Machine Learning for Data Transfer Anomaly Detection

Sarah Cooper, Masud Bhuiyan, Engin Arslan
University of Nevada, Reno

Abstract

Distributed scientific applications require ever-increasing transfer rates but seldom reach promised throughput speeds in high-performance research networks for a variety of reasons that are poorly understood. Understanding the true underlying reasons for poor transfer performance is the key to mitigate them and deliver the promised transfer speed. However, the involvement of multiple end systems, dynamically changing background traffic, and the complexity of today's networking infrastructures turn it into a complicated and time-consuming process. This project builds a set of novel models and algorithms to detect anomalies for end-to-end data transfers in high-speed networks.

Problem Statement

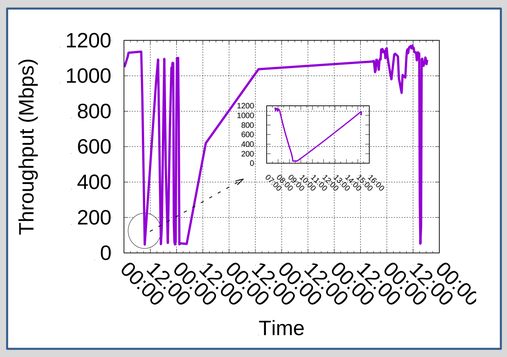

Fig 1.a: Memory-to-memory transfer

- A week-long *memory-to-memory* transfer between BlueWaters and Comet supercomputers
- 10G network bandwidth, 32 ms RTT
- Throughput fluctuated between 50 Mbps and 1200 Mbps

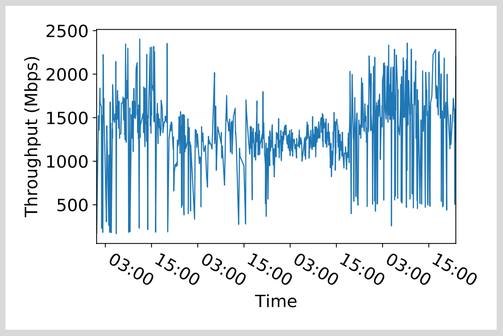

Fig 1.b: Disk-to-Disk Transfer

- A 7GB file is transferred (Fig 1.b) repeatedly between Stampede and Comet supercomputers
- 10G network bandwidth, 40 ms RTT
- Throughput fluctuated drastically between 150 Mbps and 2500 Mbps between consecutive runs

Problem: Although network congestion is typically the common suspect for low transfer performance, other reasons such as misconfigured servers, I/O congestion, and overloaded data transfer nodes can also lead to performance anomalies

Observation: Throughput of data transfers in high-speed networks are significantly less than available bandwidth and exhibit fluctuating behavior.

Objective: Can we identify the root causes of the performance problems in high performance networks in real-time such that we can make online adjustments to alleviate the impact of anomalies or plan necessary arrangements to enhance performance for future transfers.

Inspecting data transfers

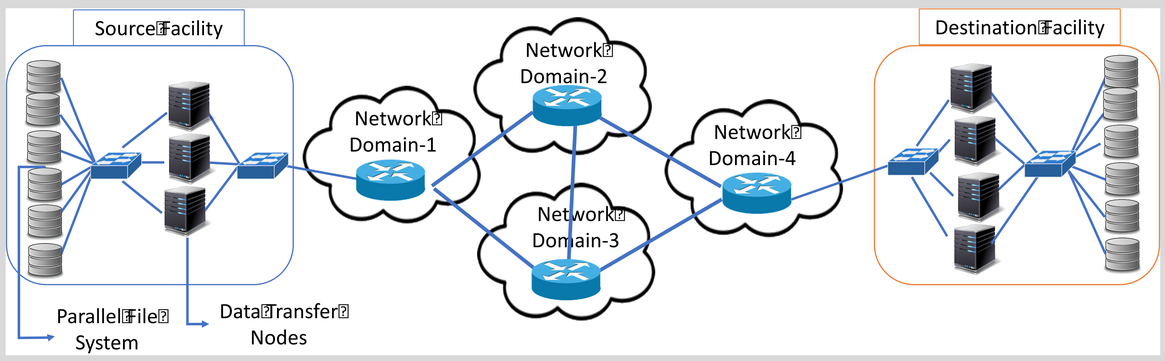

Figure 2: End-to-end transfers read/write from/to many storage servers, use multiple DTNs, and pass through many network domains before reaching their destination.

Challenge: Data transfers in high performance networks travel through many storage servers, data transfer nodes, and network domains, necessitating data collection from many systems to identify the underlying reasons for performance anomalies.

Proposed Research

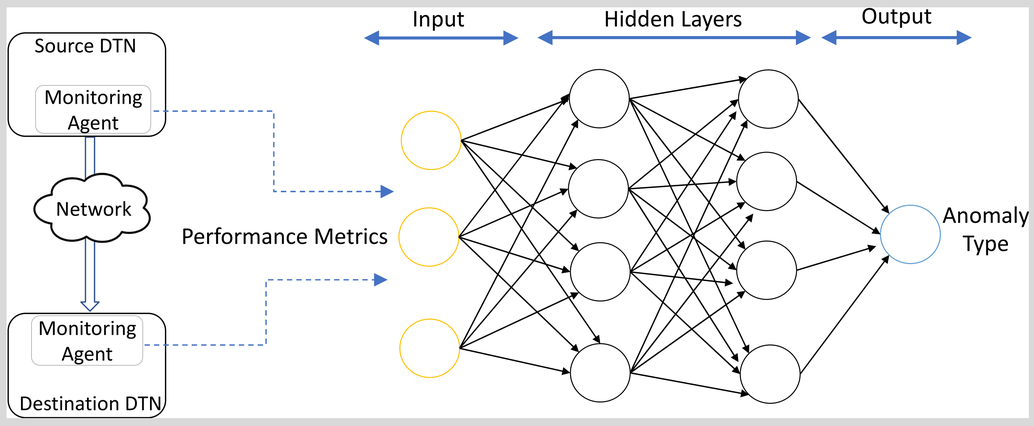

Figure 3: Proposed system architecture

- **Monitoring Agents** gather real-time performance metrics for data transfer node, network, and file system in real-time.
- ***Performance Metrics:*** Network: RTT, pacing rate, congestion window size, retransmission timeout and retransmitted packet count. Data Transfer Node: CPU, memory, and NIC utilization, TCP buffer size and cache statistics. File System: I/O read/write rates, the number of I/O requests per second, and I/O wait time

Data Collection and Model Training

- Used Chameleon testbed to run file transfers between Chicago, IL and Austin, TX sites where network bandwidth is 10 Gbps and RTT is 32ms
- Reproduced and labeled 10 common anomalies at 3 different severity levels; e.g., packet loss rate of 0.1%, 0.5%, and 1%.
- Test dataset includes performance metrics for disk, network and transfer nodes collected at 100ms intervals using *sar, ss*, and *iostat.*
- Neural Network model processes the gathered parameters to predict the underlying reasons of performance issues.
- Trained a 5-layer NN with 50 neurons using Rectified Linear Unit (ReLU) activation sequence for the hidden layers and a SoftMax activation sequence for the output layer and Adam optimizer.

Results

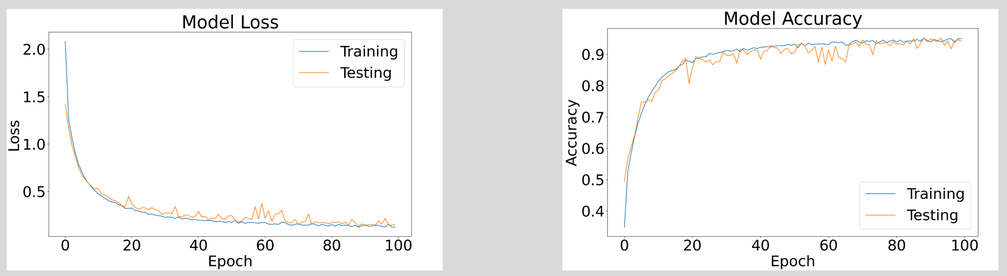

Fig 4: Accuracy and loss rate of neural network model over training time

- Model accuracy reaches to 90% after 20 epoch and 93% after 100 epoch
- Model loss decreases to 0.5 after 20 epoch and 0.3 after 100 epoch

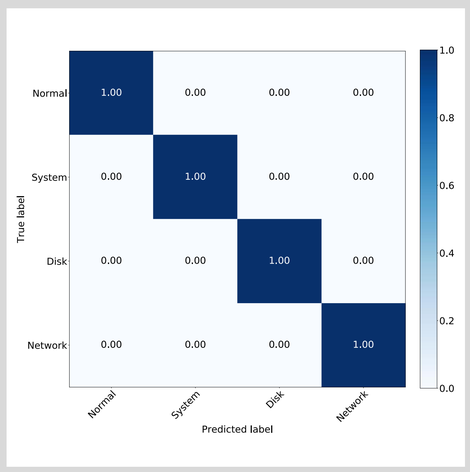

Fig 5: Evaluation of NN model

- 93% accuracy in predicting the anomaly type with correct severity level
- 96% accuracy in finding the right anomaly category
- **99.96% accuracy** in finding the anomaly type that are combined into disk, network, and transfer node (i.e., system) categories.
- Future work is to apply other machine learning models to improve the accuracy rate.

Acknowledgement

This project is in part sponsored by the National Science Foundation (NSF) under awards 1850353 and 2007789. Some of the results presented in this paper were obtained using the Chameleon testbed supported by the NSF.