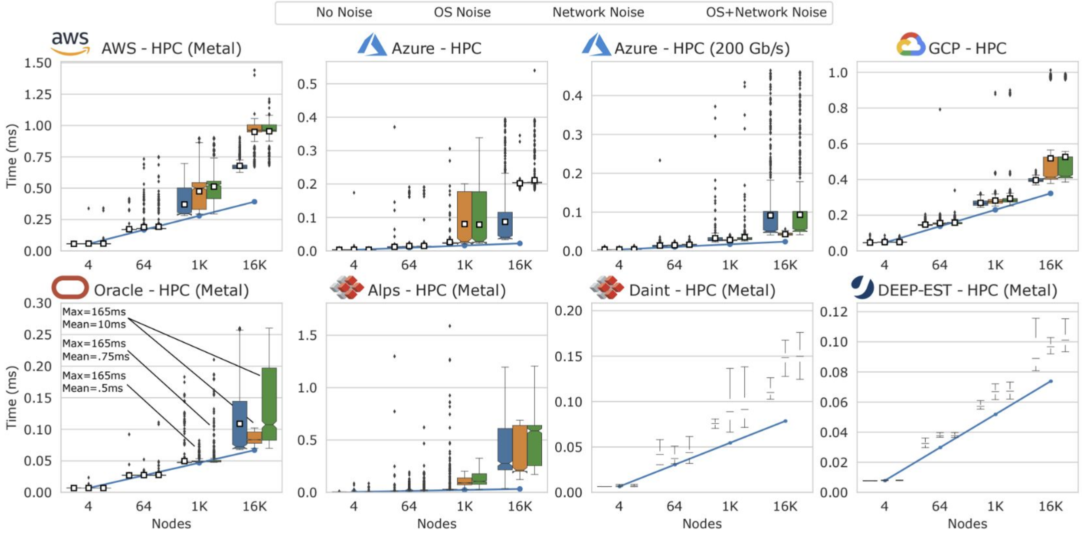

Fig. 14. Simulation of the scalability of a 16B dissemination algorithm, with and without OS and network noise. Y-axes have different scales.

More on rings

- Only considered HPC systems from cloud providers
- Effects on dissemination algorithms

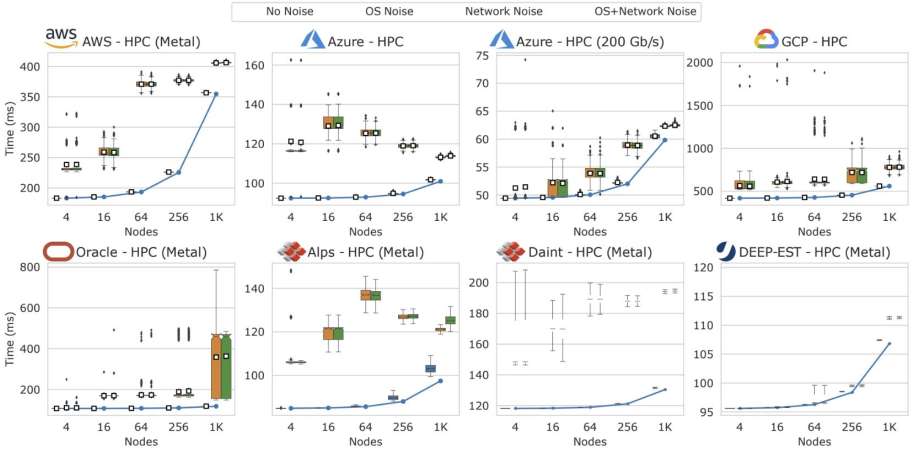

Fig. 15. Simulation of the scalability of a 512MiB ring allreduce collective on HPC instances.

Cost

- Measuring the cost of noise as the number of nodes increases

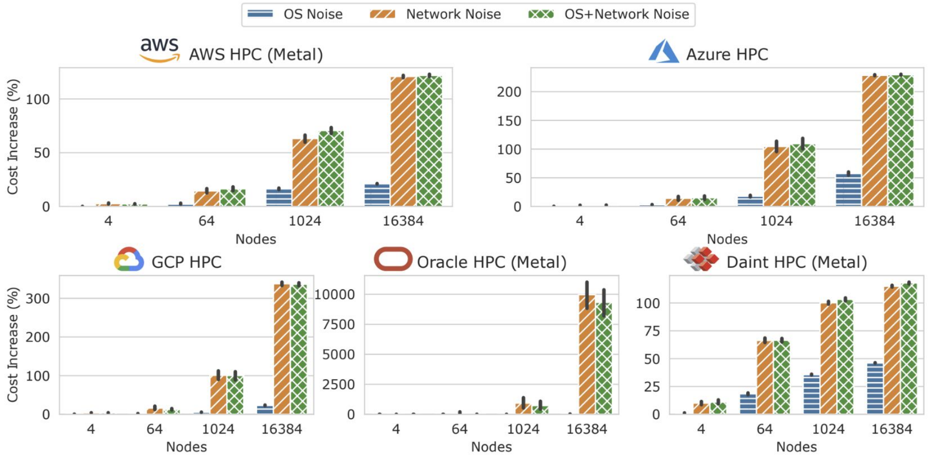

Fig. 16. Simulation of the cost 128 × 128 double-precision matrix multiplications followed by a 128KiB dissemination, on different node count. Black vertical lines at the top of the boxes represent the 95% confidence interval.

The end

Questions?