Network Bandwidth Utilization Forecast Model on High Bandwidth Networks

Scientific Data Management Group
Computational Research Division
Lawrence Berkeley National Laboratory

Wucherl (William) Yoo, Alex Sim

Motivation

- Increasing Data Volume
  - Efficient resource management and scheduling data movement
    - Predict the network bandwidth utilization between two HPC sites
- Challenge
  - Accurate and fine-grained performance model
    - Computational complexities and variances/burstiness

SNMP Data

- Simple Network Management Protocol (SNMP) Data
  - Collected by ESNet in 2013 and 2014 on each router
  - Connect a pair of large data facilities
    - P1 and P2 between NERSC and ORNL
    - P3 and P4 between NERSC and ANL
    - P5 and P6 between ORNL and ANL
- Univariate time series with bandwidth utilization size and time-scale at 30s interval

Bandwidth Utilization

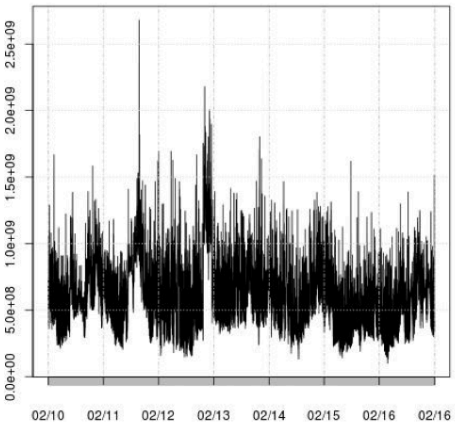

NERSC ➔ ANL (P1)

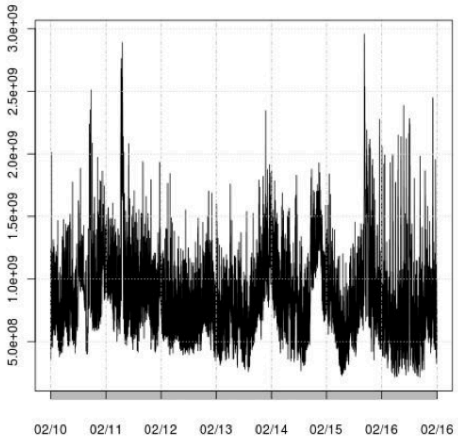

ANL ➔ NERSC (P2)

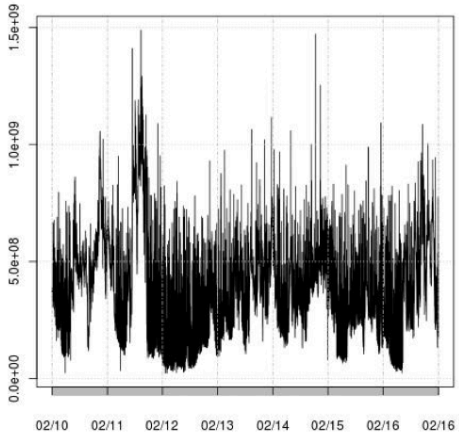

NERSC ➔ ORNL (P3)

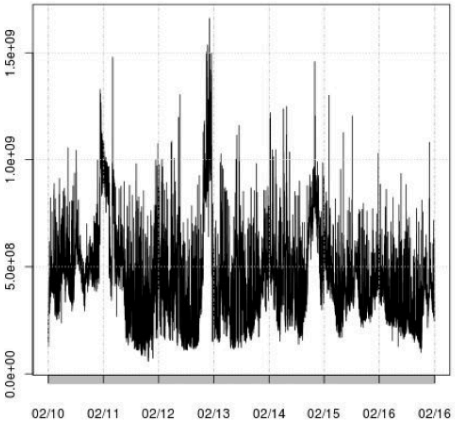

ORNL ➔ NERSC (P4)

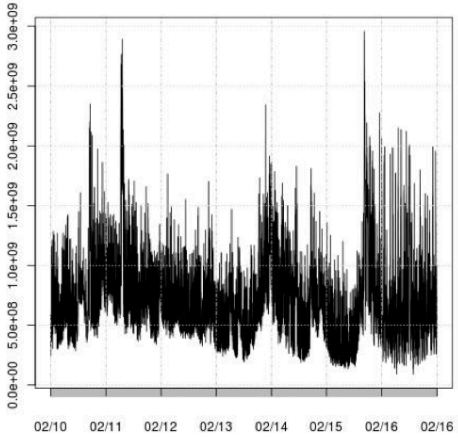

ANL ➔ ORNL (P5)

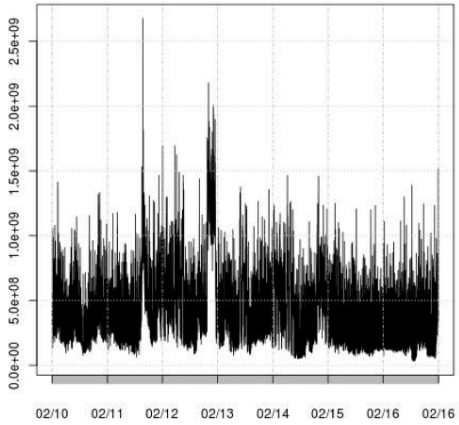

ORNL ➔ ANL (P6)

Our Approach

- Seasonal Adjustment
- Logit Transformation
- Stationarity
- Delayed Model Update

Prediction Model

- Forecast Error

$$e_n(h) = x_{n+h} - \hat{x}_n(h)$$

- Logit Transform
  - lower bound **a**, upper bound ***b***

$$y = \text{logit}(x) = \log\left(\frac{x-a}{b-x}\right)$$

$$\hat{x}_h = (b-a)\cdot\frac{\exp(\hat{y}_h)}{1+\exp(\hat{y}_h)} + a$$

- Prediction Models
  - **ARIMA**, ETS, and Random Walk

Seasonal Adjustment

- STL
  - A sequence of smoothing from Loess (Locally Weighted Regression Fitting)
  - Decomposes the logit transformed SNMP time series into the components S, T, and R.
    - Seasonality S
    - Trend T
    - Remainder R

$$y = y_t = S_t + T_t + R_t$$

Stationarity

- Stationarity
  - The mean or variance of time-series does not change over time and does not follow any trends
- Burstiness
  - When there is a sudden bandwidth utilization change, the time series can be looked as non-stationary
  - Keeping the stationary assumption made less prediction error in our model
    - bursty change may not be a long-term change

Delayed Model Update

- Based on the stationarity, keeping the same model and delaying model updates
  - Instead of refitting, the minimal parts such as auto-correlation and moving averages are updated from the initially fitted ARIMA model

Evaluation

- Forecast Model Comparison
- Logit Transformation
- Training Set Size
- Stationarity
- Delayed Model Update
- Forecast Results

Forecast Model Comparison

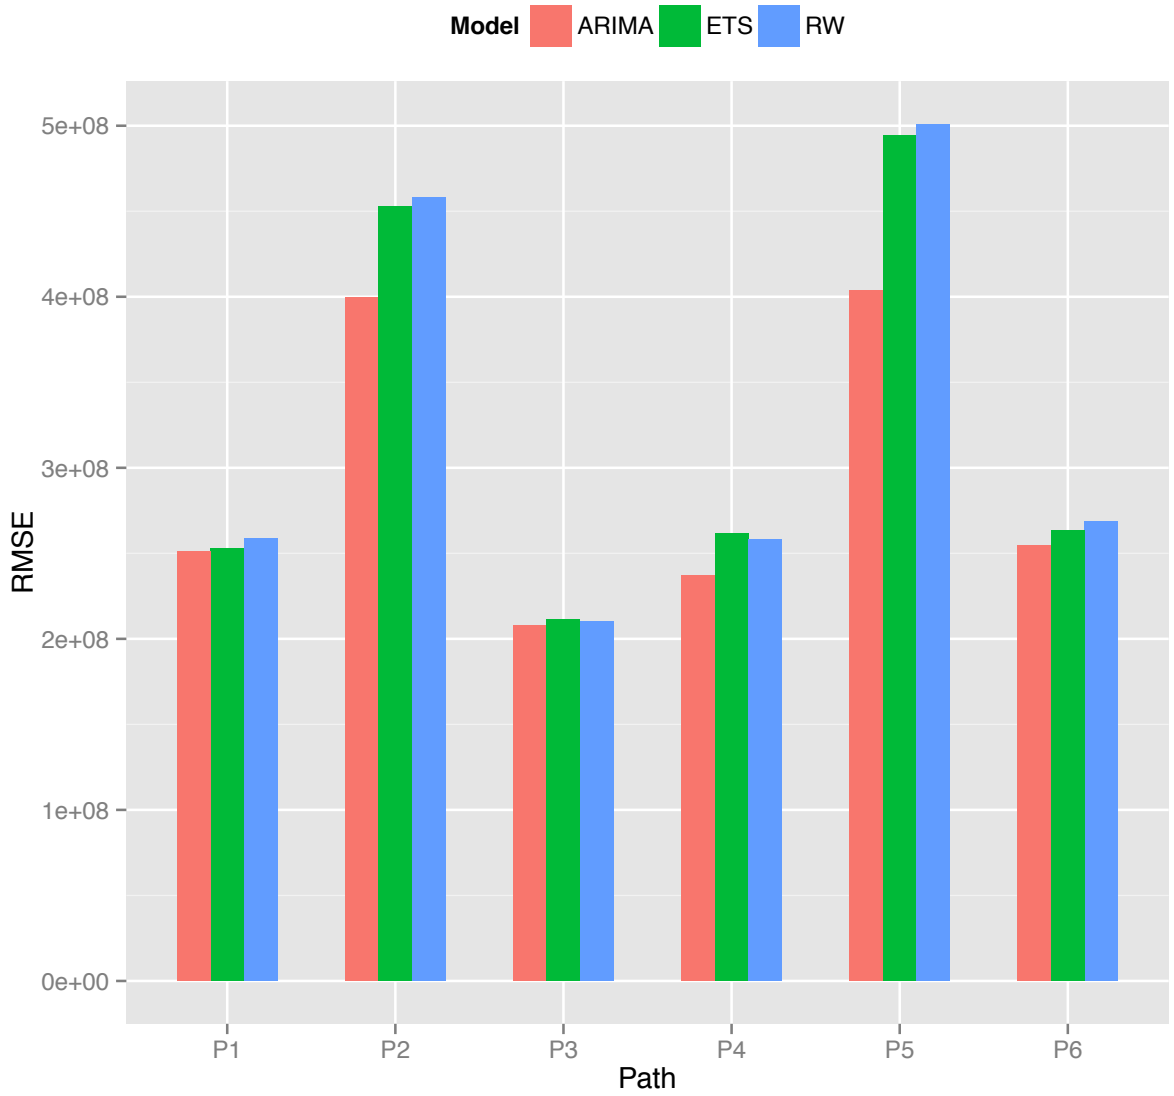

Logit Transformation

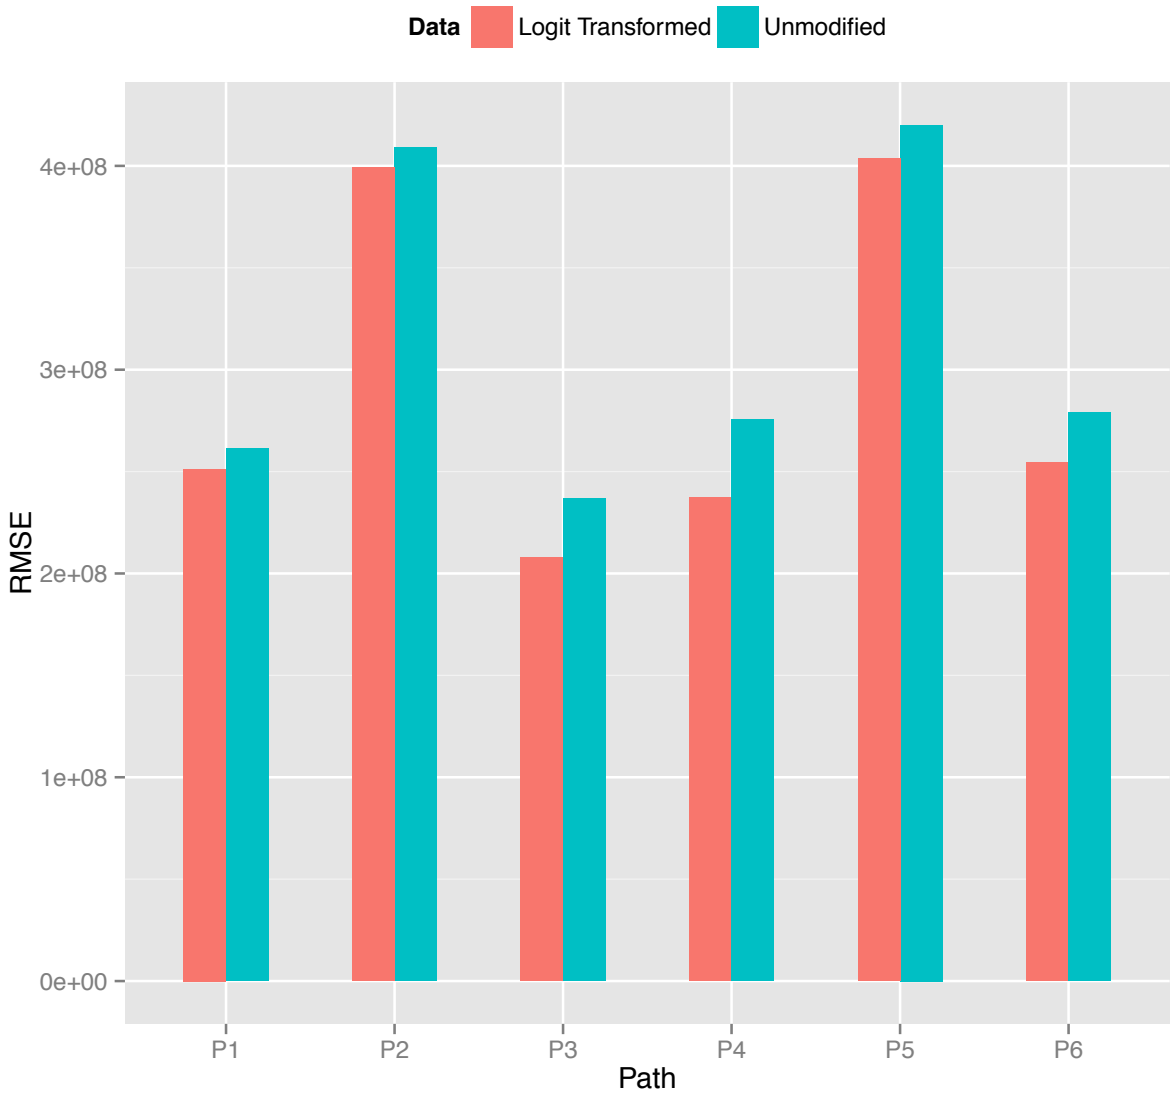

Training Set Size

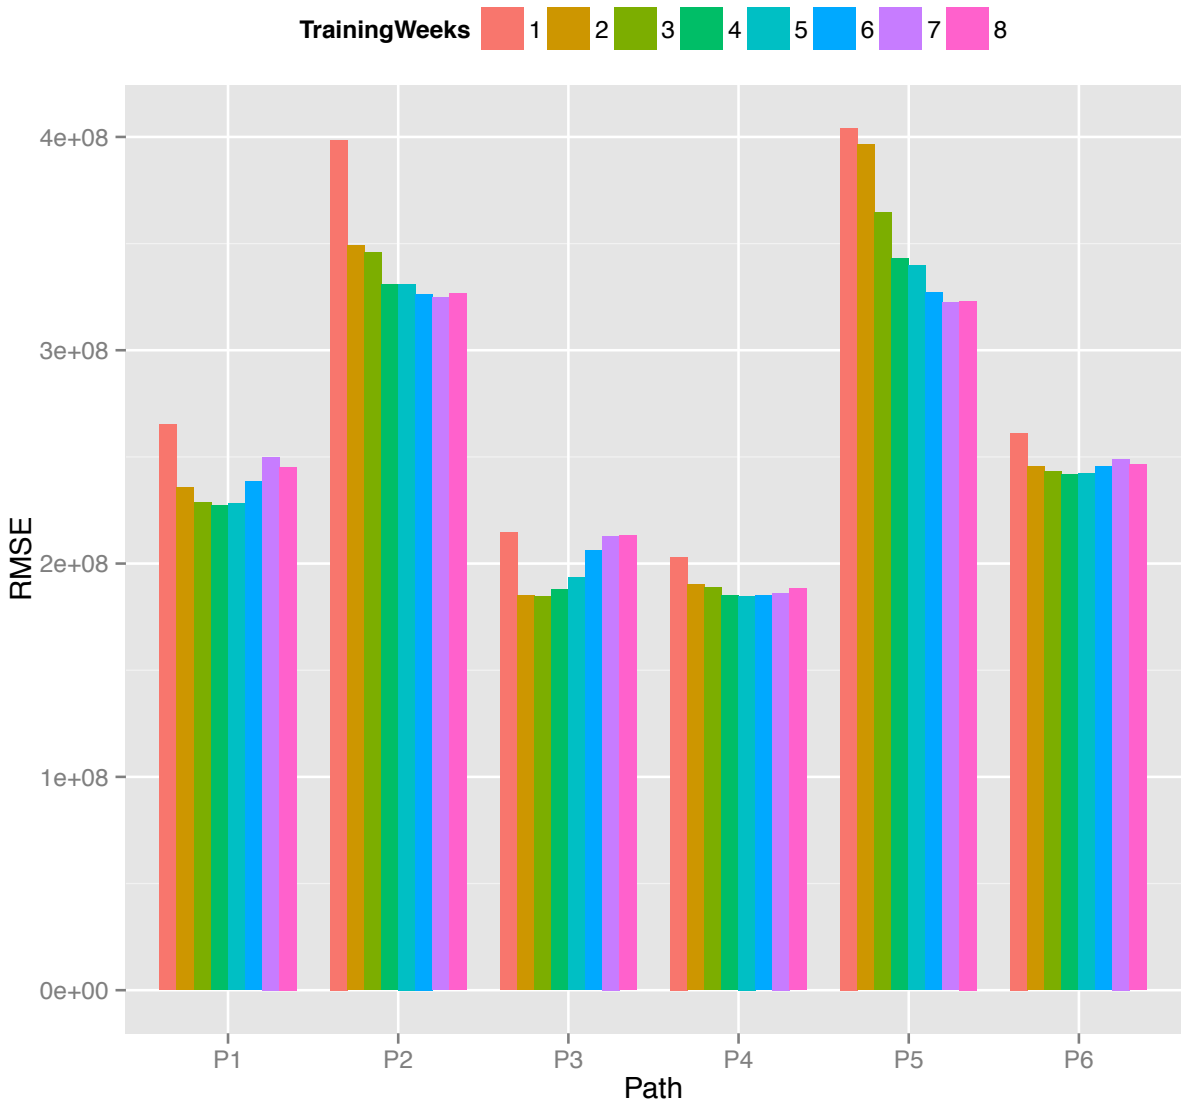

Stationarity

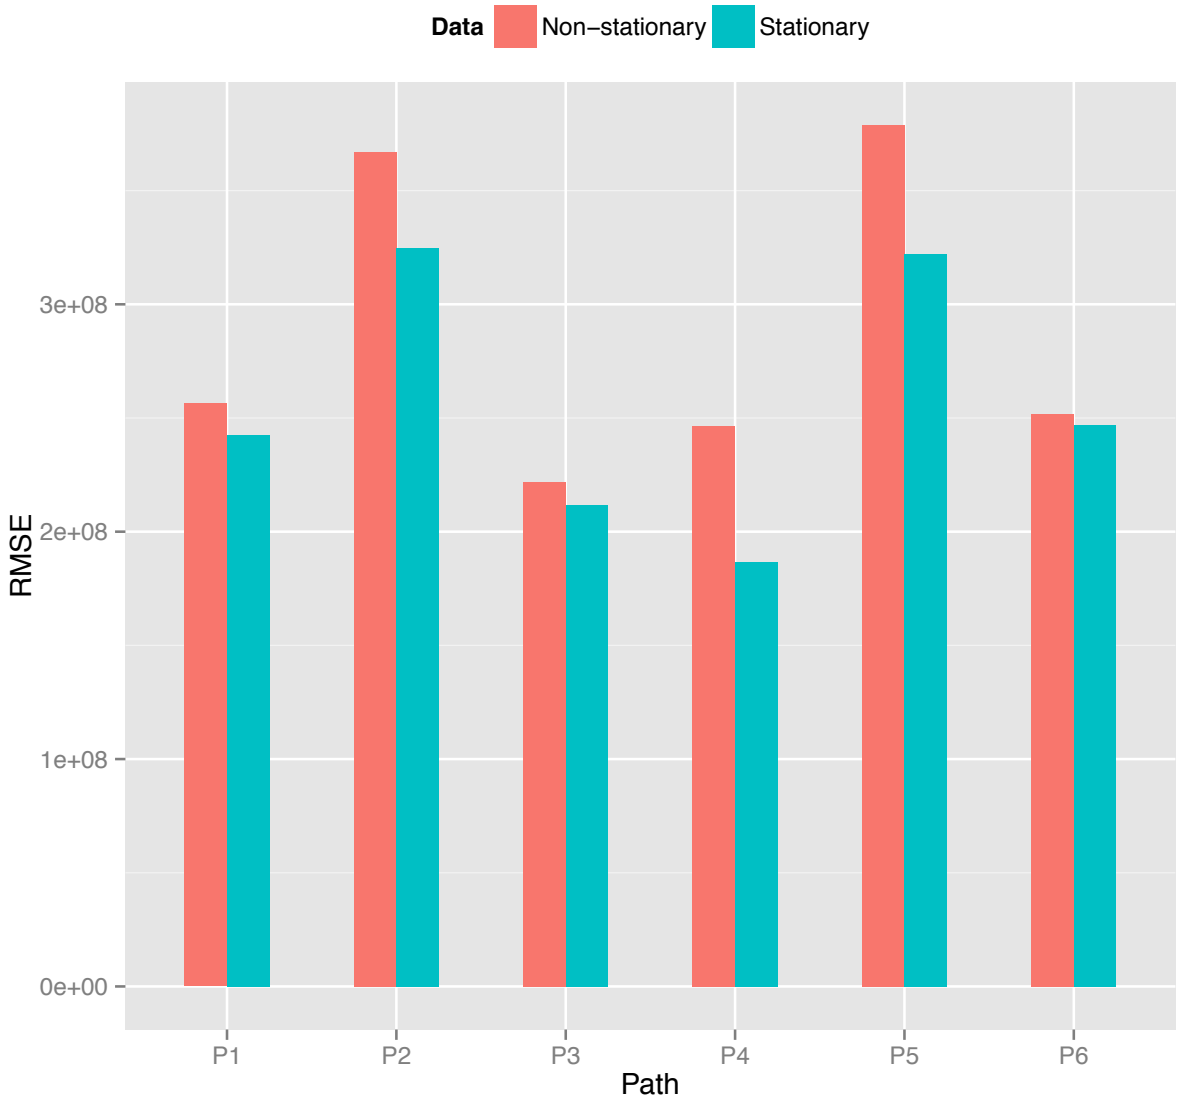

Forecast Results

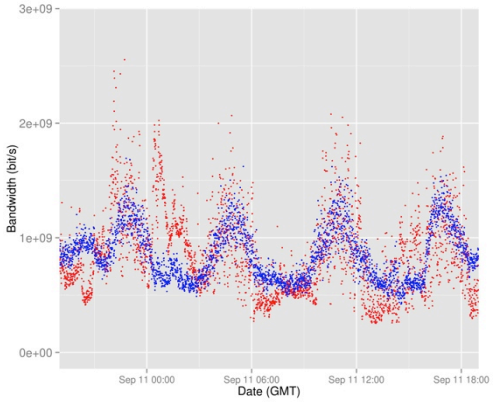

NERSC ➔ ANL

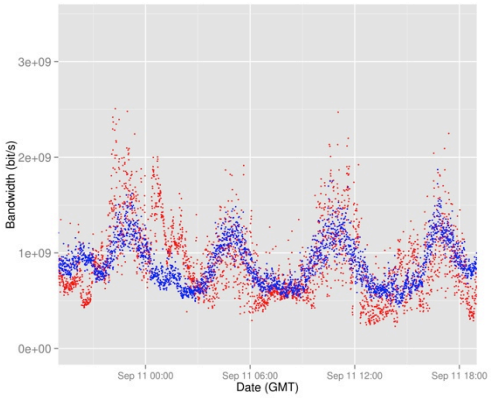

ANL ➔ NERSC

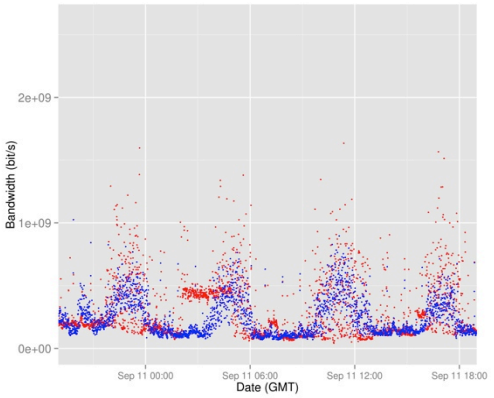

NERSC ➔ ORNL

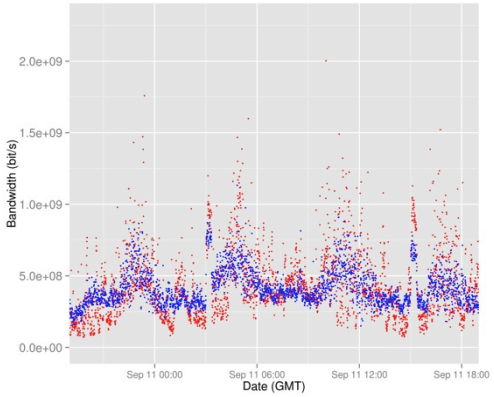

ORNL ➔ NERSC

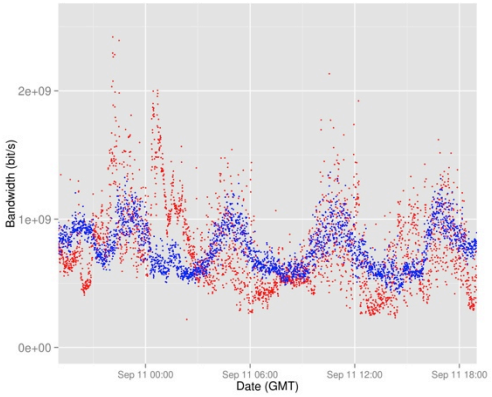

ANL ➔ ORNL

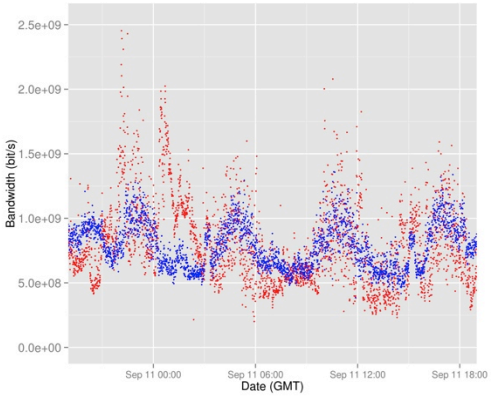

ORNL ➔ ANL

Historical Forecast Results

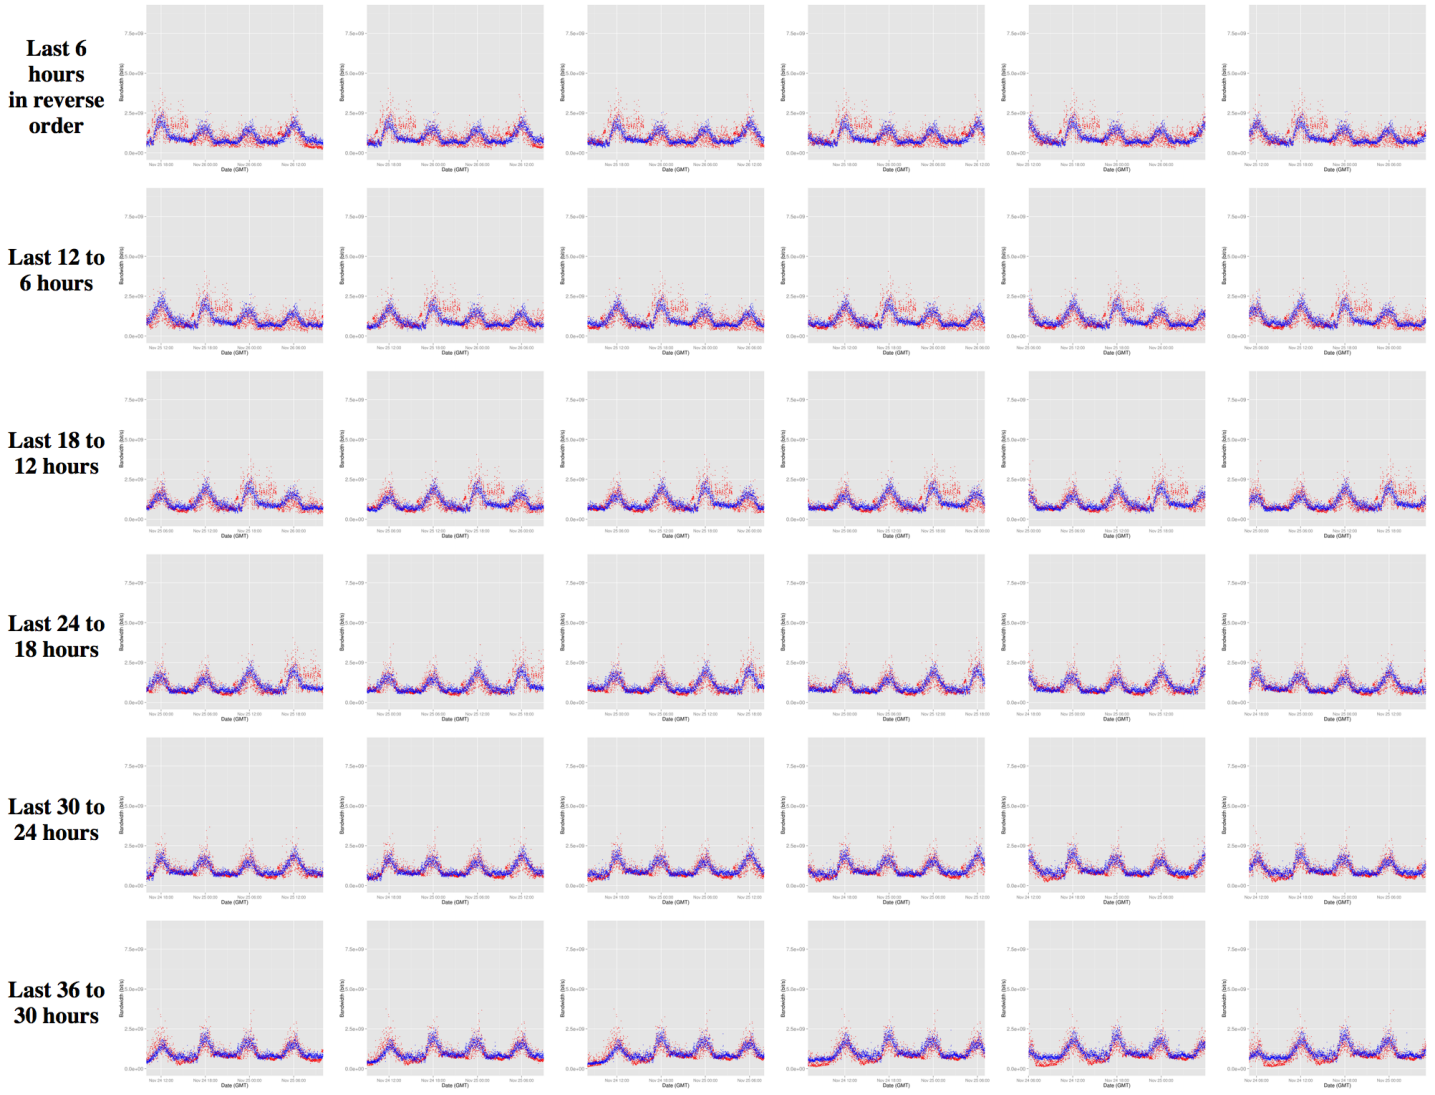

Forecast Results - RMSE

PID	SD_Train	SD_Test	RMSE
P1	4.13	2.36	2.27
P2	4.51	3.37	3.31
P3	4.01	2.07	1.88
P4	3.03	2.06	1.85
P5	4.64	3.40	3.42
P6	4.00	2.54	2.42

Conclusion

- Forecast Model
  - ARIMA with STL, logit transformation, and stationarity
  - Forecast errors were within the variances of observed data
  - Logit transform reduced prediction error by 8.5%
  - Stationarity assumption reduced prediction error by 10.9%
- Contact
  - Wucherl (William) Yoo, wyoo@lbl.gov

Backup - Future Work

- Adaptive Model
  - To adapt the long-term trend changes
- Multivariate Performance Prediction Model
  - To extend the analysis to multivariate data

Backup - Seasonal Adjustment

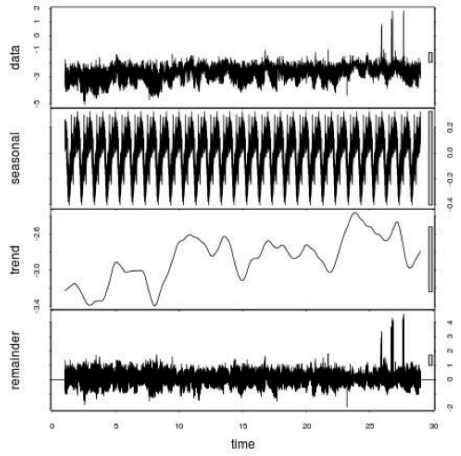

(a) P1

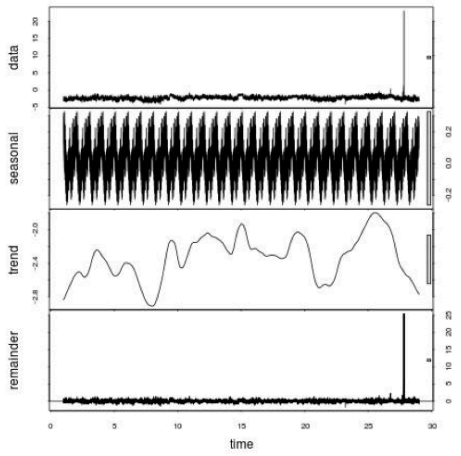

(b) P2

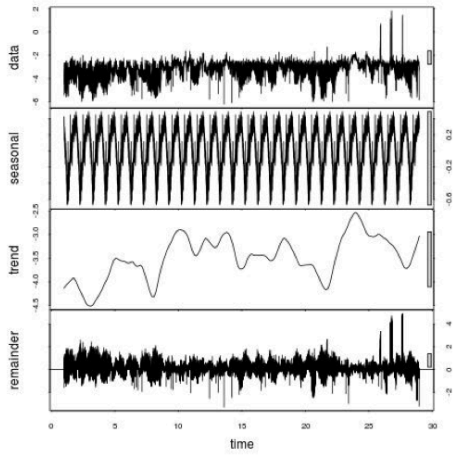

(c) P3

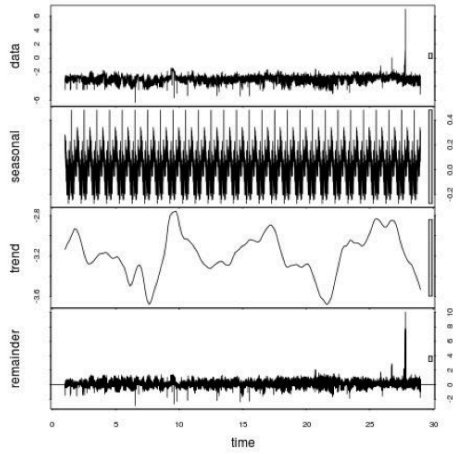

(d) P4

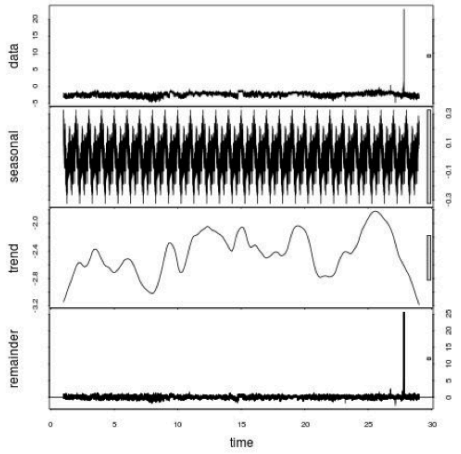

(e) P5

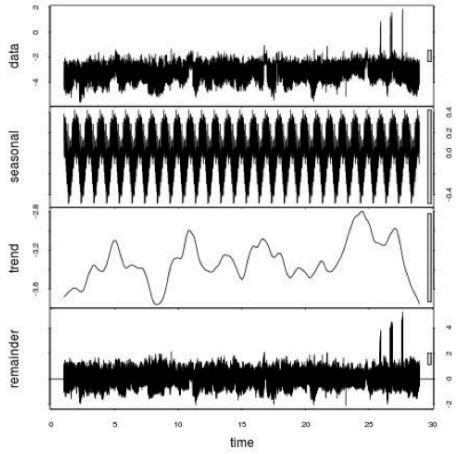

(f) P6