# Spatio-temporal Analysis of HPC I/O and Connection Data

July 2, 2018
(Presented at SNTA 2018)

Jinoh Kim
Computer Science Department
Texas A&M University, Commerce, TX 75429, USA

TEXAS A&M UNIVERSITY COMMERCE

BERKELEY LAB

# Introduction

- HPC systems advanced significantly in capacity and scale
- The complexity of HPC has also increased
  - Making it hard to identify and predict performance bottlenecks and anomalies
- In HPC, data access and sharing has become more important for many HPC applications
- Identifying I/O and network bottlenecks is one of the first-class requirements to improve efficiency and scalability in HPC systems

# Challenges

- Various logs needed to be analyzed, in a set of layers of software and hardware for I/O and networking, for example:
  - LMT (Luster Monitoring Tool): I/O activities logs
  - MDS (Meta-Data Server): CPU load logs
  - DTN (Data Transfer Node): TCP connection logs (tstat)
- Question: Would there be certain correlations among the different types of activities, which can be used to identify anomalies?

# Data flow to HPC storage (NERSC)

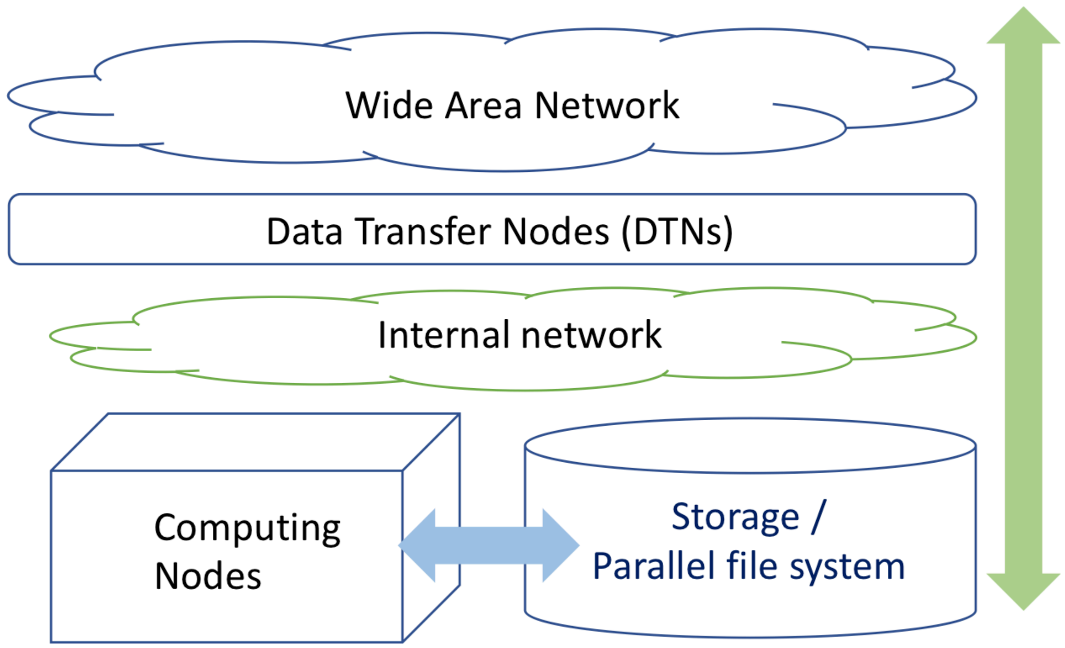


- Data Transfer Nodes (DTNs) as a part of the Science DMZ are dedicated systems for wide area data transfers with special configuration and proper transfer tools for high throughput performance.
- All DTNs at NERSC have four 10-gigabit Ethernet links for wide area network and two FDR IB connections to the each underlying file system.

# View of Lustre file system (NERSC)

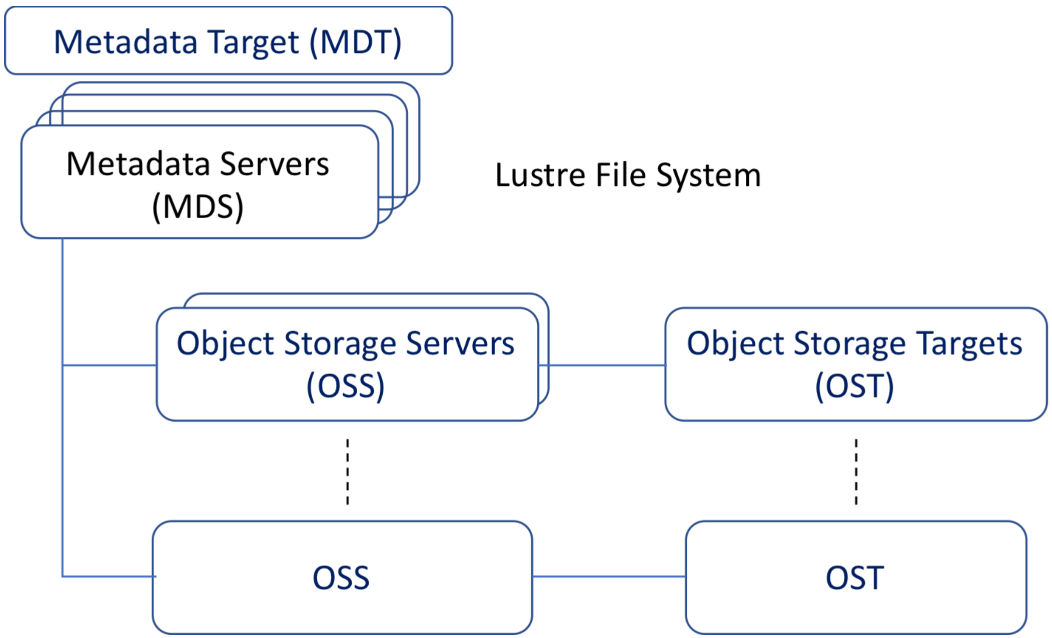


**NERSC**: National Energy Research Scientific Computing Center (http://www.nersc.gov/)

- Lustre file system is composed of Object Storage Servers (OSSs) and Object Storage Targets (OSTs)
- Lustre parallel file system at NERSC is with about 30PB of storage with 248 Object Storage Targets (OSTs)
- File is usually striped for read and write operations accessing multiple OSTs concurrently to increase I/O performance.

# In this presentation …

- Share our preliminary efforts for multi-level log analysis of HPC I/O and connection data
- Manually analyzed the log data to find out potential behavioral correlations
- Report our initial observations from the analysis of a 6-day logs  (January 1–6, 2018) collected from cori
- Data sets include the metadata server CPU load, file system logs sampled once every 5 seconds, and TCP connection records

# Storage RW and traffic tput

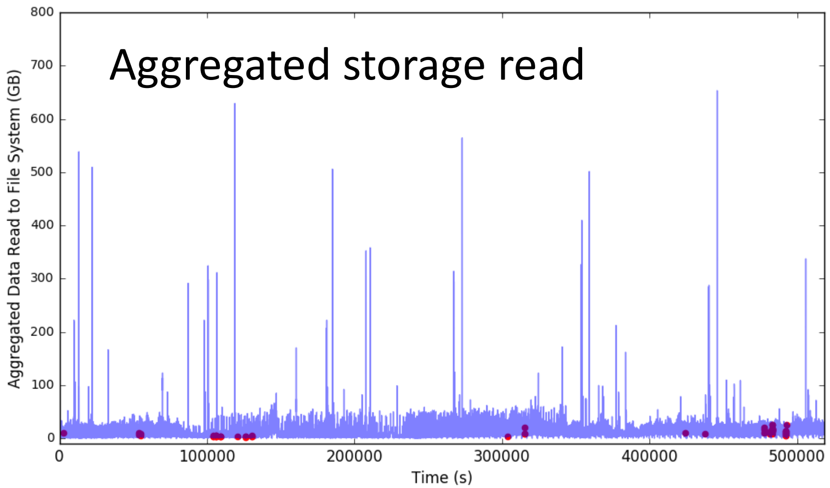


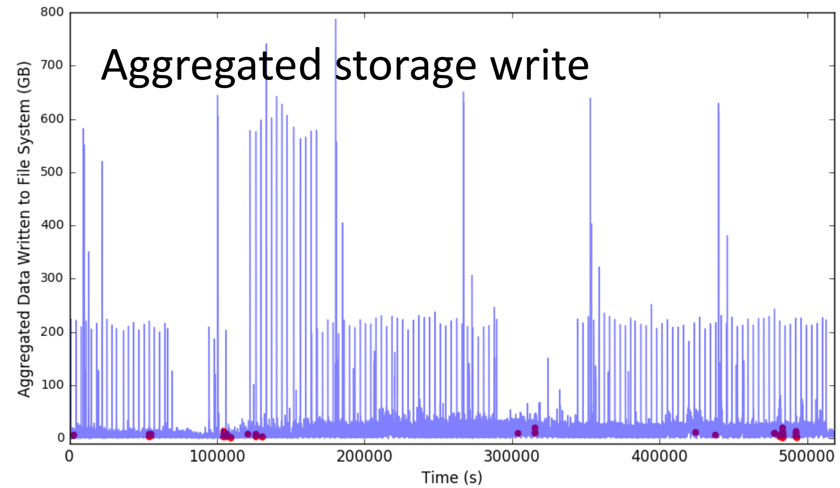


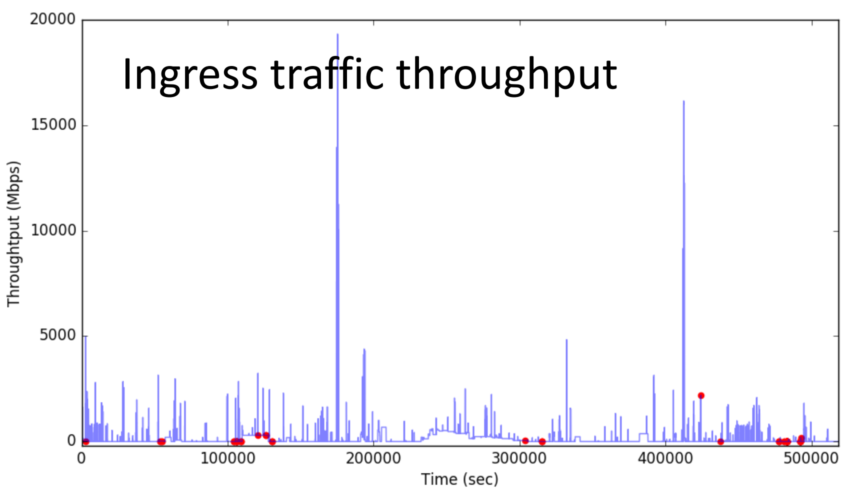


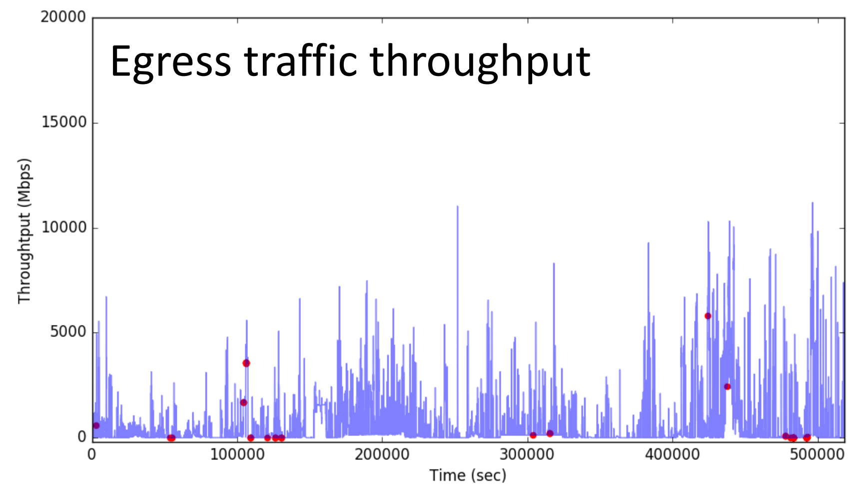


Red dots indicate the time points at which the MDS CPU load is equal to 60% or above.

# MDS CPU load

- Less than 60% over 99% of the collection period
- Storage read/write rates are quite small when MDS is busy
- Network tput does not completely correlated with MDS load
- Thus we focused on I/O and CPU load logs

TEXAS A&M UNIVERSITY COMMERCE

BERKELEY LAB

# t=55100–55200

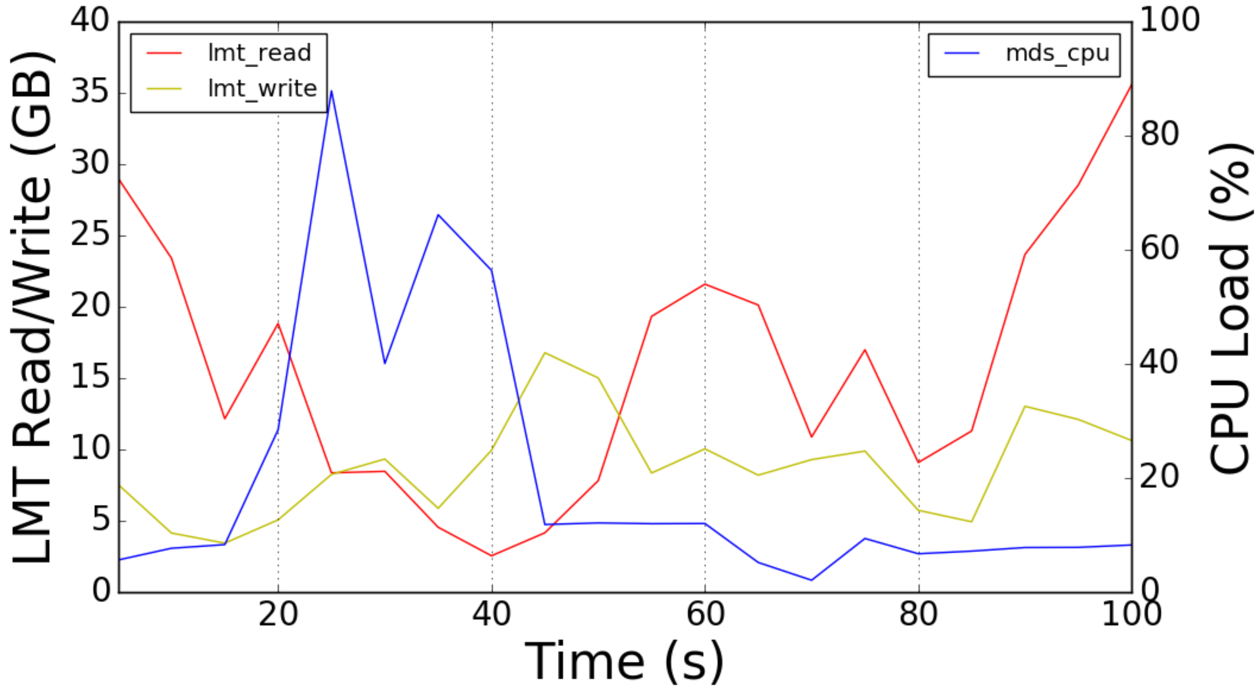


- MDS is busy in t =20–35
- Within that busy interval, storage read rate decreases quickly but gets back once the MDS load goes down at t =40
- Write pattern also shows the access rate goes higher at t =35 which is the starting point that the MDS load goes down

# OST Read (Delta)

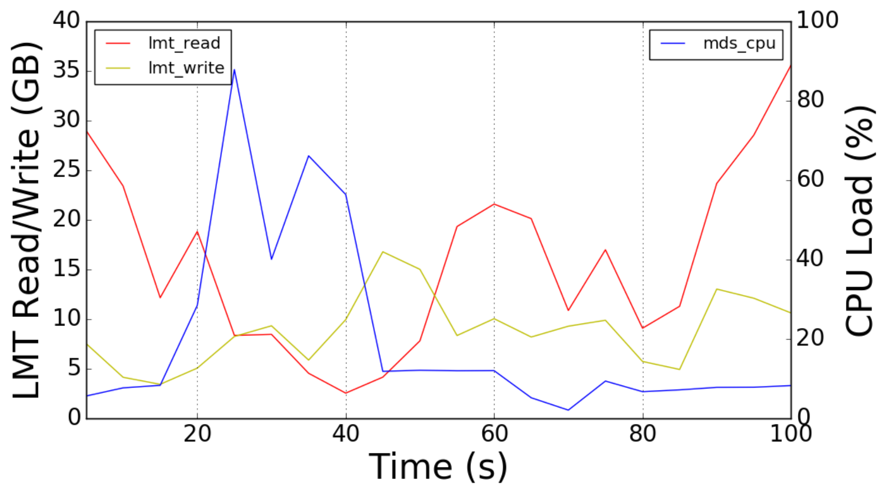


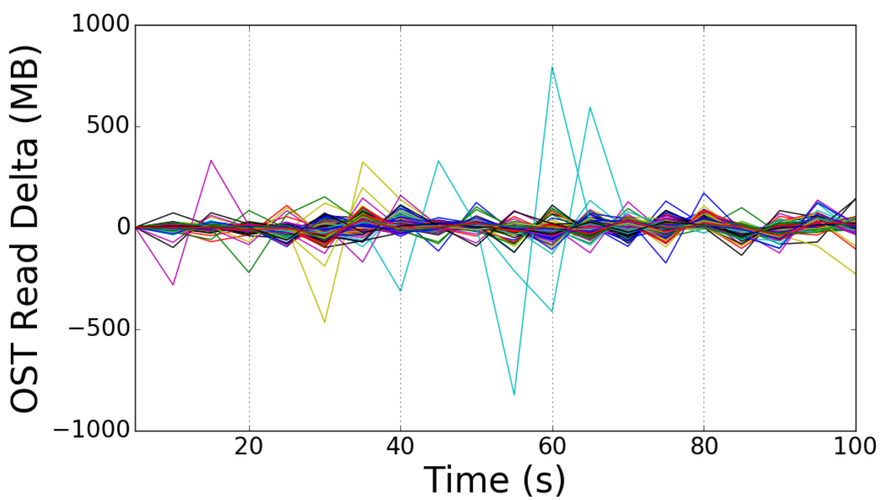


- Delta  = current window data rate - past window data rate (window size = 5 seconds).
- Positive indicates the increase of data rate compared to the past window, and vice versa
- Majority of OSTs changed to a downturn at t=25 and t =35, respectively, at which the CPU usage is relatively high.

# OST Read (Delta)

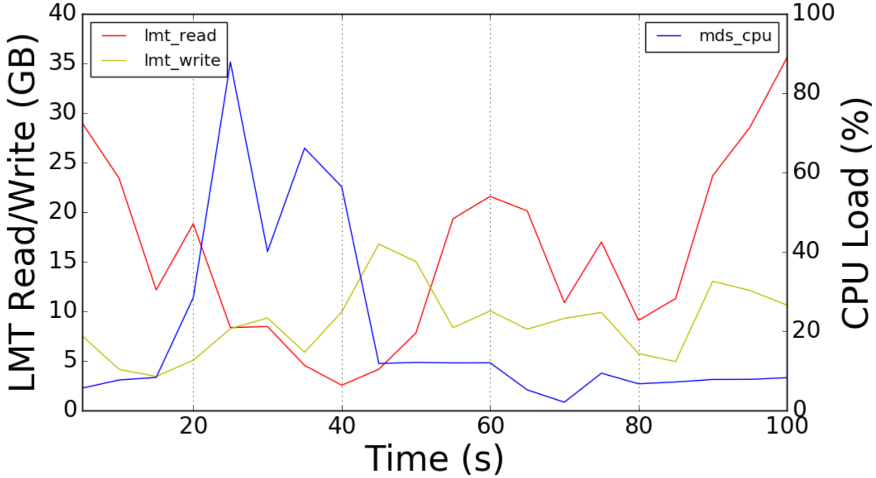


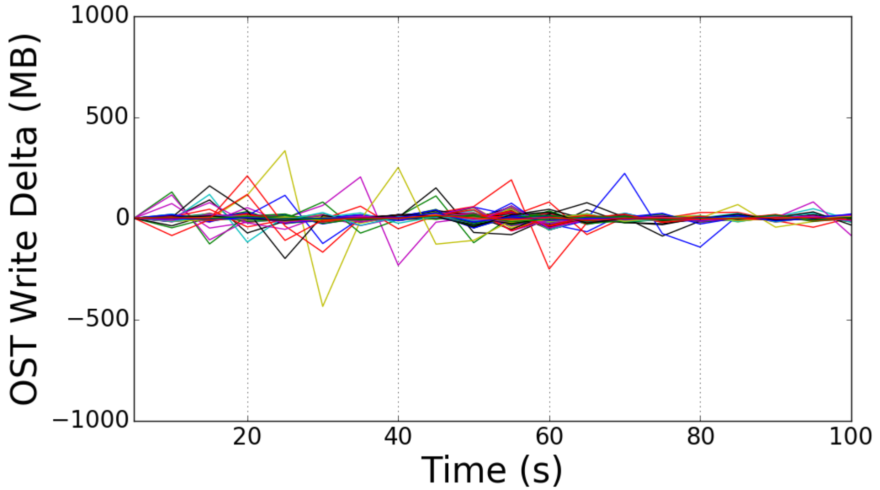


- Delta = current window data rate - past window data rate (window size = 5 seconds).
- Positive indicates the increase of data rate compared to the past window, and vice versa
- Majority of OSTs changed to a downturn at t=25 and t =35, respectively, at which the CPU usage is relatively high.

# t=126220–126320

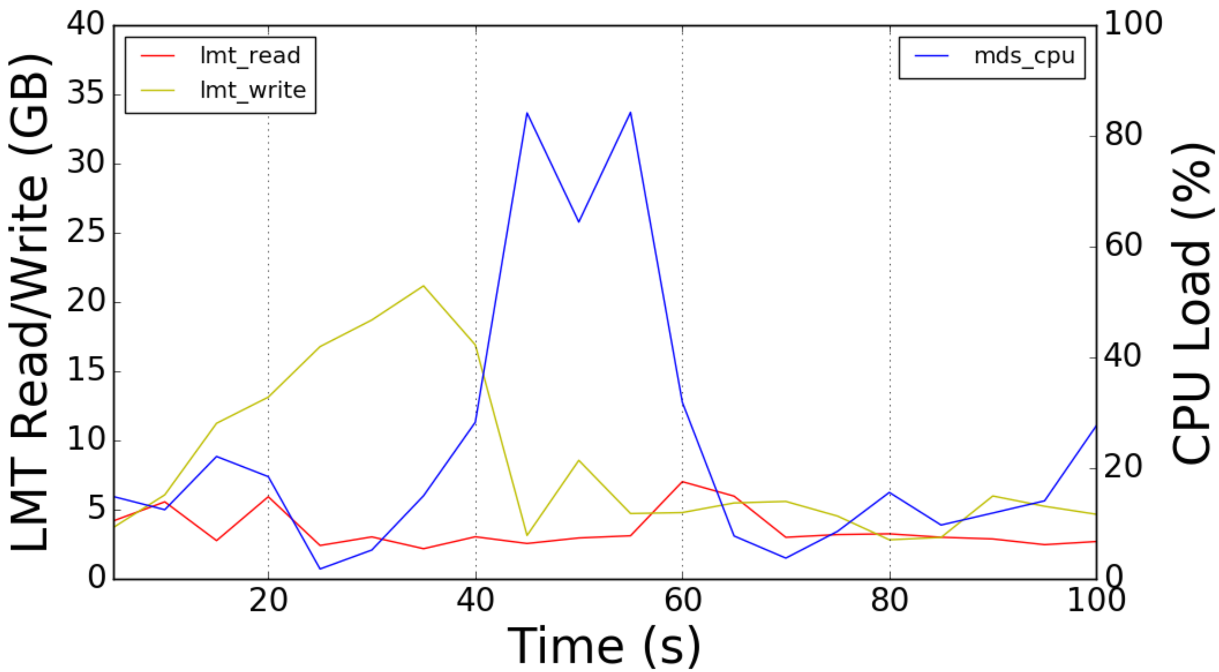


- CPU load is high between t=45-55
- Write throughput changes inversely against the CPU usage in the peak time interval (t =40–60)
- That is, write throughput significantly decreased at the first peak interval, goes up at t =50, and then goes down again at the second peak.
- Read throughput is also reduced after the second peak

TEXAS A&M UNIVERSITY COMMERCE

BERKELEY LAB

# Summary and Directions

- Identifying and predicting performance bottlenecks in the HPC system are the crucial concern
- Our initial work analyzed HPC I/O and connection data to understand the potential correlation among MDS CPU usage, I/O rates, and network traffic throughput
- Observed that the MDS load may affect the I/O rates from the temporal patterns, while network throughput does not show any clear correlation
- Also observed some OSTs showed a degree of correlations with the MDS load patterns in an inversely proportional fashion
- Plan to extensively analyze a long term data spanning months to discover the indication of correlations among different logs and data sources if any
- Using machine learning would ne necessary to analyze long-term data sets to better identify correlations

# Questions?

This work was supported by the Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center.

Contact: Jinoh.kim@tamuc.edu, asim@lbl.gov