# Advanced Performance Modeling
## with
## Combined Passive and Active Monitoring

Alex Sim, Jaesik Choi, Kejia Hu, **SDM, CRD, LBNL**
Constantine Dovrolis, Demetris Antoniades, **Georgia Tech**

**apm@hpcrd.lbl.gov**

# Advanced Performance Modeling

## Project Goals

- Develop performance estimation models and software tools for high-bandwidth networks.
- Develop a performance prediction tool data throughput, for a given time window.

Use of SNMP counters to infer traffic transfers magnitude and path

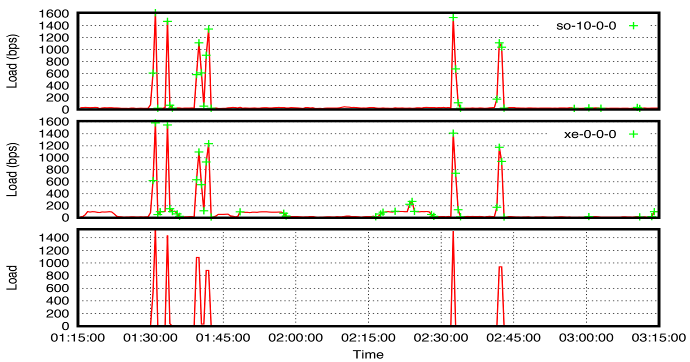


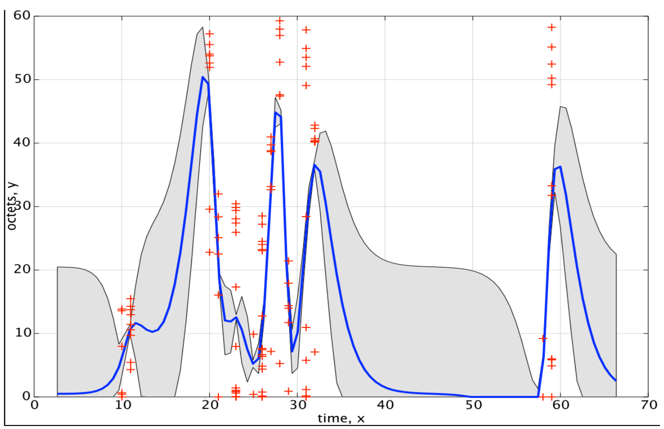


Nonparametric Bayesian models to infer a model size/complexity from the data automatically.

## Current Accomplishments

- Developed overall performance inference and prediction framework for this project.
- Inference of end-to-end network traffic
  * Enabling prediction, tracing and quantifying the network traffic with partial observations
  * Poster at PAM'13 (this week)
  * A paper in preparation for IMC'13
- Prediction models
  * Seasonal changes adjustment: decomposing and quantifying the network traffic
  * Improved accuracy of prediction by linear models and non-linear models
  * TIP2013 talk
  * A paper submitted to MLDM'13
  * A paper in preparation for SC'13

## Impacts

- Enable scientific collaborations to utilize the resources offered by high-bandwidth network infrastructures more effectively.
  * Improve network usage and enable predictable data throughput
  * Long-term capacity and traffic engineering planning of network infrastructures.

# Conceptual web page images

## NERSC/PDSF->BNL

En Español

**7.5 Gbps**
**throughput**

| | |
|---|---|
| Total Bandwidth | 20Gbps |
| Availability | 67.5% |
| Predictability | 92% |

Last Update on 19 Mar 9:18 am PDT

Current conditions at
**EW0993 Oakland (E0993)**
Lat: 37.86033 Lon: -122.23483 Elev: 960ft.

More Local Wx | 3 Day History | Mobile Weather

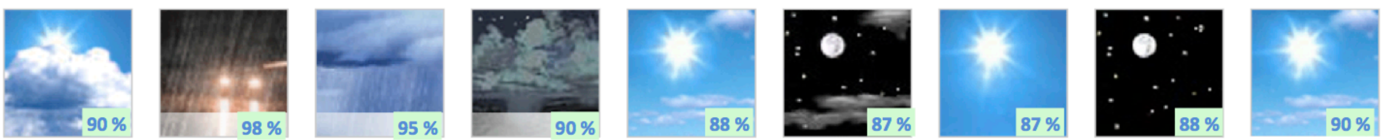


| TODAY | TONIGHT | WEDNESDAY | WEDNESDAY NIGHT | THURSDAY | THURSDAY NIGHT | FRIDAY | FRIDAY NIGHT | SATURDAY |
|---|---|---|---|---|---|---|---|---|
| Mostly Available | Busy Mostly | Mostly Available | Slightly Busy | Mostly Available | Mostly Available | Mostly Available | Mostly Available | Mostly Available |
| High: 5 Gbps | High: 18 Gbps | High: 5 Gbps | High: 13.5 Gbps | High: 4 Gbps | High: 4 Gbps | High: 3.5 Gbps | High: 4 Gbps | High: 4 Gbps |

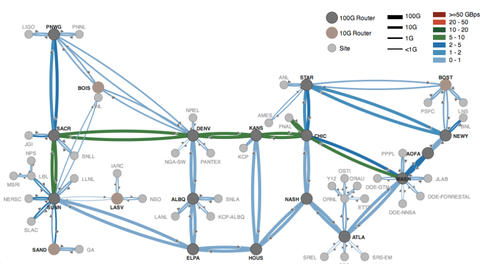

**Link Forecast:** NERSC/PDSF->BNL
37.86N 122.26W (Elev. 190 ft)

Last Update: 3:01 am PDT Mar 19, 2013

**Hourly Weather Forecast Graph**

[dashes/dots] | [b/w] | [hide menu]

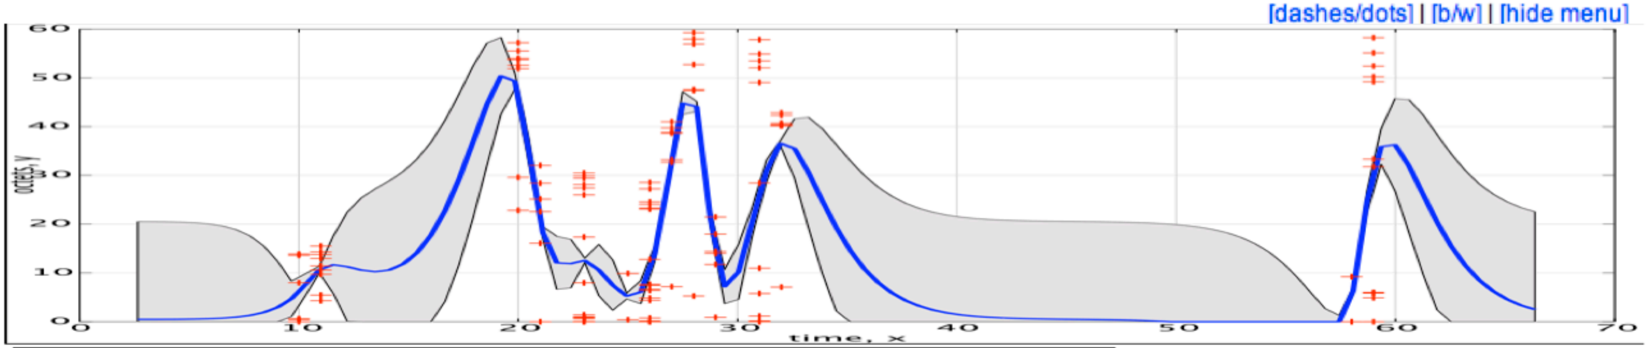

# Inference

**Demetris Antoniades**
**Georgia Tech**

# What SNMP data can tell us about edge-to-edge network performance

- **Need for historical transfer records as input to a TCP throughput prediction method**
- **NetFlow data**
  - Limited availability
  - Extensive sampling
  - Major user privacy concerns

- **Simple Network Management Protocol (SNMP)**
  - Widely used to provide aggregated link usage data from network components
  - Valuable source of information for network administrators

# This work

- **Propose a method that uses SNMP data to:**
  - **Identify network transfers** by observing variations in the aggregated throughput
  - Follow these transfers through the network and **identify source and destination routers}**

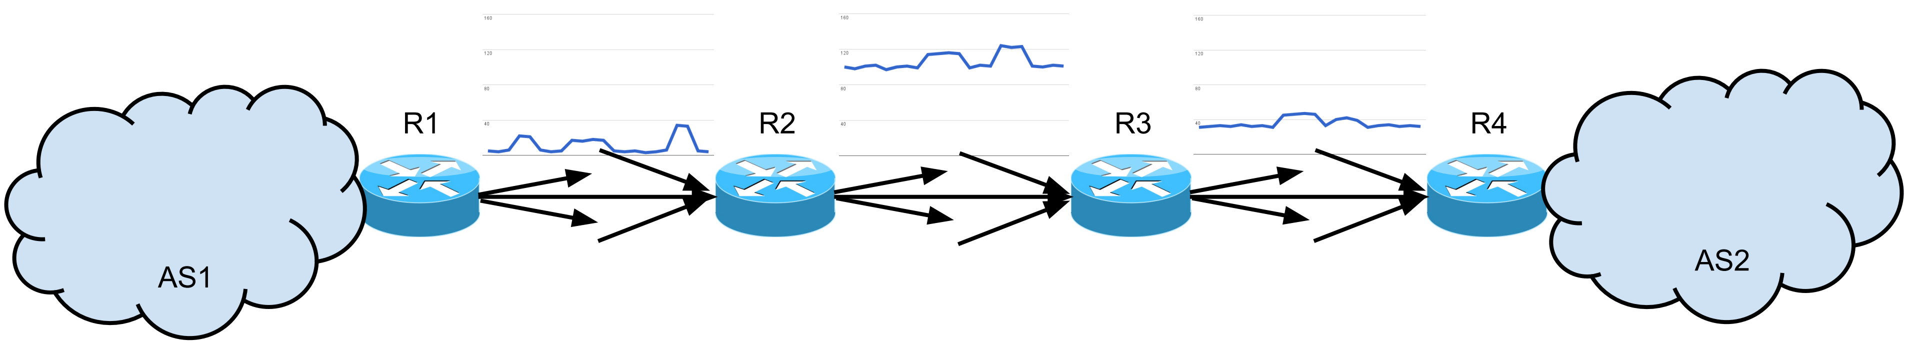

# Two main observations

- Deviations in link throughput show beginning and ending of network transfers

- Transfer path can be inferred by matching deviations from an incoming link to an outgoing link of the same router

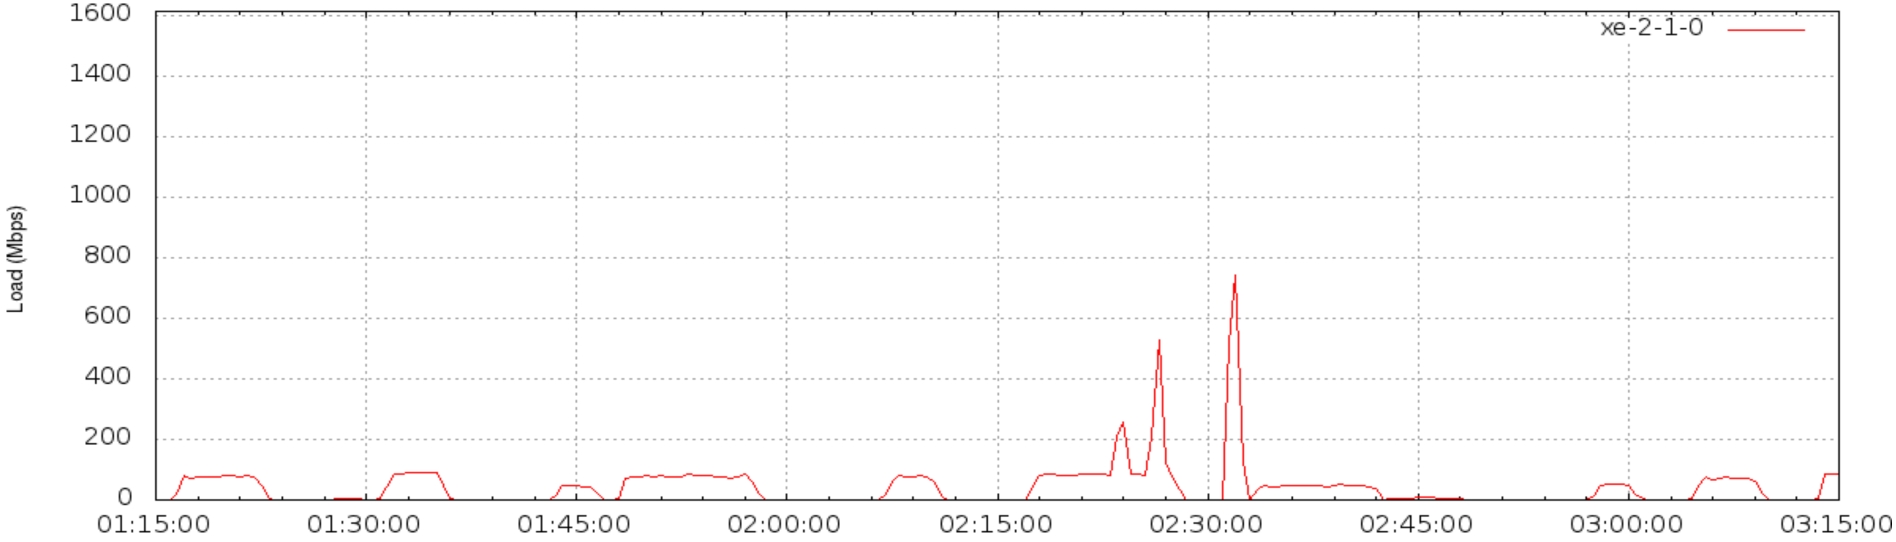
1600
1400
1200
1000
800
600
400
200
0
Load (Mbps)
xe-2-1-0
01:15:00
01:30:00
01:45:00
02:00:00
02:15:00
02:30:00
02:45:00
03:00:00
03:15:00

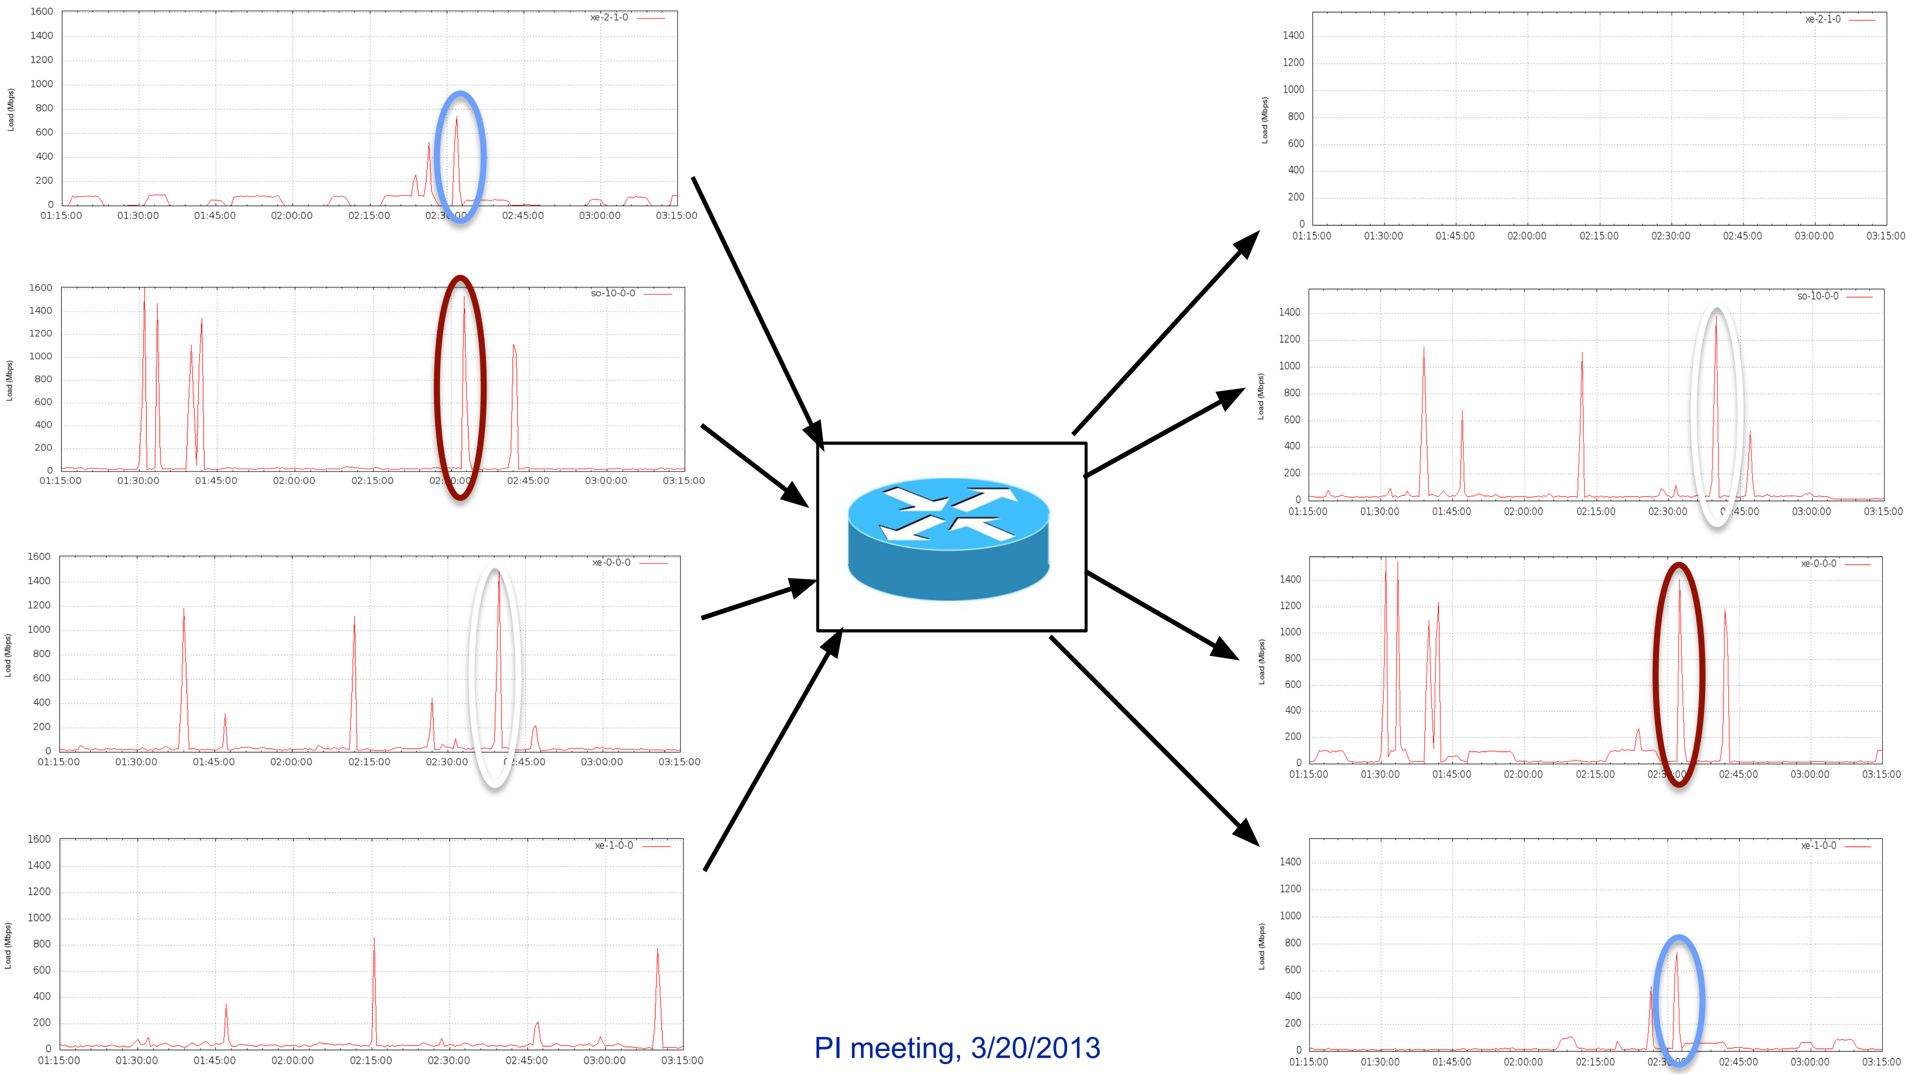

xe-2-1-0
so-10-0-0
xe-0-0-0
xe-1-0-0
Load (Mbps)
01:15:00 01:30:00 01:45:00 02:00:00 02:15:00 02:30:00 02:45:00 03:00:00 03:15:00

# Identify transfer starting and ending points

- Deviations in link's throughput can be considered **outliers** from link's normal behavior

- Outlier detection method should be:
  - **Robust** to the link's **variability**
  - **Robust** to any **periodicity** in the time-series
  - Does **not assume** any **predefined distribution** over the time-series
  - Able to detect outliers **online** as data becomes available

# MAD: Median absolute deviation from the median

- Using a moving window $V_n$, $n = 1 \ldots N$ of size $N$
  - Calculate absolute difference from the $median(V_N)$ for each value in $V_n$
  - MAD equals the median of these absolute differences
- Value $V_i$ is an outlier if
  - $V_i - median(V_N) \geq c * MAD$

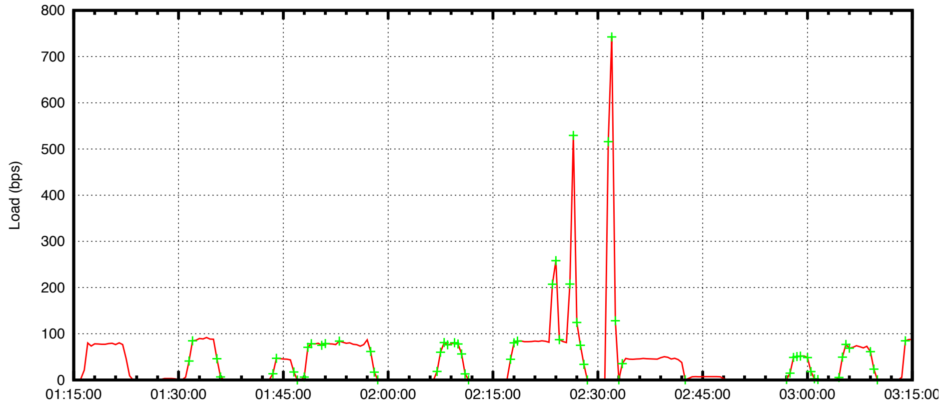

# Transfer route inference

- Identify outgoing interface *O* for each incoming interface *I* outlier $Outlier_I(t)$
  - Find outgoing interfaces $O_1$, $O_2$, ..., $O_k$ with traffic deviations $V_{O1}(t)$, $V_{O2}(t)$, ..., $V_{Ok}(t)$$ in range

$$Outlier_I(t) \pm D\%$$

  - Select $min(V_{O1}(t), V_{O\_2}(t), ..., V_{Ok}(t))$ as the outgoing interface *O*
- Iterate the procedure through all subsequent routers

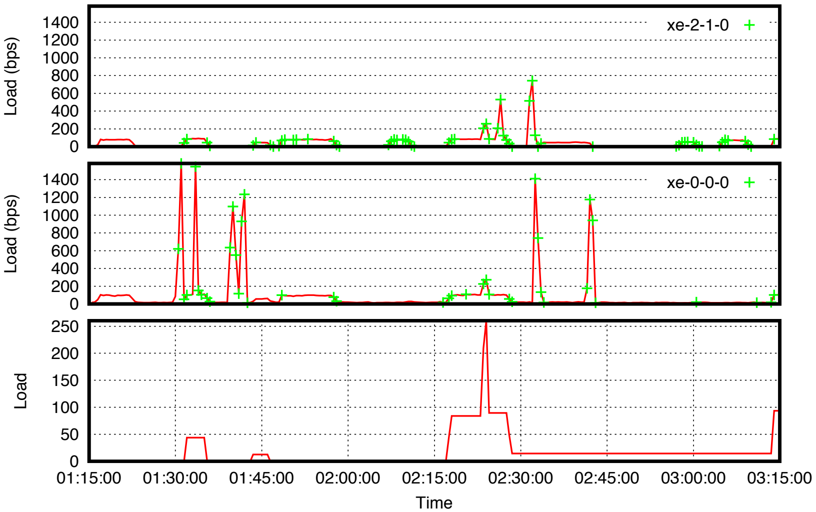
xe-2-1-0
xe-0-0-0
Load (bps)
Load
Time
01:15:00
01:30:00
01:45:00
02:00:00
02:15:00
02:30:00
02:45:00
03:00:00
03:15:00

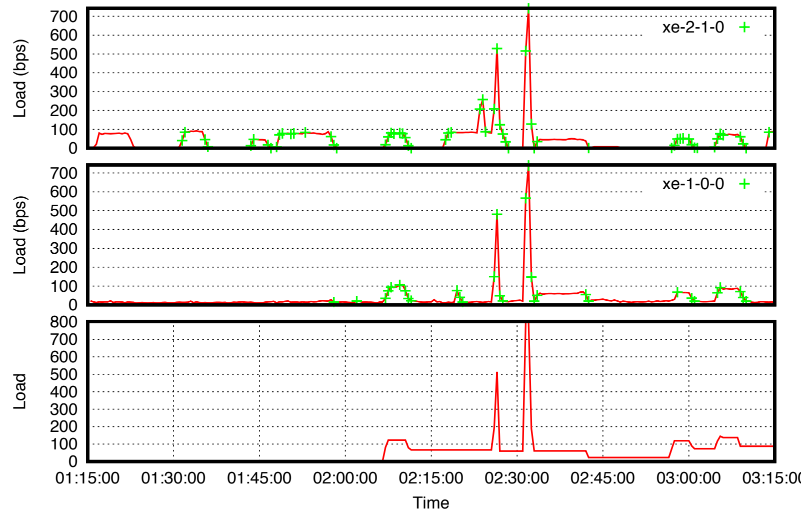
xe-2-1-0
xe-1-0-0
Load (bps)
Load
Time
01:15:00
01:30:00
01:45:00
02:00:00
02:15:00
02:30:00
02:45:00
03:00:00
03:15:00

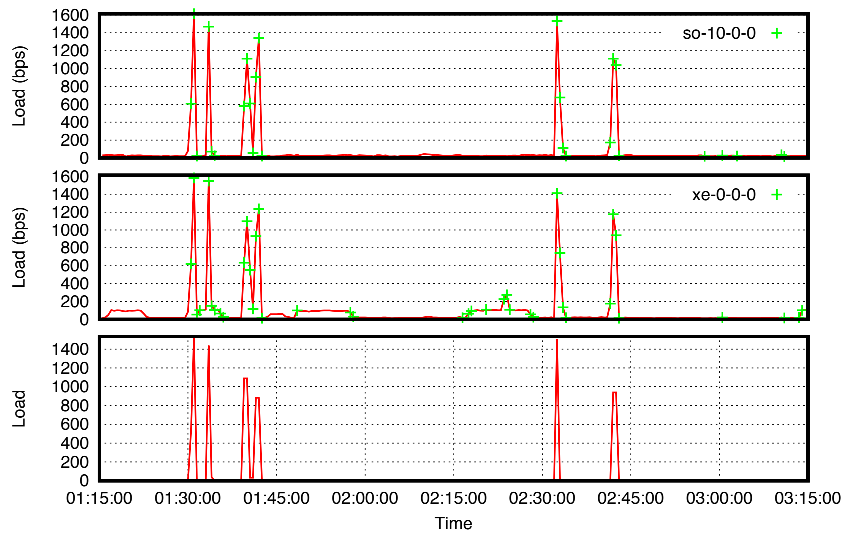
so-10-0-0
xe-0-0-0
Load (bps)
Load
Time
01:15:00
01:30:00
01:45:00
02:00:00
02:15:00
02:30:00
02:45:00
03:00:00
03:15:00

# Preliminary Evaluation

- Using a number of self-generated transfers
- Examine
  - Throughput of inferred transfers
  - Duration of inferred transfers

# Inferred transfer throughput

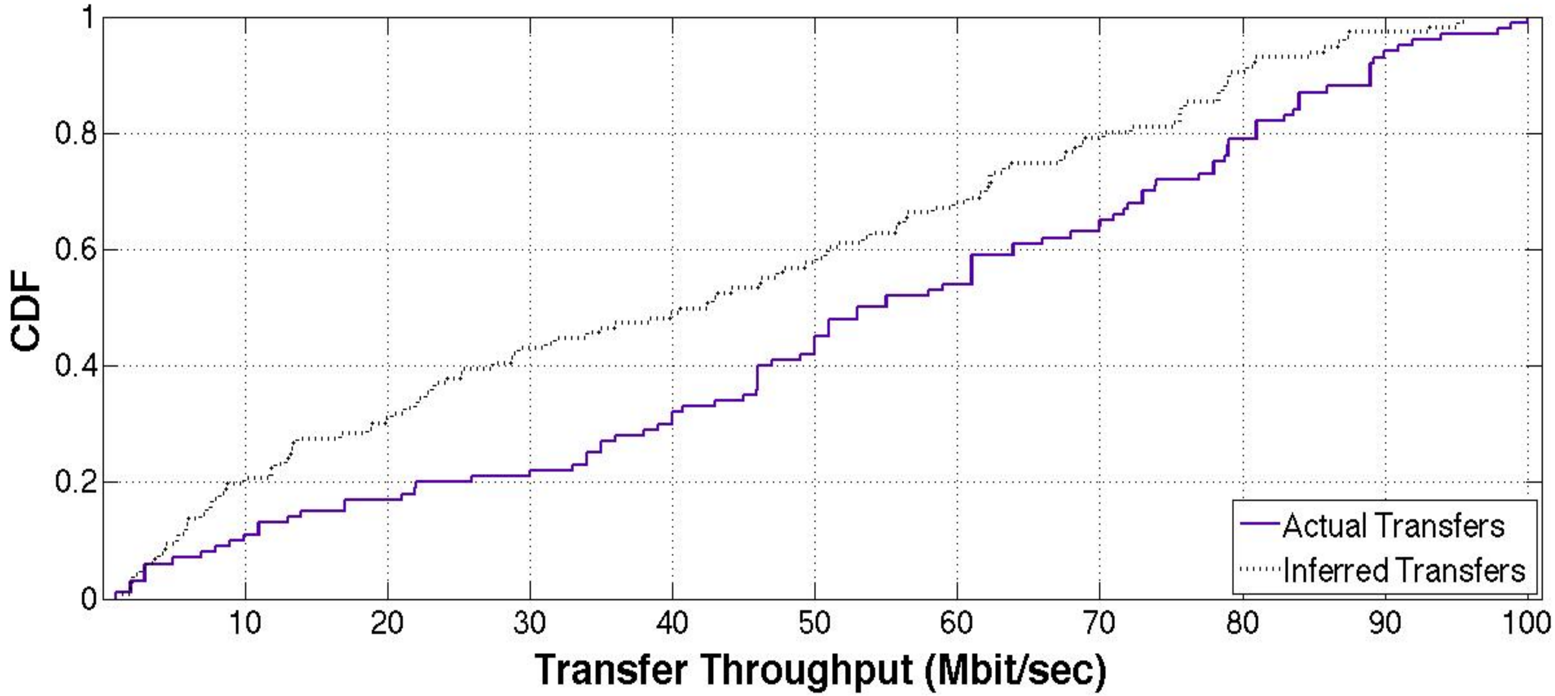

# Inferred transfer duration

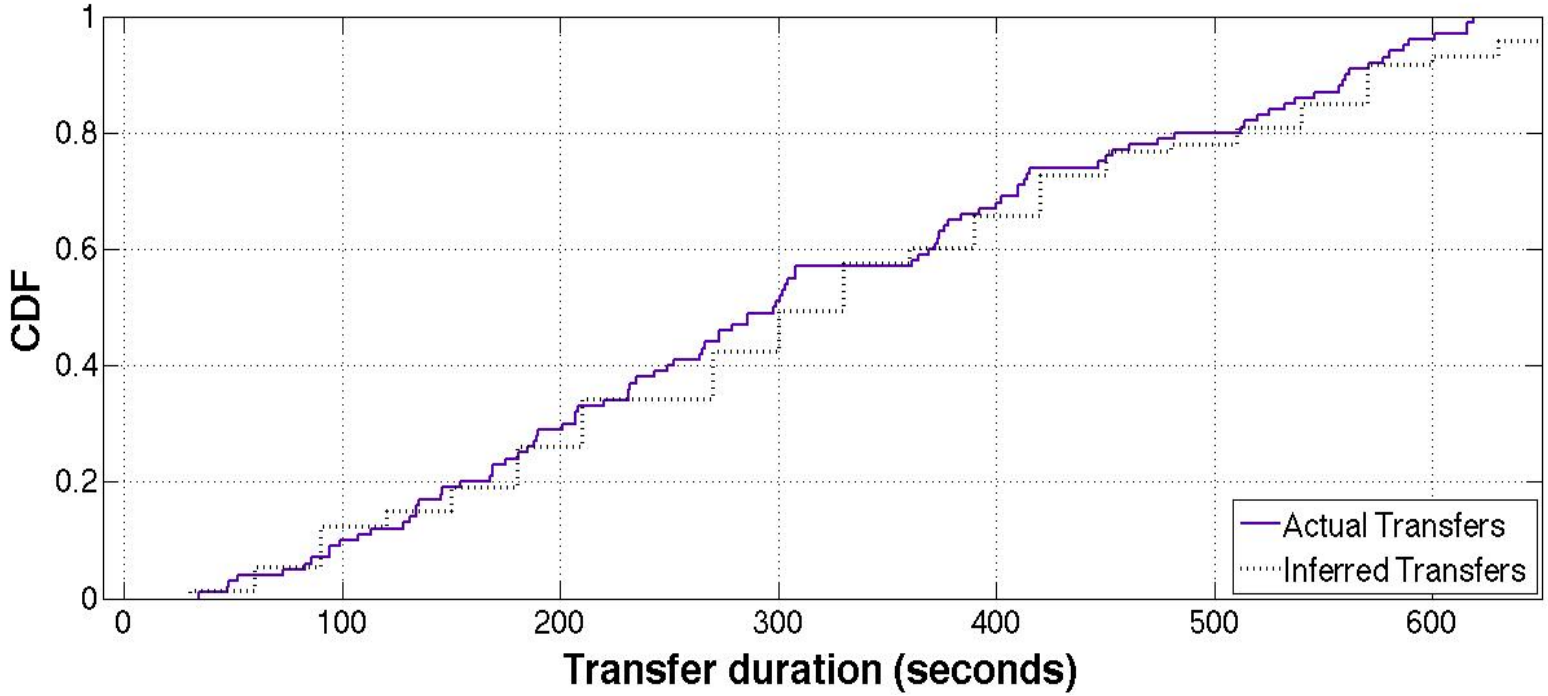

# Ongoing/Future Work

- **Further evaluation of our method**
  - **Using NetFlow data from real transfers**
  - **Multipath transfers: how to identify transfer splits over different outgoing interfaces**
- **TCP throughput prediction**
  - **Use inferred transfers to assist prediction**
    - **In the absence of in addition of NetFlow transfers**

# Statistical prediction models

**Jaesik Choi**
**SDM, CRD, LBNL**

# Prediction Models

- Statistical approach to the prediction models for network traffic performance based on two types of data
- SNMP → Time series model with Seasonal Adjustment
  - Analyzing network traffic patterns
  - By decomposing into seasonal, trend and random components
  - To enable prediction, tracing and quantifying the network traffic
- Netflow → Generalized Linear Mixed Model
  - Analyzing variation with the network conditions
  - By considering fixed effects, random effects and error term
  - To improve accuracy of prediction by involving both universal variance caused by randomness and variance by changes in the network traffic

# Seasonal Adjustment Performance Prediction

- Data: ESnet SNMP from May 2011 to June 2012.
- STL: A Seasonal-Trend Decomposition Procedure Based on Loess*.

Data = Seasonal + Trend + Residual

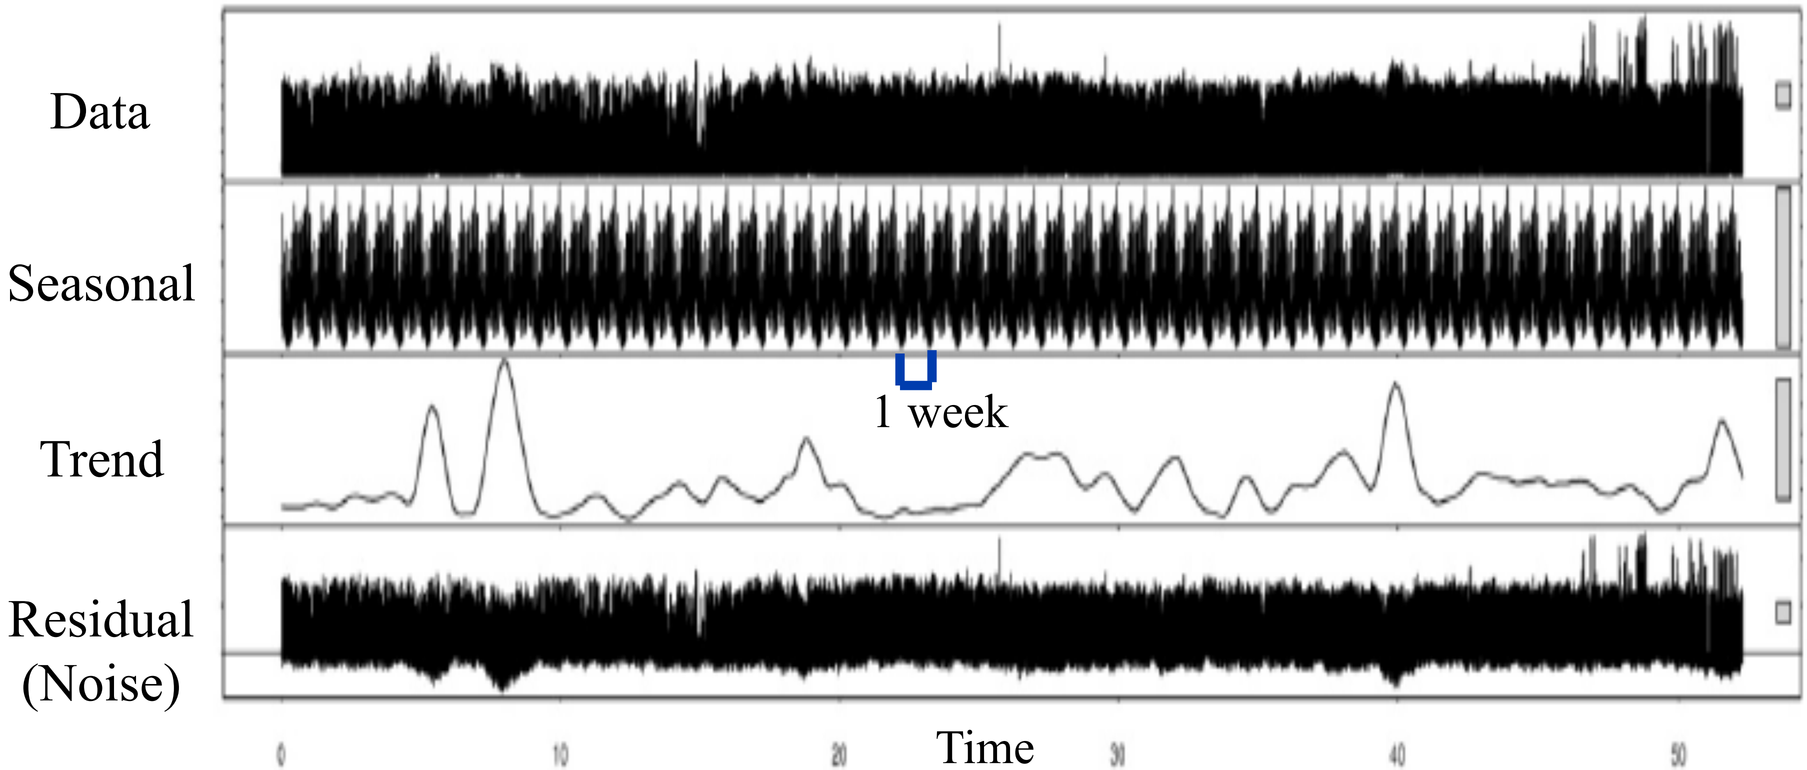


*Loess: locally estimated scatterplot smoothing

apm@hpcrd.lbl.gov

# Seasonal Adjustment Performance Prediction

Prediction (w/o seasonal trends) | Prediction (w/ seasonal trends)

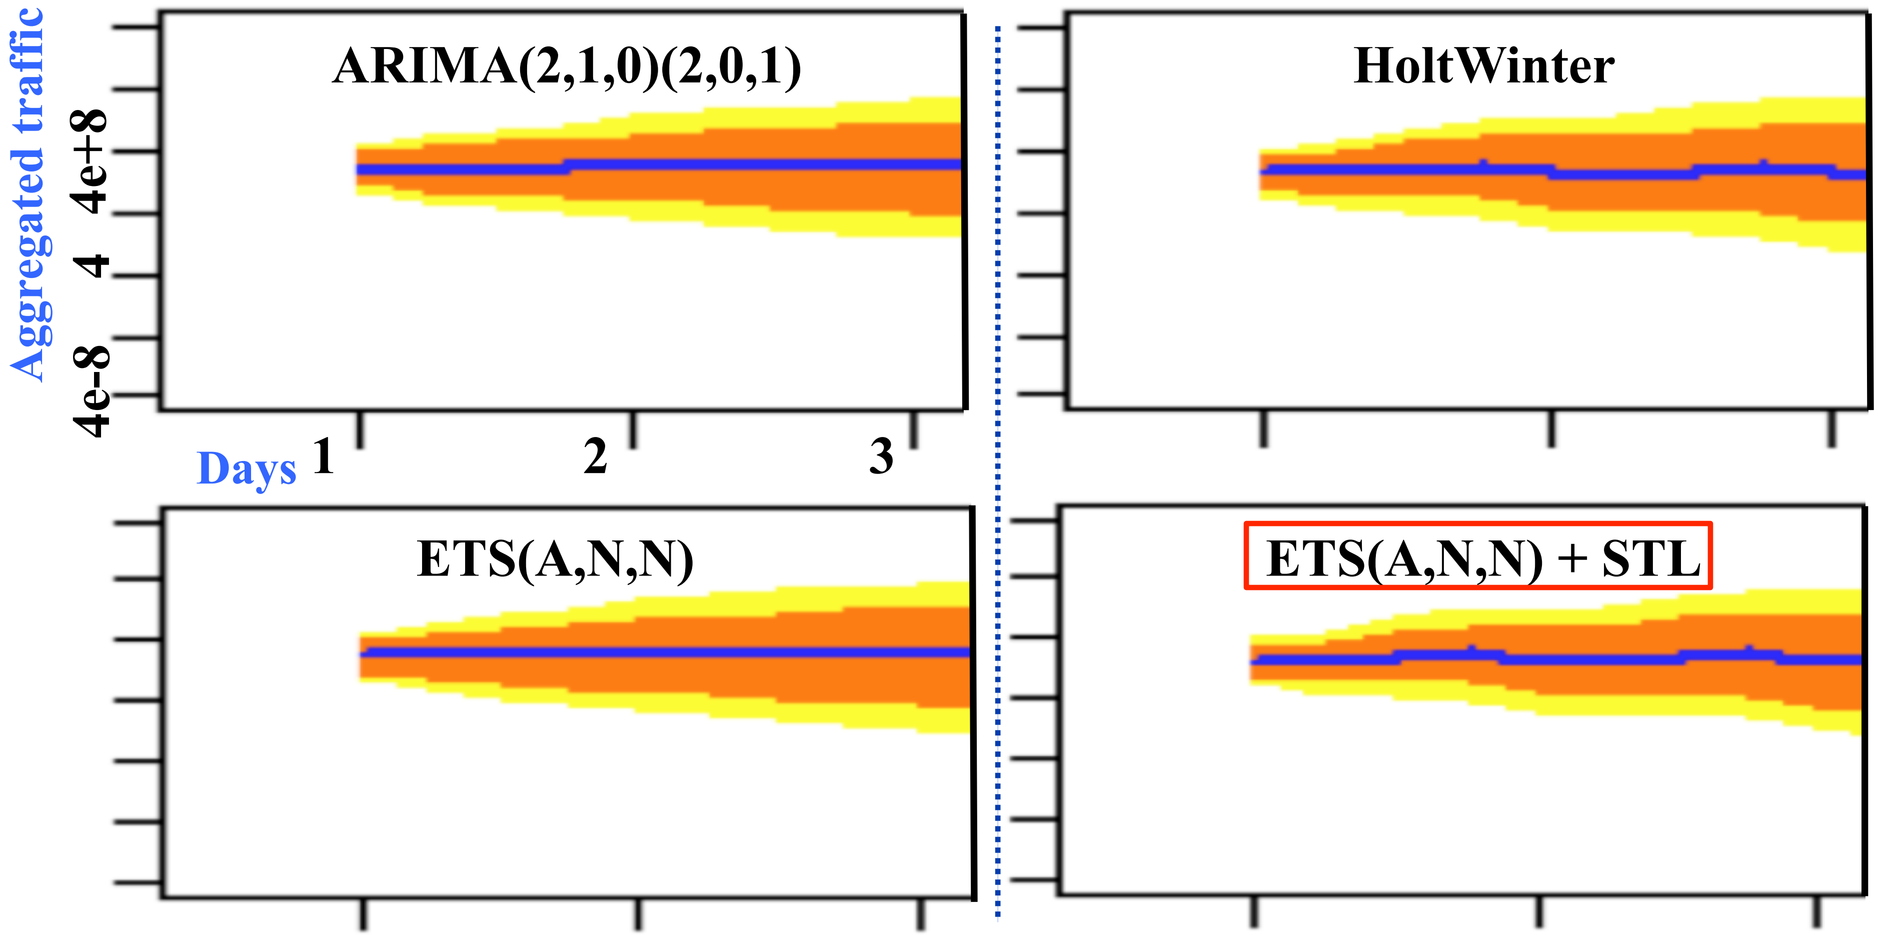


ARIMA: Autoregressive integrated moving average; ETS: Exponential smoothing state space model

# Generalized Linear Mixed Models

- **GLMM** : a set of predictors (linear Mixed Models) w/ shared coefficients and individual random effects.

$$Y = X\boxed{\beta} + Z\alpha + e$$

  - where X, Z are known matrics
  - Random effects $\alpha \sim N(0, G)$, $e \sim N(0, \Sigma)$
  - $\alpha$, e are uncorrelated

- **Estimation using LASSO (Maximize Log-likelihood)**

$$\beta_{est}(\lambda) = argmin\left|\left|y - Z\tilde{D}\tilde{\Gamma}\alpha - X\beta\right|\right|^2 + \lambda|\beta|$$

  - $G = D\Gamma\Gamma'D$ where D diagonal, $\Gamma$ lower triangular matrix; $\tilde{D}, \tilde{\Gamma}$ is kronecker product

- **Prediction (Minimized Mean Squared Prediction Error)**

$$\beta_{est}(\lambda) = argmin\left|\left|(y - X\beta)B\right|\right|^2 + \lambda|\beta|$$

  - $B = GZ'V^{-1}$ and $V = \Sigma + ZGZ'$

New method

LASSO: Least Absolute Shrinkage and Selection Operator

# Generalized Linear Mixed Models Performance Prediction

- Our new method (MMSPE) is better then MLE

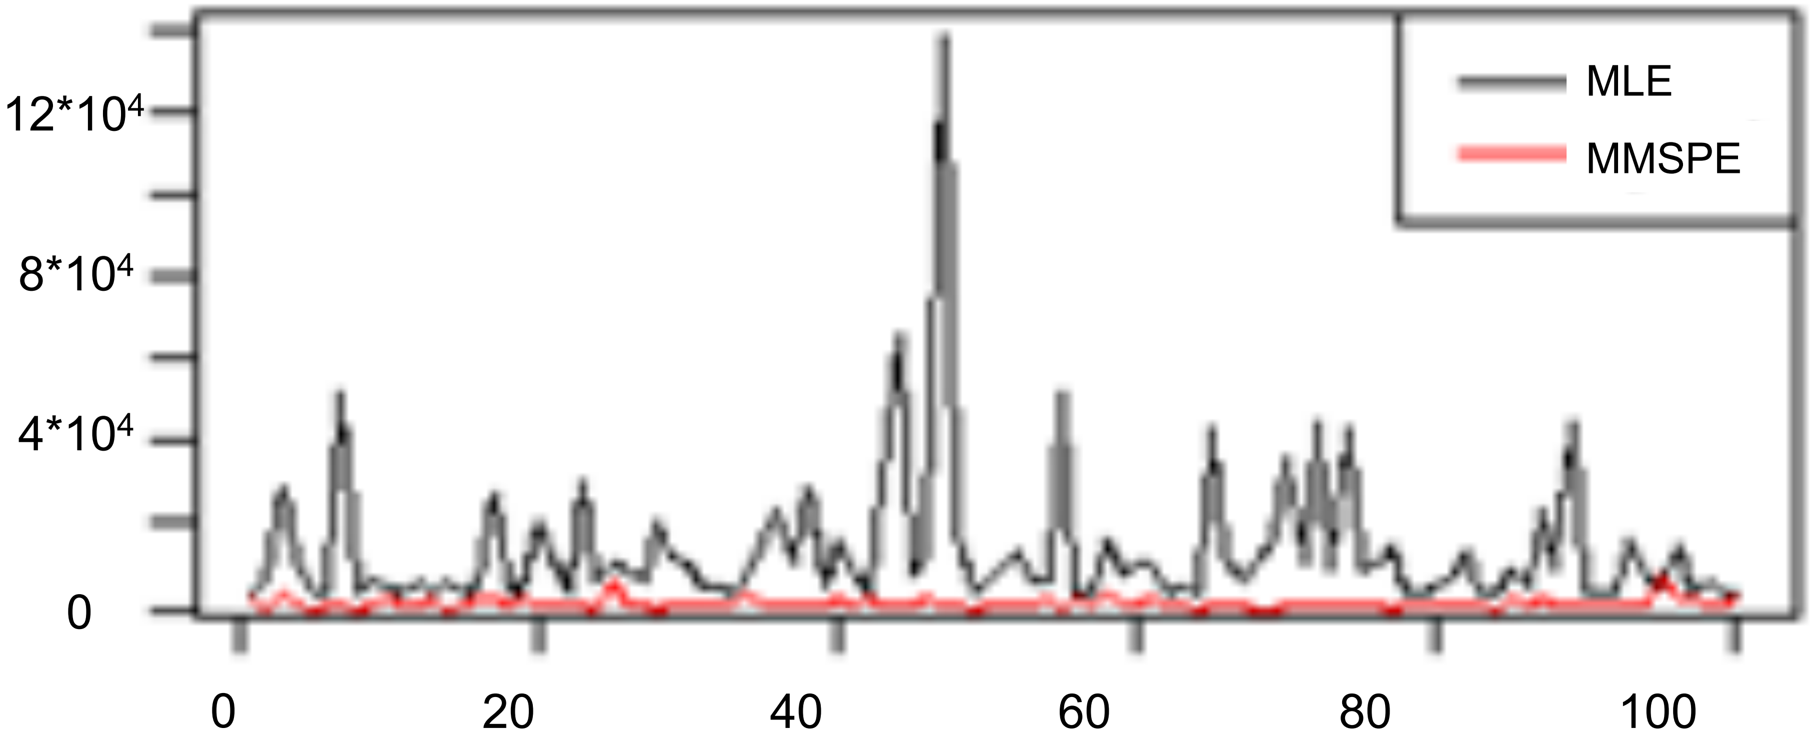

# Nonparametric Bayesian to Estimate Network Trends

Sampling from different Gaussian processes with different c $f \sim GP(\cdot|0,c)$

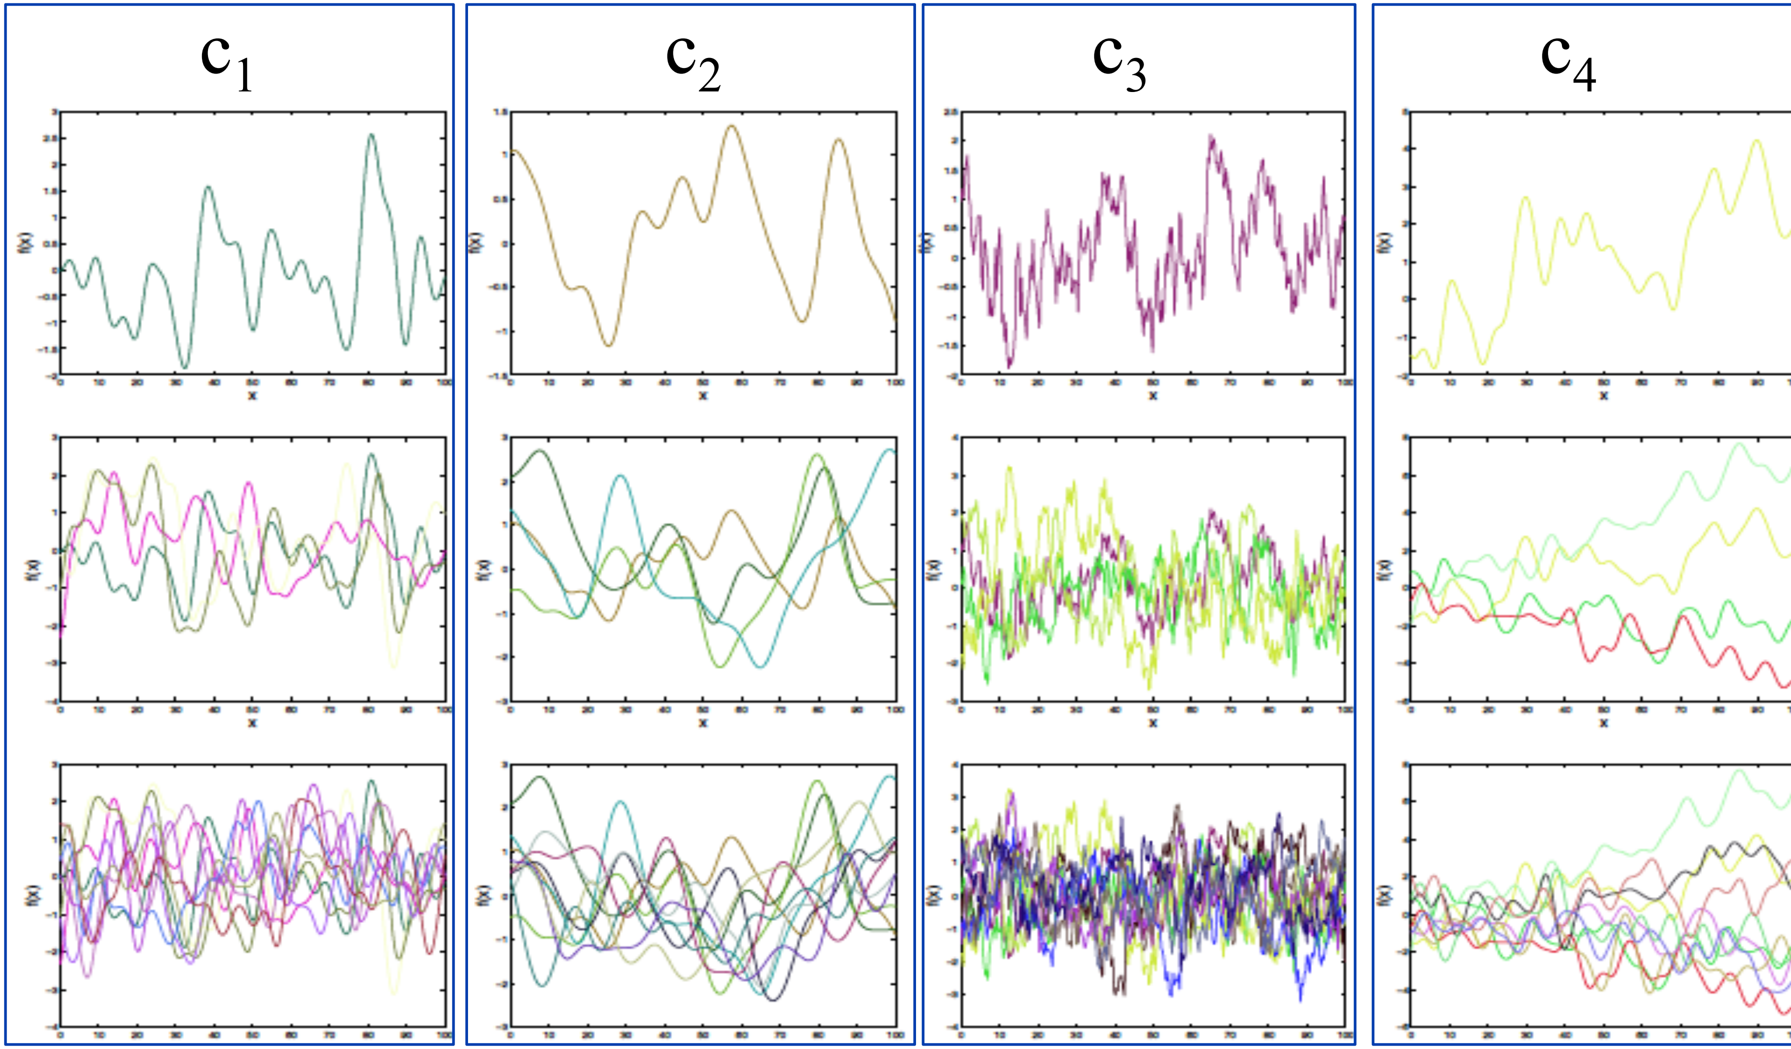




Courtesy of Zoubin Ghahramani

# Nonparametric Bayesian to Estimate Network Trends

- **Nonparametric Bayesian**
  - Flexible to infer an adequate model size/complexity from the data (e.g., no predefined # of components)
- **Nonlinear Regression with Gaussian Process**

**Model**

$y_i = f(x_i) + e_i$ (= function + noise)

$f \sim GP(\cdot \mid 0, c)$ (=function space)

$e_i \sim N(\cdot \mid 0, \sigma^2)$ (=noise)

**Prediction**

$$P(y' \mid x', D)$$

$$= \int df P(y' \mid x', f, D) P(f \mid D)$$

D: data

x': Points of interests

y': outputs

# Nonparametric Bayesian to Estimate Network Trends

- **Predict network bandwidth with Gaussian Process**

Input: Octets/sec from A to B

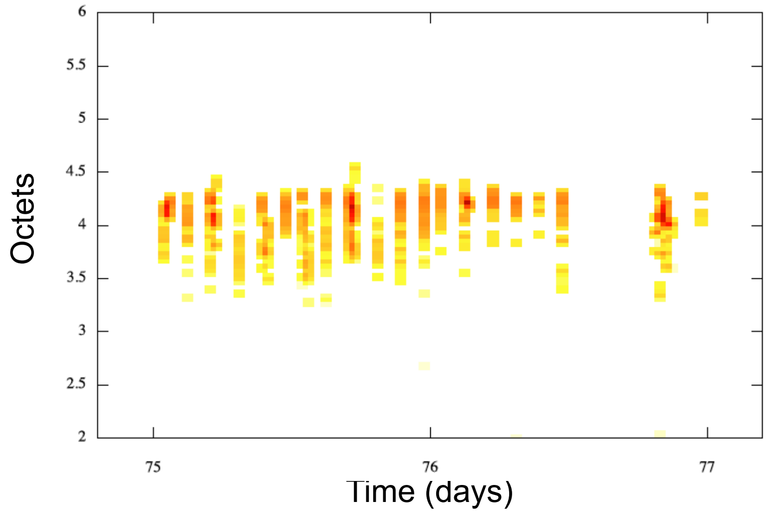


Output: Posterior traffic trends (functions)

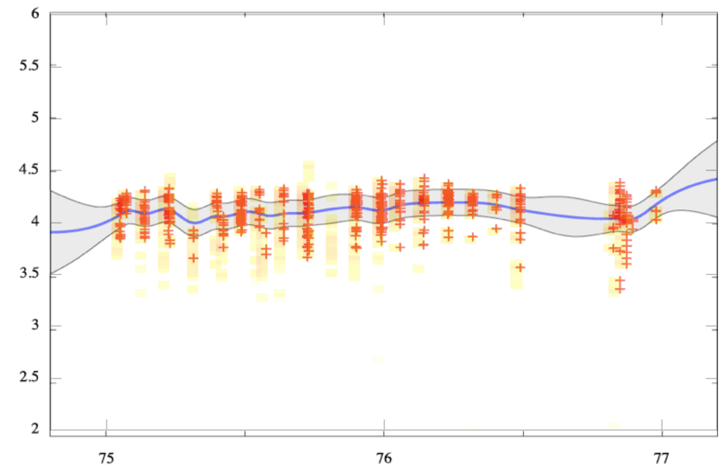


Other examples

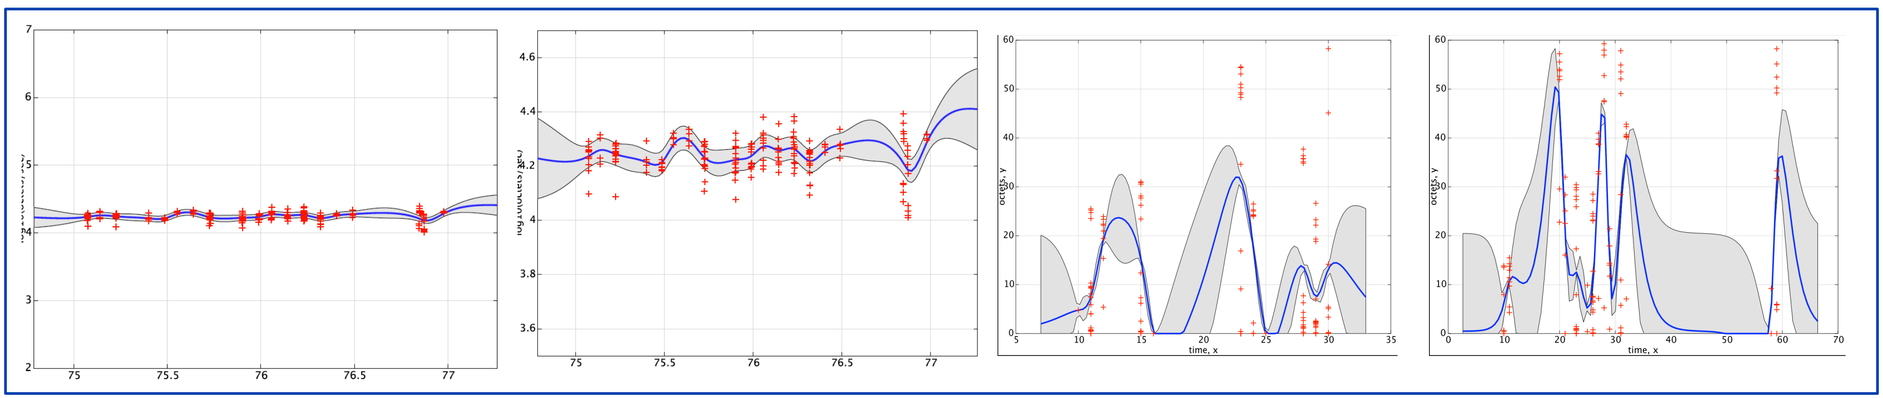


apm@hpcrd.lbl.gov

# Next steps

- **Inference**
  - Further evaluation of the method on NetFlow data
  - Investigate multipath transfer issue
  - Integrate inferred transfers with the performance prediction
- **Statistical prediction models**
  - Further investigation on linear and non-linear models
  - Study hybrid models, adaptive models
- **Integration with ESnet portal**
  - Collaboration with communities
- **Other research issues**
  - E.g. how much measurement data is needed for a "good" prediction
- **Questions: apm@hpcrd.lbl.gov**