Understanding performance of applications based on 100GbE and RDMA

Jichun Wu

4 July 2024

Motivation

RDMA:

- 70% of traffic in Azure is RDMA[1]
- Meta GenAI 24K-GPU clusters with RDMA [2]

Challenges:

- 100 Gbps and above
- Nature of RDMA

[1] Bai et al., “Empowering Azure Storage with RDMA”, 2023.

[2] https://engineering.fb.com/2024/03/12/data-center-engineering/building-metas-genai-infrastructure/

Introduction

Communication delay affects application performance:

- **ML**: GPU idle time from 11% to 70%[3].
- **HPC**: Latency variation slow down by 3.5x[4].
- **Datacenter**: 50 µs latency degrades Memcached by 50%[5].

- How does each layer contribute to the end-to-end latency?
- How and why are different applications affected differently?

[3] Gebara, Ghobadi, and Costa, "In-network Aggregation for Shared Machine Learning Clusters", 2021.
[4] Underwood, Anderson, and Apon, "Measuring Network Latency Variation Impacts to High Performance Computing Application Performance", 2018.
[5] Zilberman et al., "Where Has My Time Gone?", 2017.

Approaches

Measurement

- **Software**: instrumentation in OpenMPI and rdma-core library

Emulation

- **FPGA**: NRG to control latency and bandwidth

Calibration

- **FPGA**: mini-OSNT to validate NRG and network setup

Instrumentation

- Tracing framework: KUtrace
- Target: OpenMPI, rdma-core package

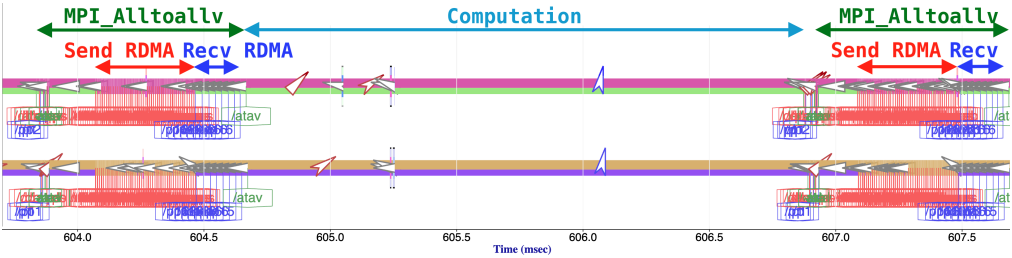

Figure: Different phases of a HPC benchmark

NRG and validation

NRG[6]:

- Line-rate delay injection
- Rate control
- Statistics collection

mini-OSNT:

- Line-rate packet generation
- Timestamping

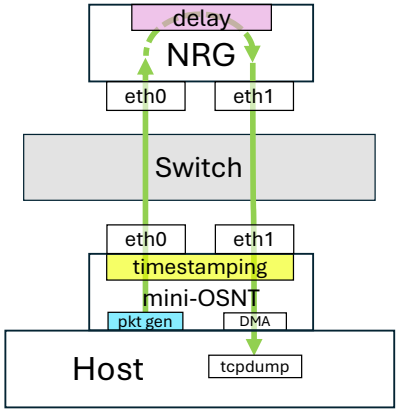

[6] Zilberman et al., "NRG: A network perspective on applications' performance", 2021.

NRG latency verification results

- 500 ns insertion
- 655.36 µs maximum
- 9 ns variation

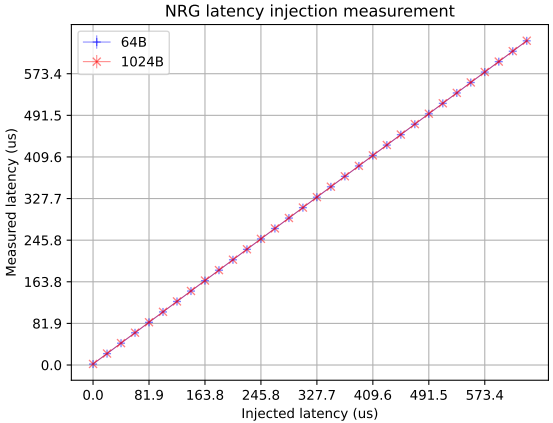

Motivation experiment

- NASA Parallel Benchmark

	Operation types	Specification
ft	floating point	Discrete 3D fast Fourier Transform, all-to-all communication
is	keys ranked	Integer Sort, random memory access
cg	floating point	Conjugate Gradient, irregular memory access and communication
mg	floating point	Multi-Grid on meshes, long- and short-distance communication
lu	floating point	Lower-Upper Gauss-Seidel solver

Table: NASA Parallel Benchmark Specifications

Motivation exp. results

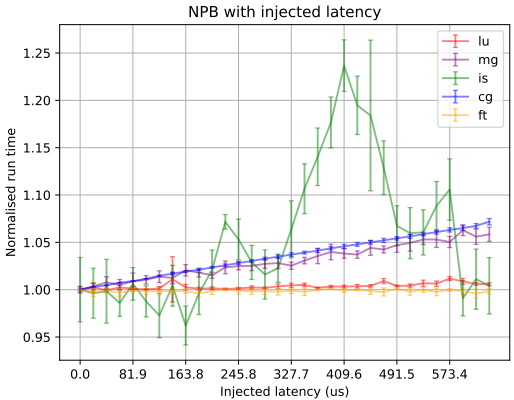

Figure 1: Our results with latency injected by NRG

[7] Underwood, Anderson, and Apon, “Measuring Network Latency Variation Impacts to High Performance Computing Application Performance”, 2018.

Motivation exp. results

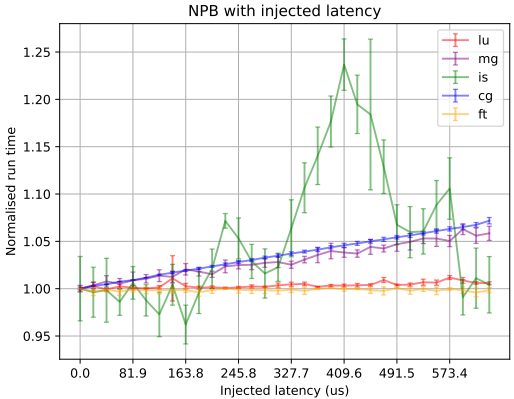

Figure 1: Our results with latency injected by NRG

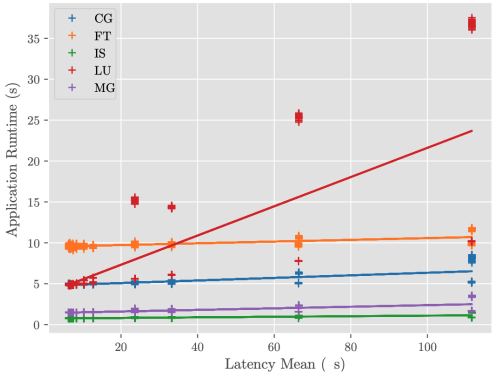

Figure 2: Previous results by Underwood et al.[7]

[7] Underwood, Anderson, and Apon, "Measuring Network Latency Variation Impacts to High Performance Computing Application Performance", 2018.

Plans

- Investigate the integer sort (IS) benchmark
- Benchmark more applications
  - File systems
  - Disaggregated memory
  - Key-value store
  - ML training
- Use application traces, explain the benchmark results