Compressive Network Monitoring

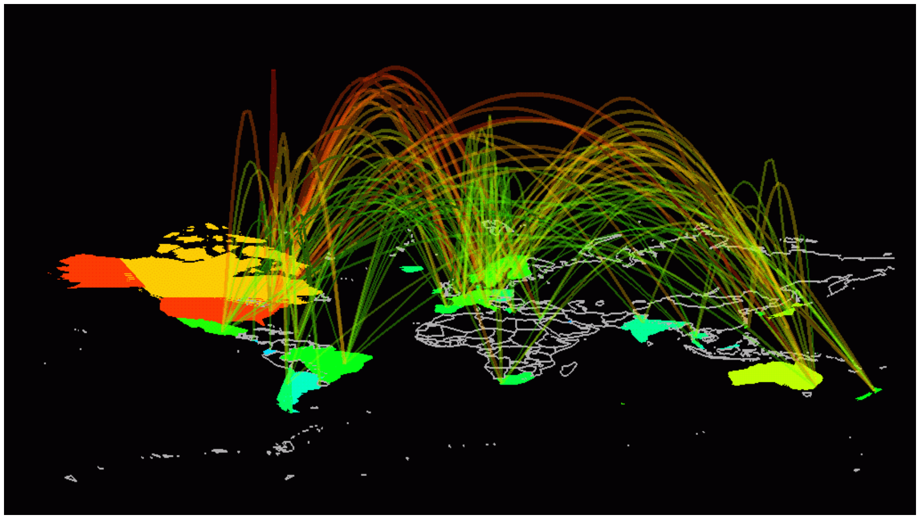

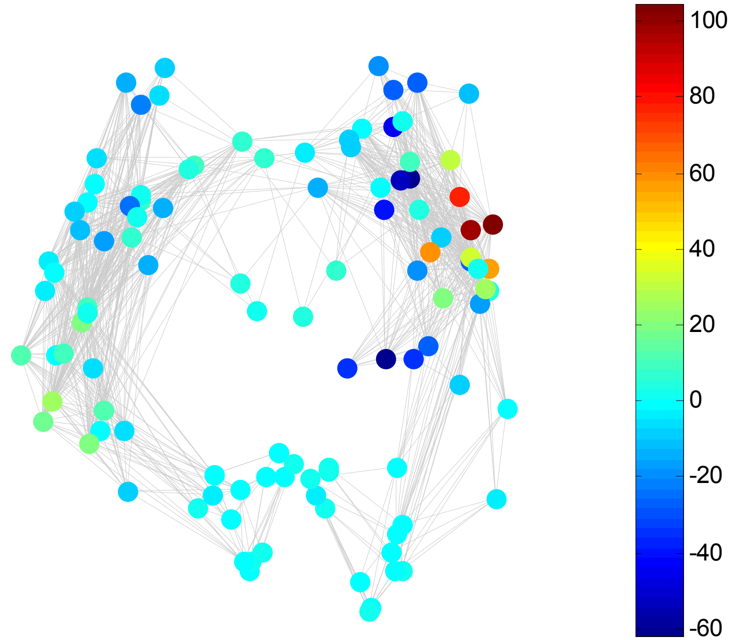

Mike Rabbat

Yvan Pointurier & Mark Coates

IPAM MultiRes Internet Analysis Workshop, 23 September 2008

Sponsored by the Natural Sciences and
Engineering Research Council of Canada

McGill

Path-Level Performance Monitoring

Motivating Application

- Voice-over-IP on an overlay network with redirectors
- Multiple paths to choose from
  - select paths with minimum mean delay or delay variance
- Send a small number of critical packets (vocal transitions) along multiple paths
- Use these packets to estimate the path delays

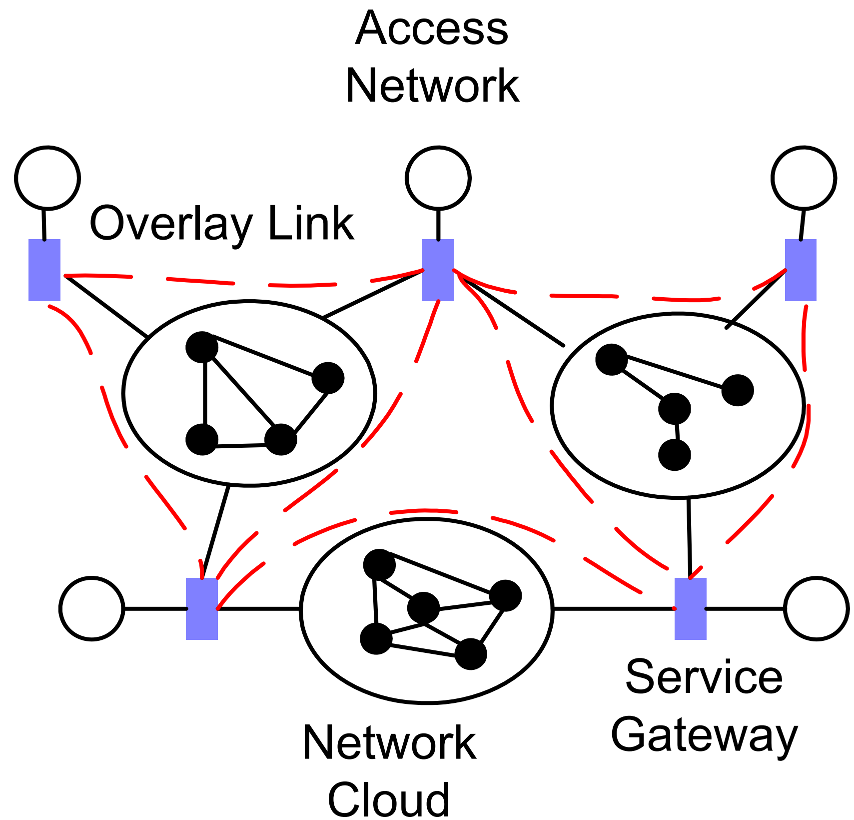

Monitoring All-Optical Networks

Routing and Wavelength Assignment

- Choose a lightpath for each incoming call

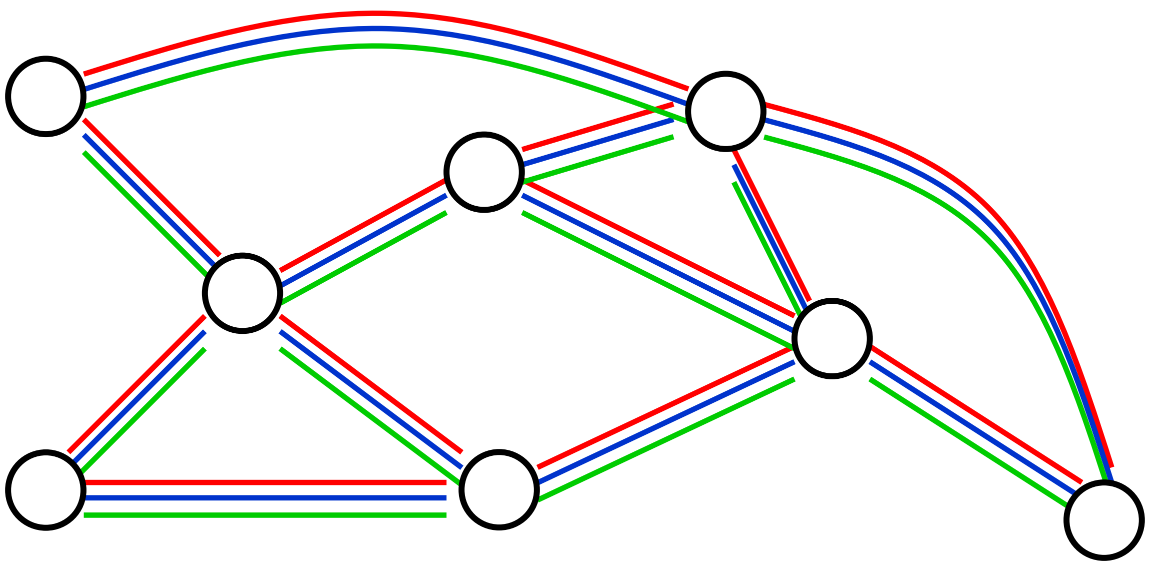

Monitoring All-Optical Networks

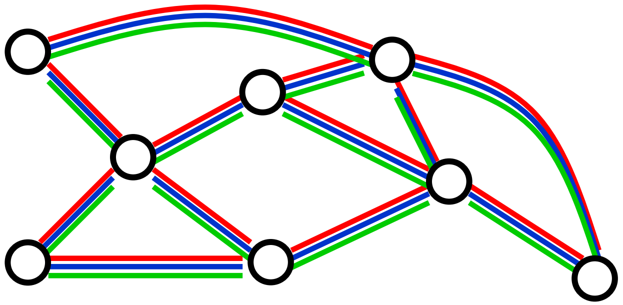

Routing and Wavelength Assignment

- Choose a lightpath for each incoming call
- Bit Error Rate reflects quality of service.
- Knowledge facilitates selection of good lightpaths
- Goals:
  - Estimate BERs of current lightpaths
  - Predict those of potential lightpaths
- Monitoring is
  - Difficult – can only measure optical properties
  - Expensive – requires high-frequency scopes

McGill

Problem Summary

- Would like performance for all paths
  - Loss rates, mean delay, delay variance, …
- BUT we only get to **measure a small subset** of paths
  - Full measurement infeasible # paths ~ $(\#nodes)^2$

→ Leads to a highly ill-posed inverse problem

- Need modeling assumptions
  - **Routing topology:** known or measured
  - **Path correlation:** link sharing
  - **Temporal correlation:** slowly varying metrics

Network (Performance) Tomography

The usual setup: $y = G\,x$

- y: path metrics
- G: routing matrix
- x: link metrics

- Notation: Routing Matrix G
  - Rows $\leftrightarrow$ paths
  - Columns $\leftrightarrow$ links
  - $G_{i,j}$ = 1 if path i transits link j

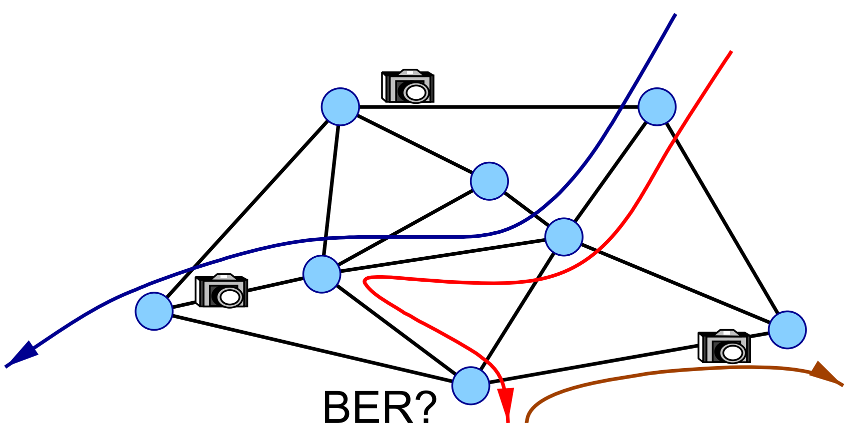

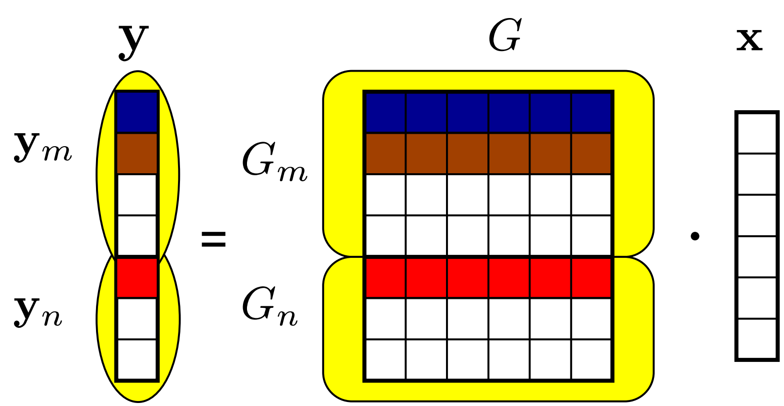

Related Work

The usual setup: $y = G\, x$

- y: path metrics
- G: routing matrix
- x: link metrics

- Cáceres, Duffield, Horowitz & Towsley – Infocom'99
  Coates, Hero, Nowak, Yu – SP Mag'02
  - Estimate performance from restricted measurements
- Chen, Bindel, Song & Katz – SIGCOMM'04
  - In practical scenarios, only need to observe $O(n \log n)$ paths in a network of n nodes
- Chua, Kolaczyk & Crovella – Infocom'05
  - Linear estimation using effective rank to approximate G

Our Approach

Our setup: $y = Gx = B\,\beta$

path metrics → y

"basis" vectors (depends on routing) → B

coefficients → β

- Build a good basis B for representing path metrics y
  - Good if β is sparse or compressible
  - Exploit correlation assumptions
- Estimate y from measurements using non-linear techniques
  - Related to compressed sensing, sparse approximation, lasso
  - Effective when we have a good representation

McGill

Compressible Signals

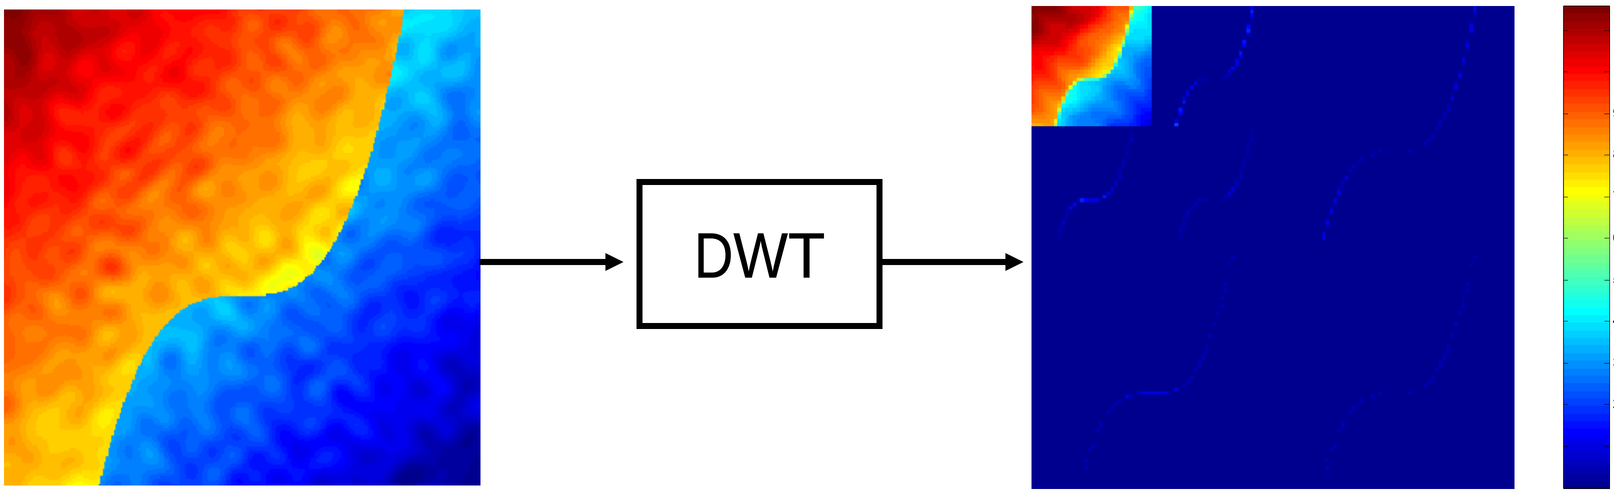

$$y = B\beta \qquad\qquad \beta$$

$|y_i|$

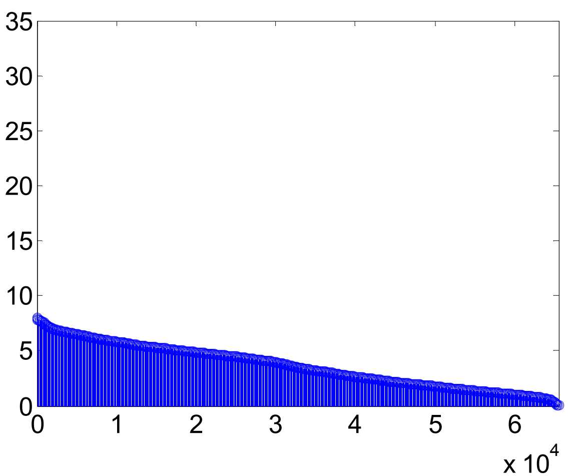

$|\beta_i|$

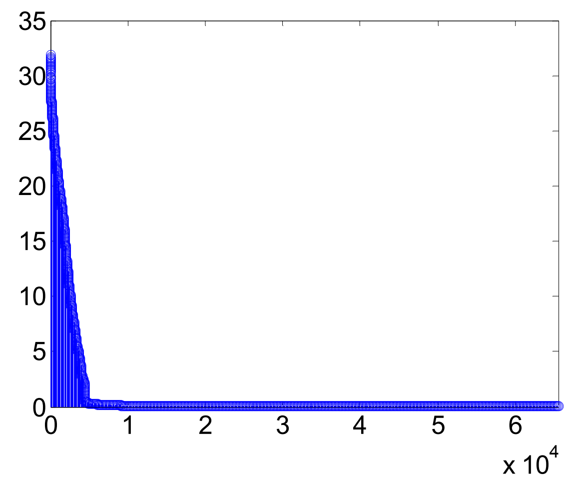

Approximating Compressible Signals

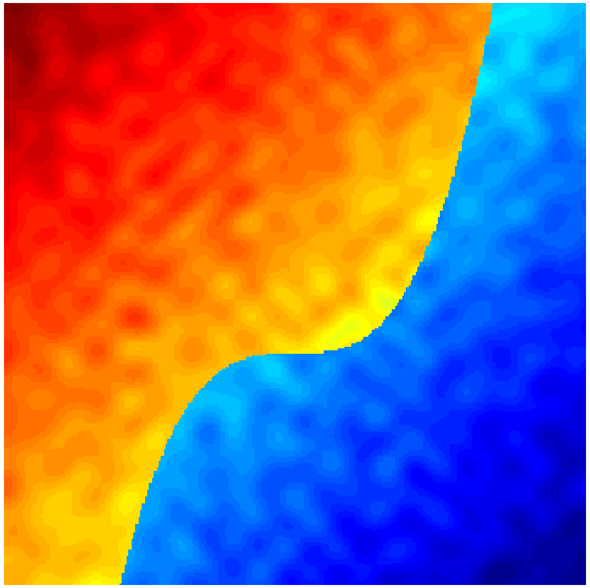

original signal
(256 x 256 $\approx$ 65,000)

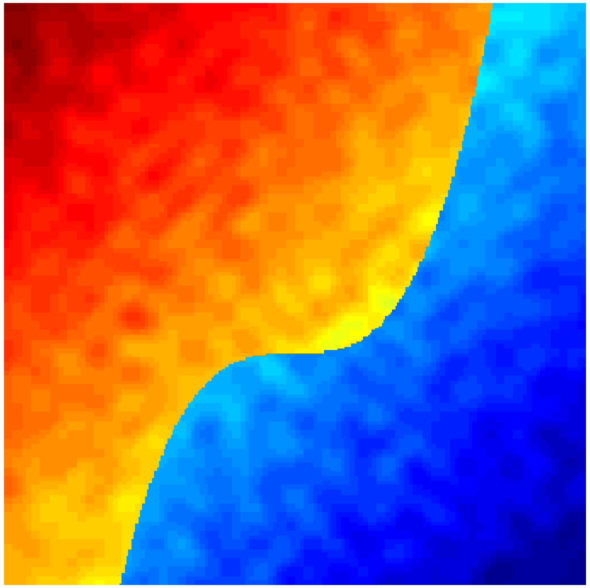

6500 largest wavelet
coefficients (10x)

Compression: Keep the large coefficients (JPEG 2000)

Compressed Sensing: Estimate the large coefficients from a few measurements [Candès, Romberg & Tao '04; Donoho '04]

Representing Data on a Graph

- Need to compress path metric "signals"
  - Network data has an irregular structure, no natural order of nodes
- Graph Wavelets [Crovella & Kolaczyk '03]
  - Extend continuous wavelets to graphs
- Diffusion Wavelets [Coifman & Maggioni '04]
  - Build wavelets from a diffusion operator
  - Define a graph relating path metric variables

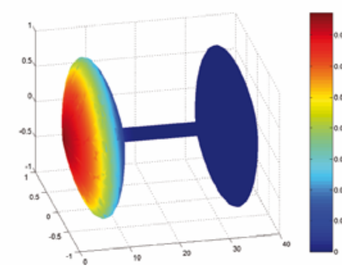

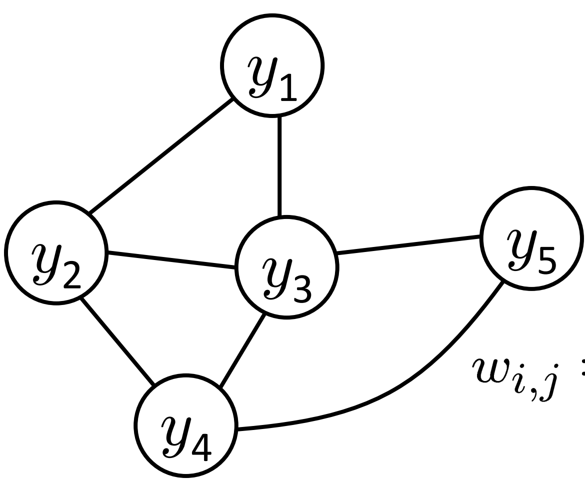

$$w_{i,j} = \frac{|\mathcal{P}_i \cap \mathcal{P}_j|}{|\mathcal{P}_i \cup \mathcal{P}_j|}$$

- One vertex for each path variable

- Place an edge between paths that share at least one link

- $\mathcal{P}_i$ is the set of links in path i
- Normalize to get diffusion operator

McGill

Compressing with Diffusion Wavelets

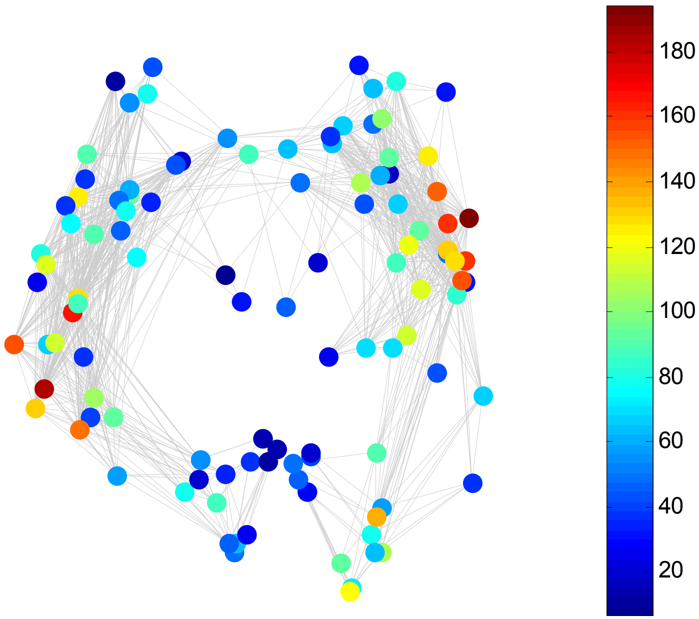

original signal
on Abilene paths

$|y_i|$

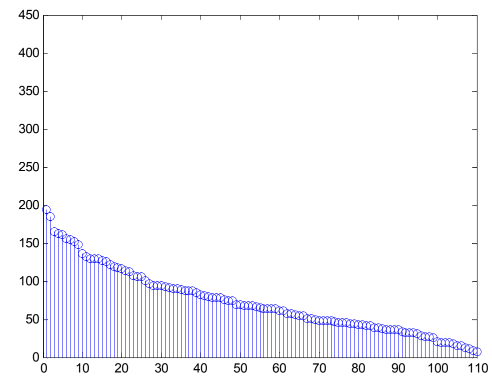

$|\beta_i|$

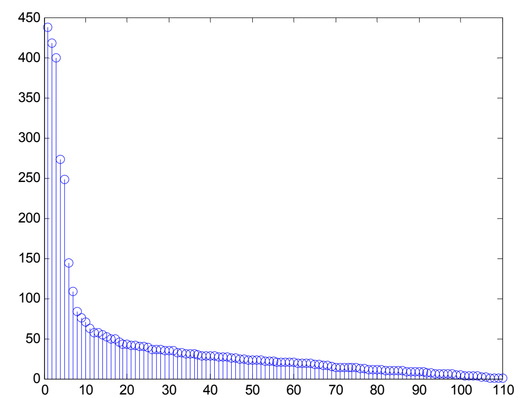

Compressing with Diffusion Wavelets

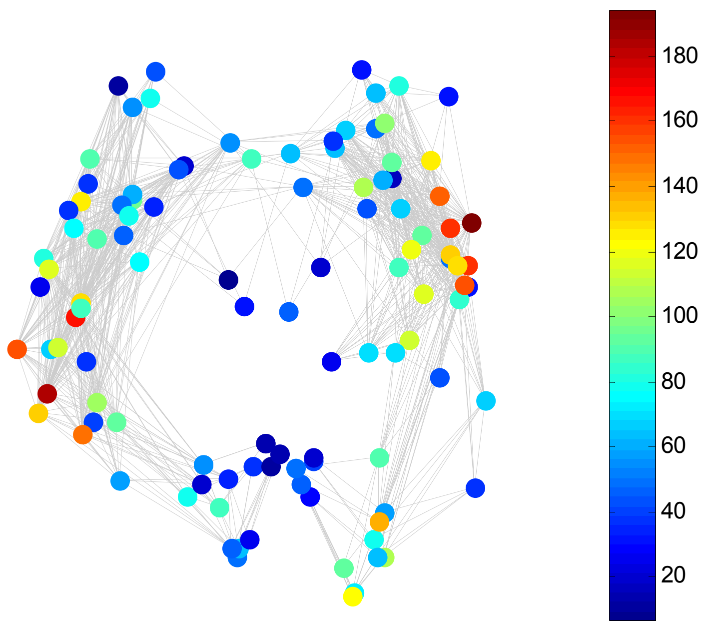

original signal
on Abilene paths

most significant
basis vector

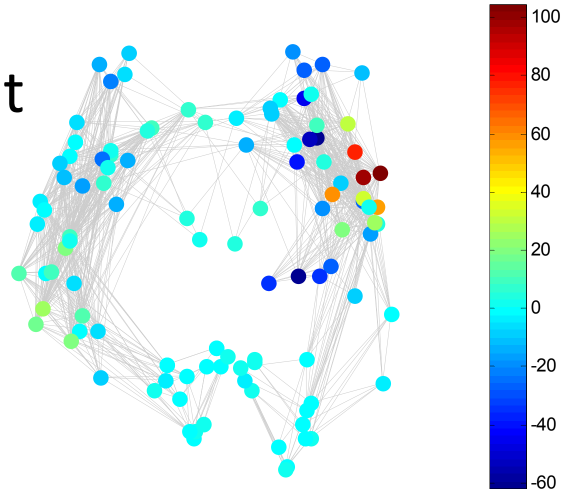

second
most significant
basis vector

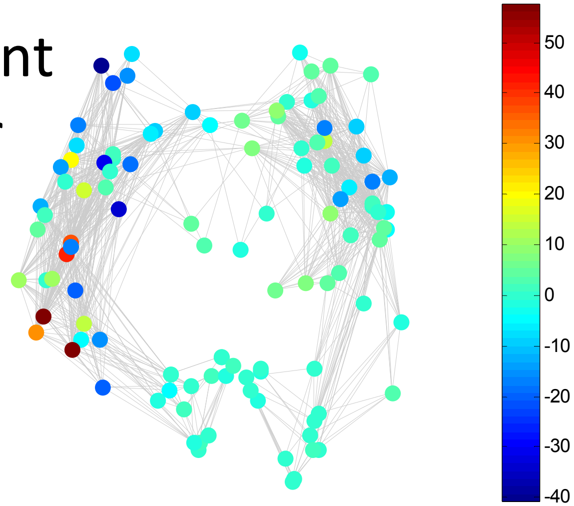

Exploiting Temporal Correlation

- Stack multiple time-steps into a vector

$$y = \begin{bmatrix} y(t) \\ y(t+1) \\ \vdots \\ y(t+\tau) \end{bmatrix}$$

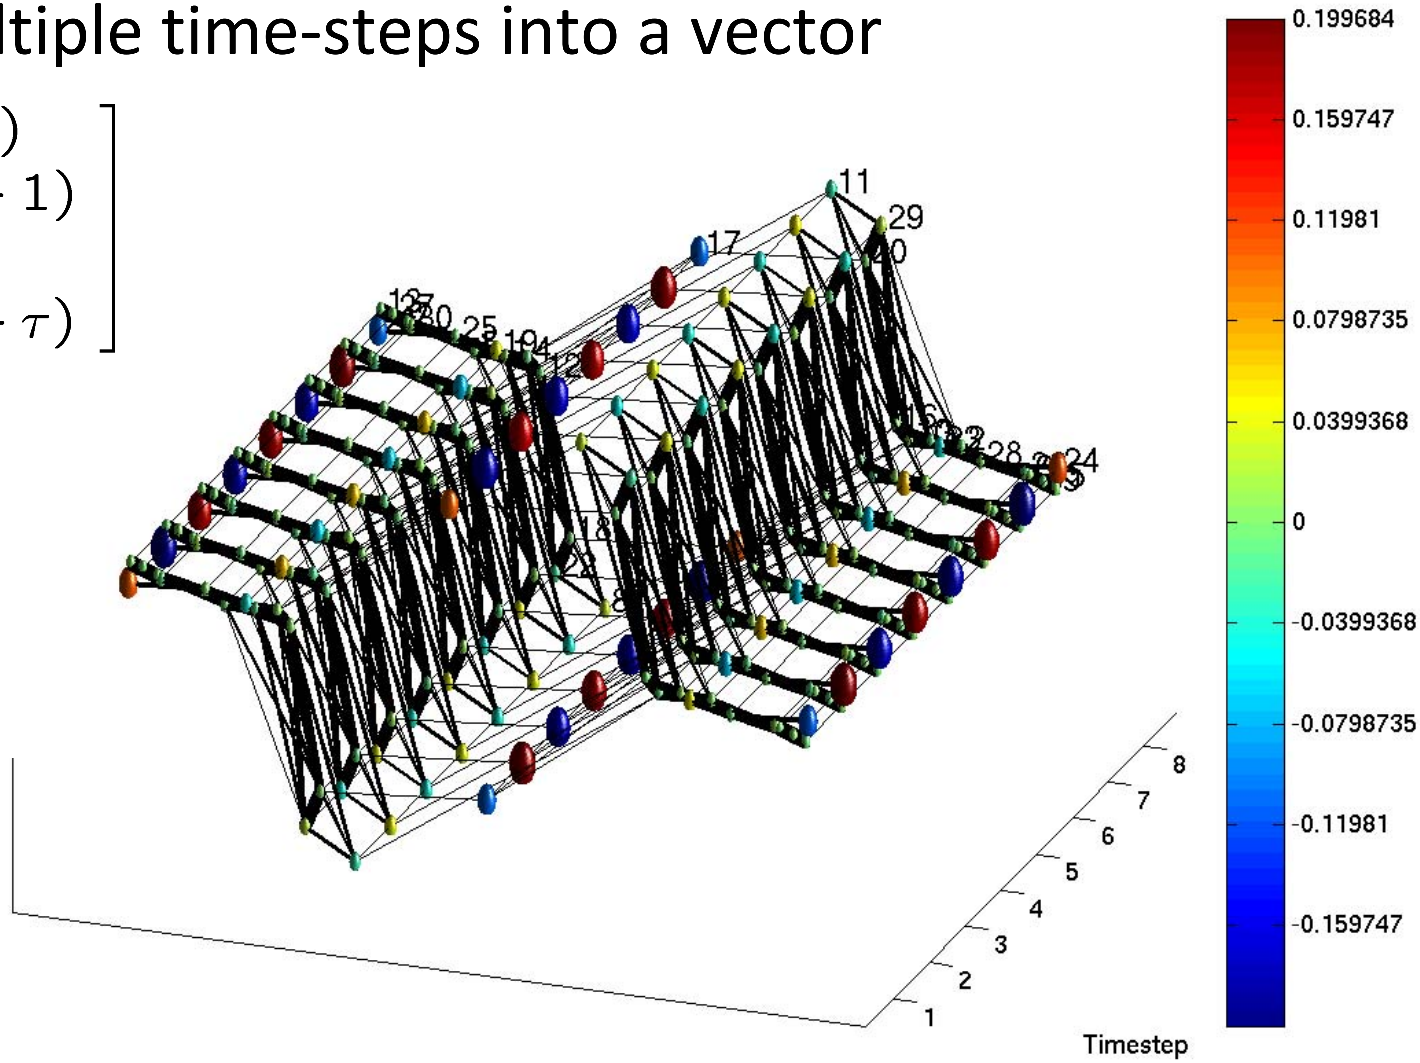

Multiple Time-Step Diffusion Matrix

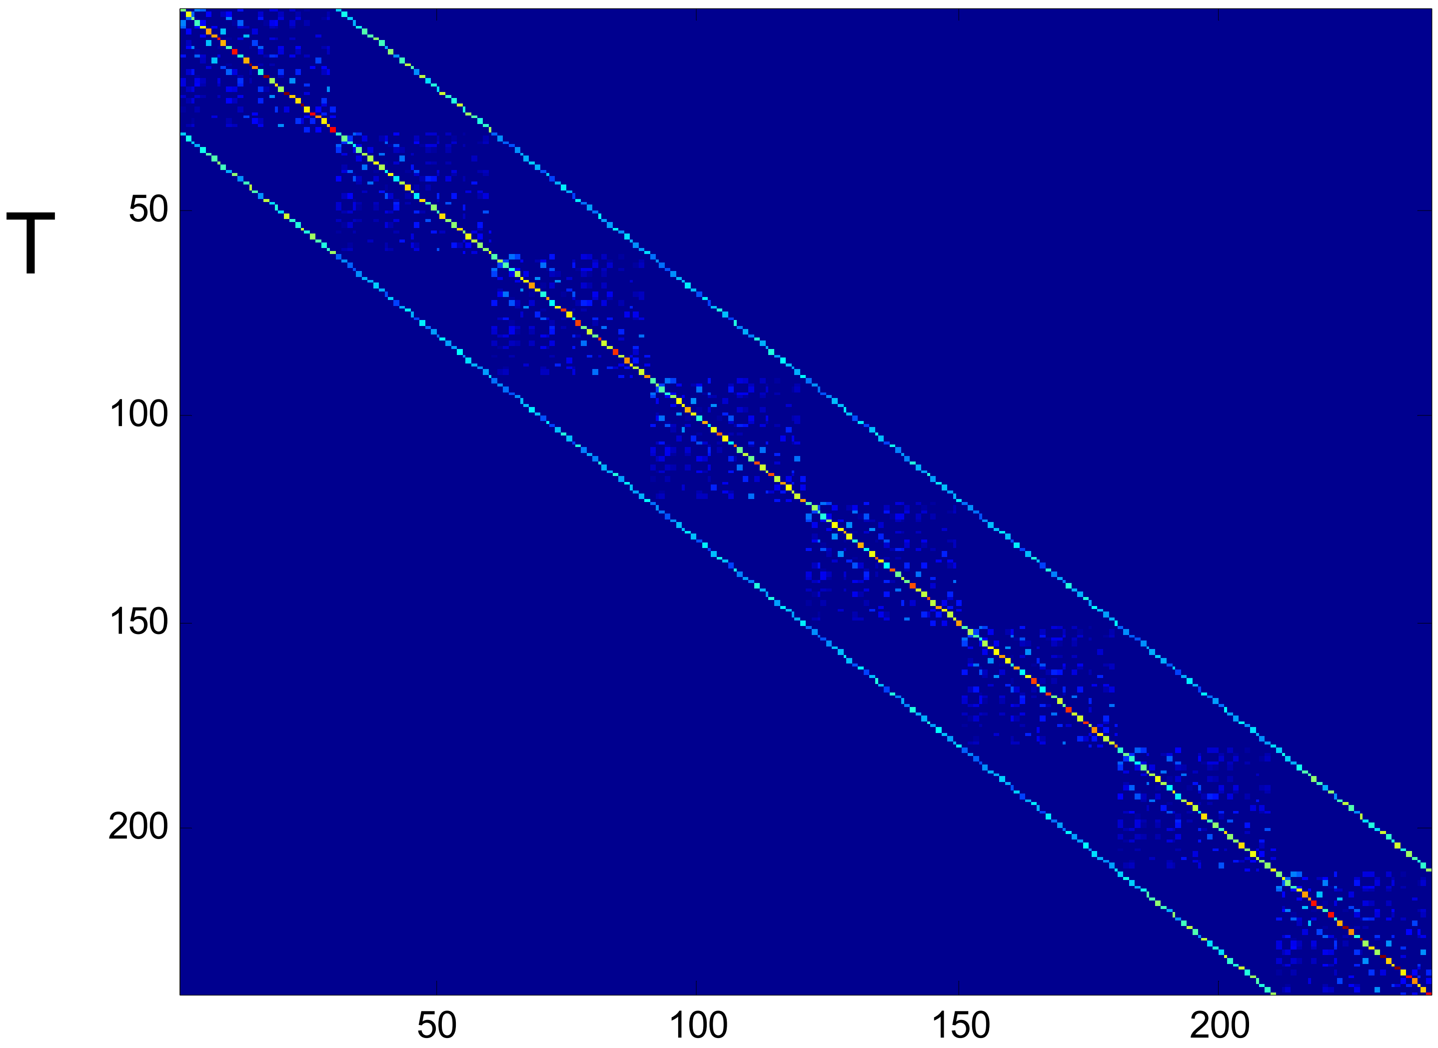

Nonlinear Estimation

1. Measure $y_s = Ay$

Binary matrix A picks off the few paths we measure

2. Solve $\widehat{\beta} = \arg\min_{\beta} \|\beta\|_1$ subject to $y_s = AB\beta$

Find a simple explanation for the measurements

3. Estimate $\widehat{y} = B\widehat{\beta}$

Plugin to recover path metrics

Recent development suggest (Candès, Tao, Donoho…)

If B is really a good representation for path metrics (i.e., β is a compressed version of y), then under appropriate conditions on A, even with few measurements, the estimation error $\|\widehat{\beta} - \beta\|$ will be small

Dataset – Abilene Network

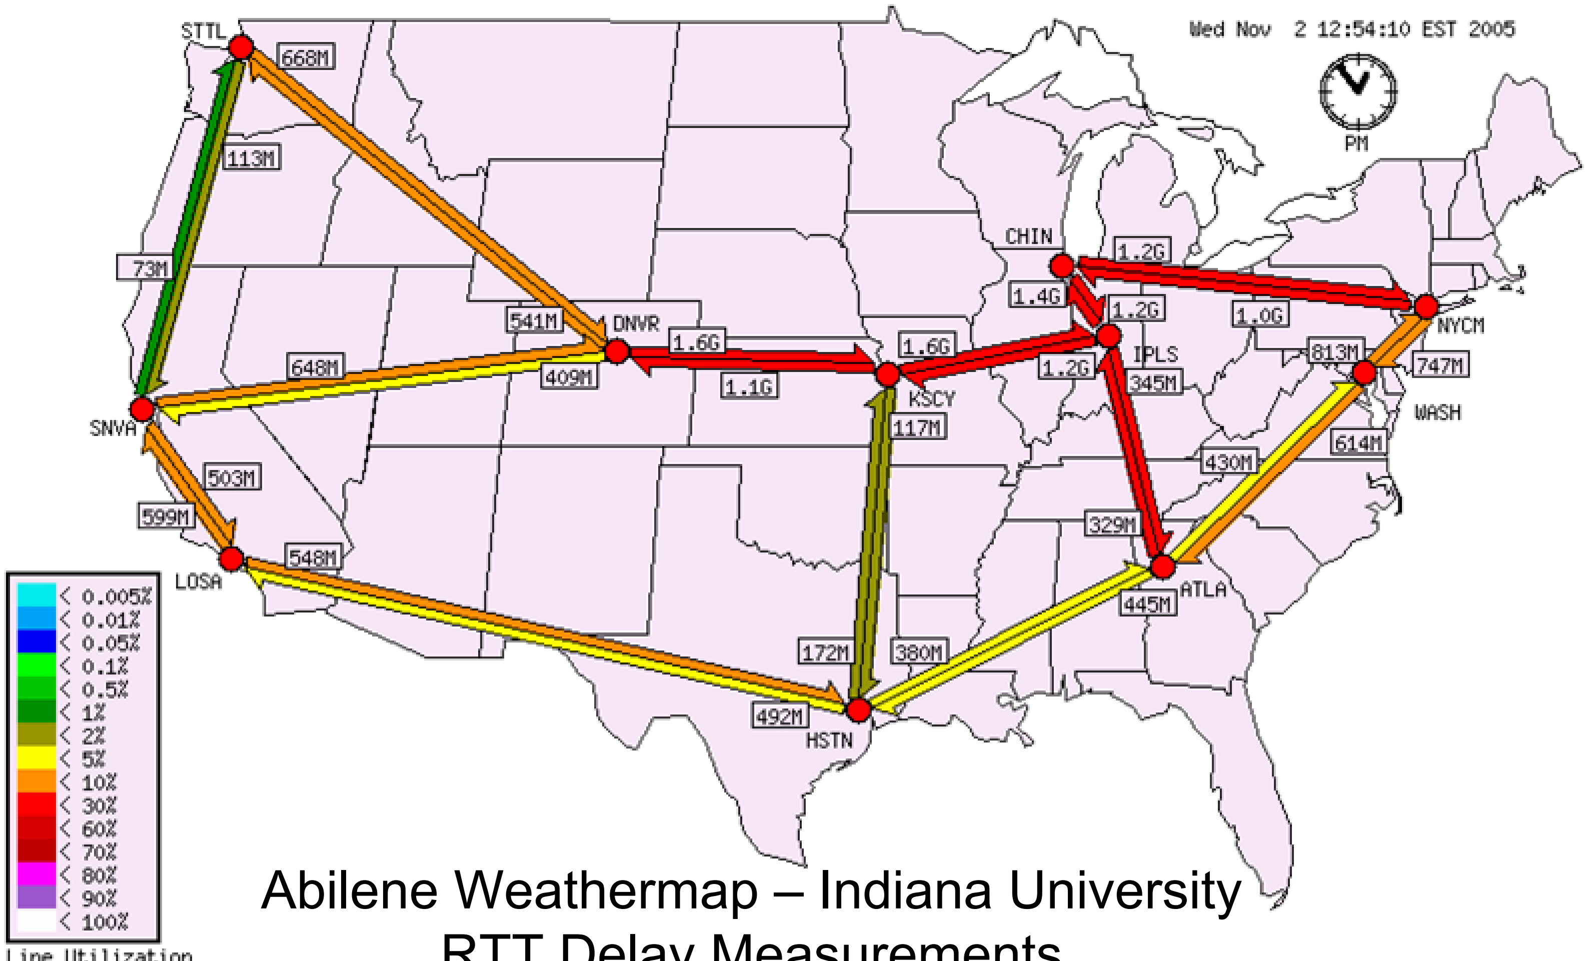

Abilene Weathermap – Indiana University
RTT Delay Measurements

Thanks to Rick Summerhill and Mark Fullmer at Abilene for providing access to the data.

Empirical Compressibility

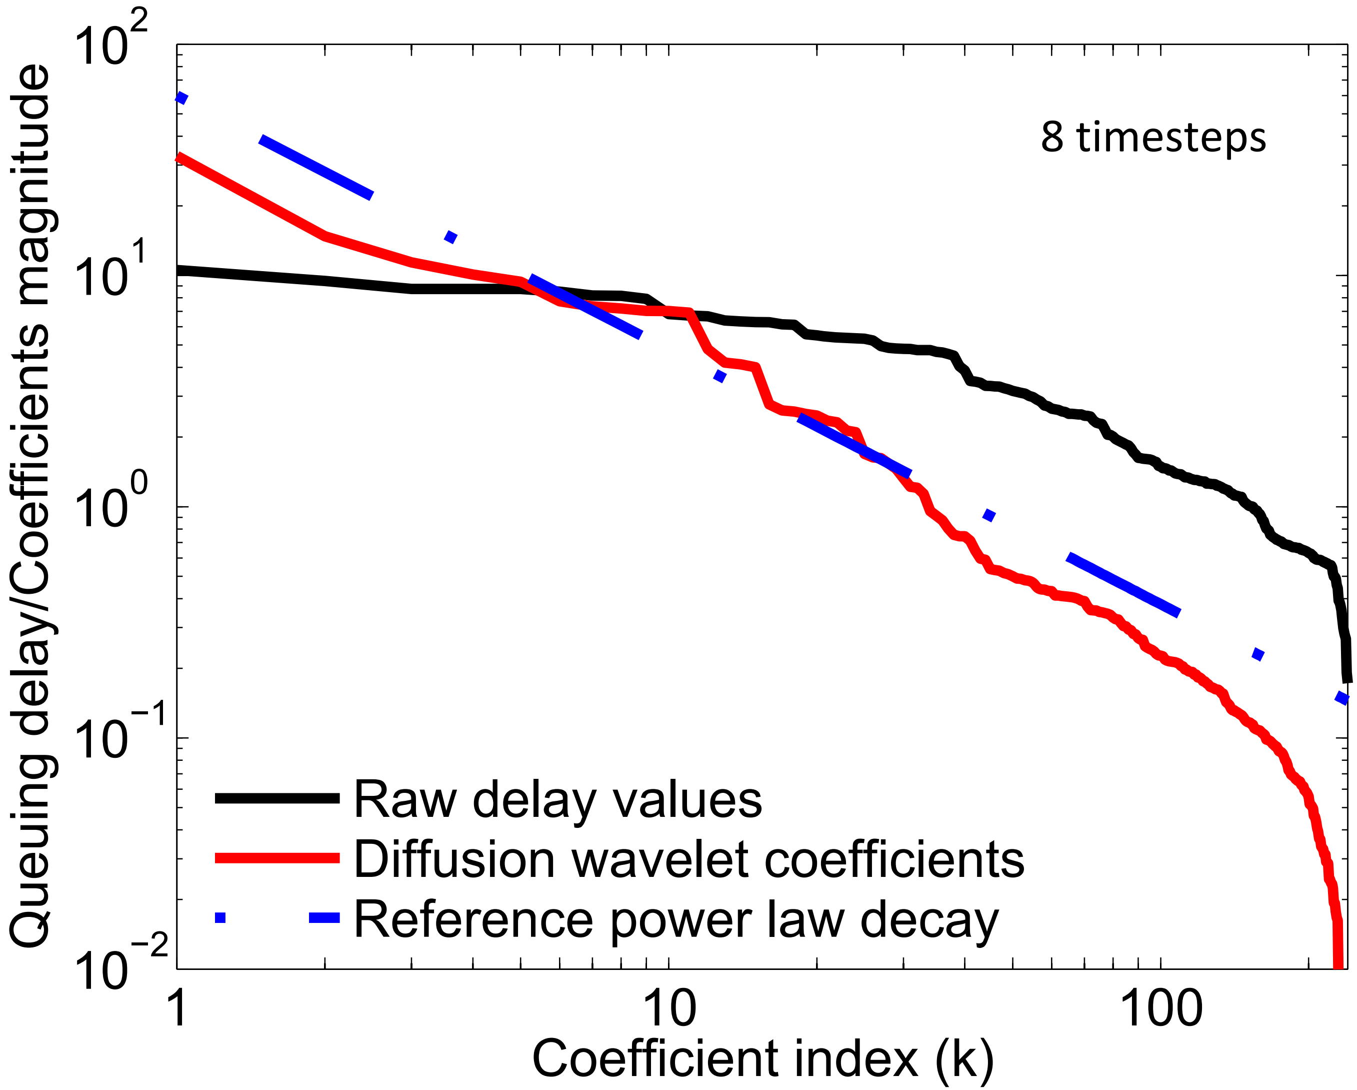

Abilene Results

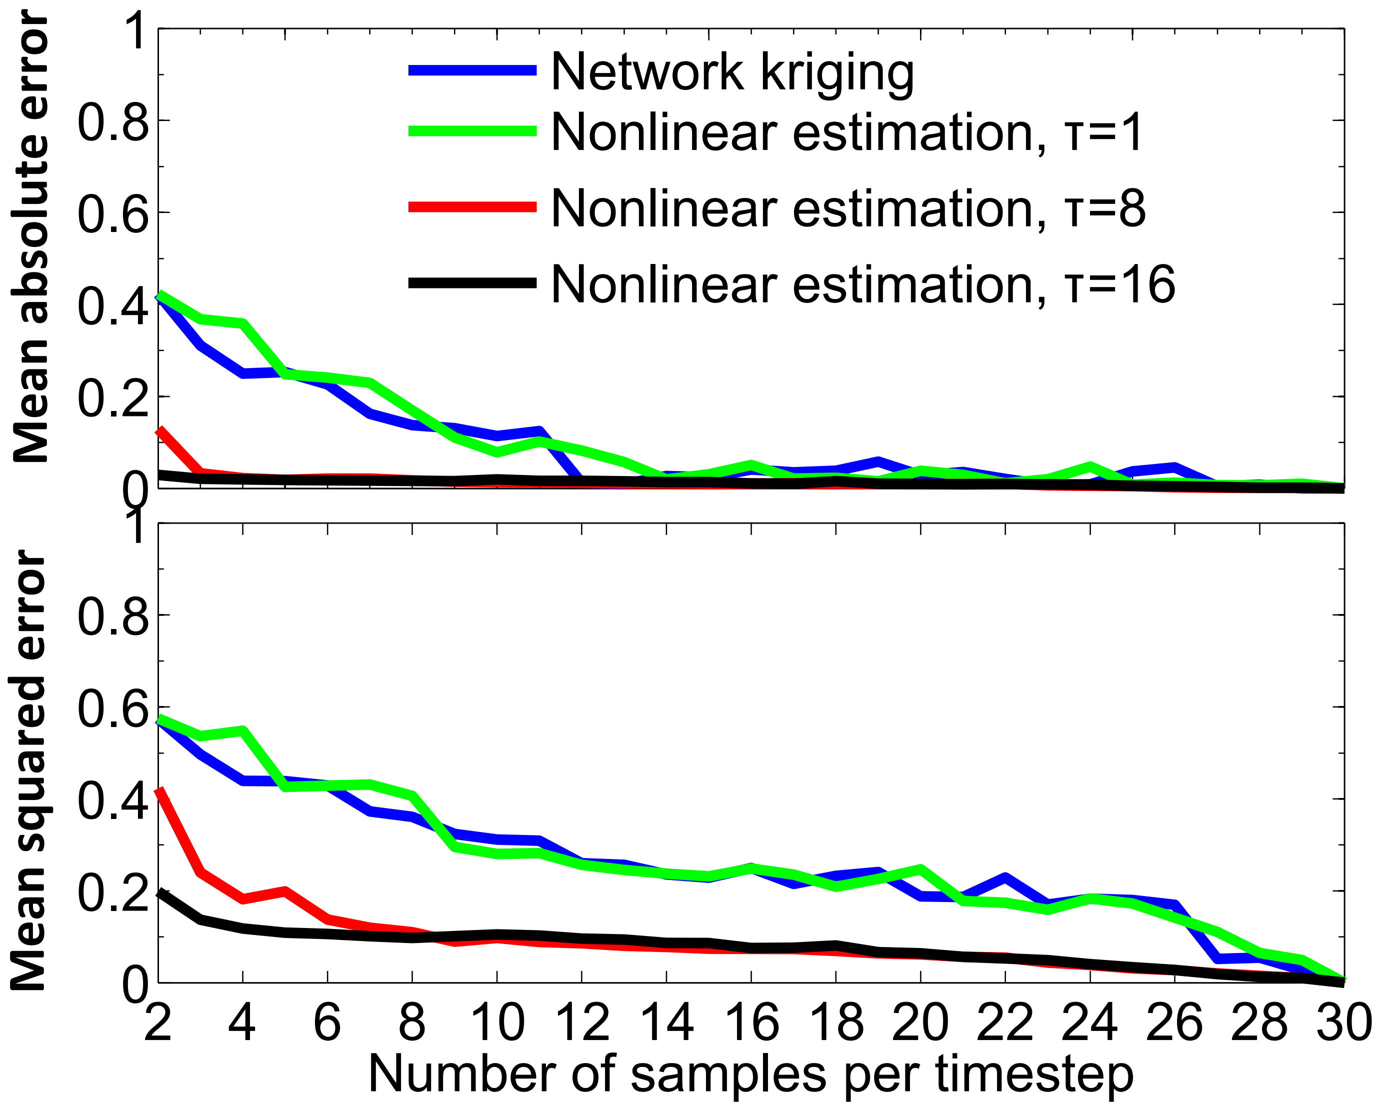

Recovered Signals

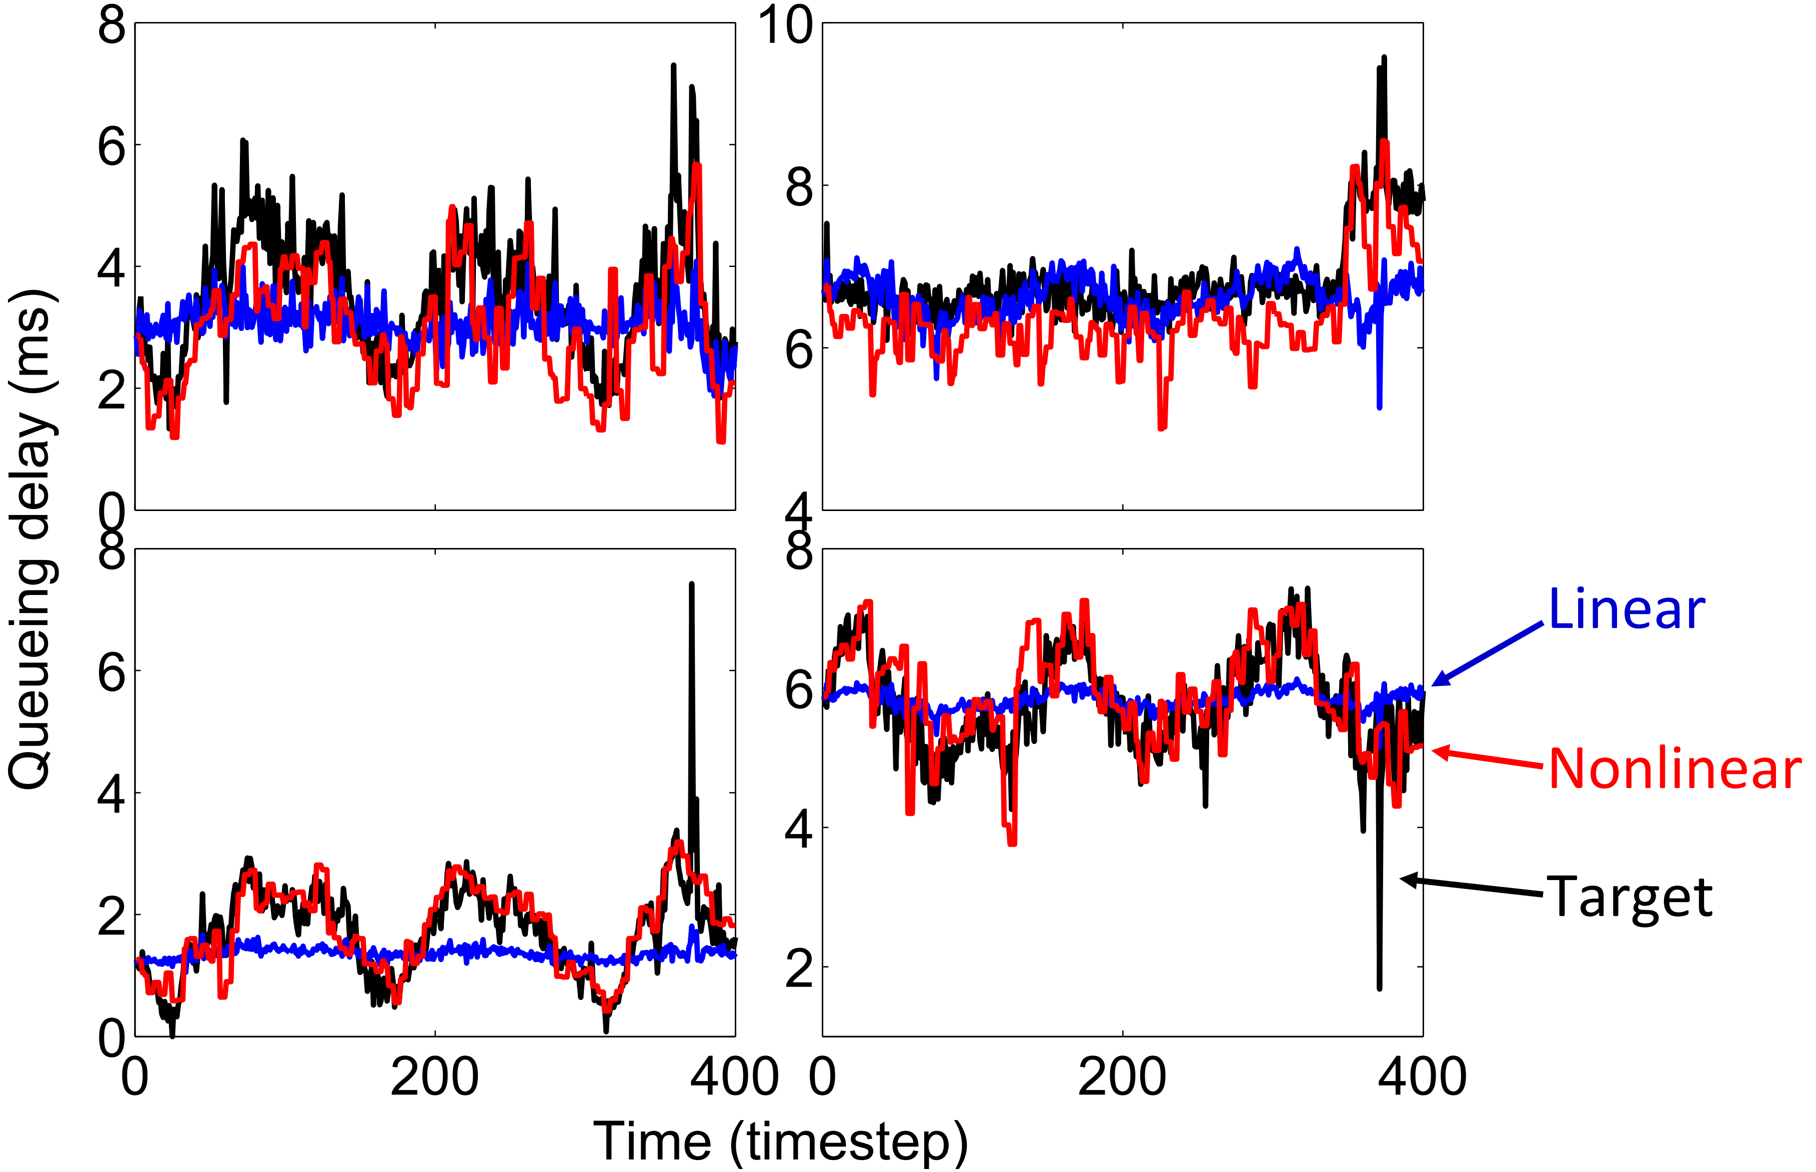

All-Optical Network Monitoring

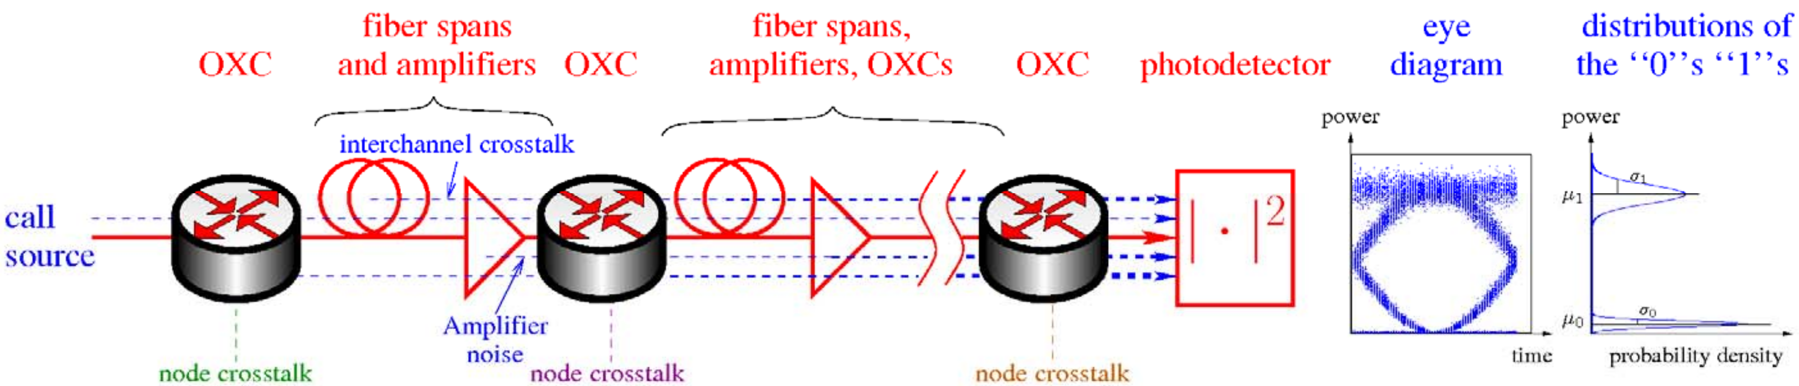

- Denote by μ_0 and μ_1 the means of 0/1
- Corresponding standard deviations are σ_0 and σ_1
- Under a Gaussian noise model

$$BER = \frac{1}{2}\text{erfc}\left(\frac{\mu_1 - \mu_0}{\sigma_1 + \sigma_0}\right)$$

- BER estimation: all four parameters for each lightpath

McGill

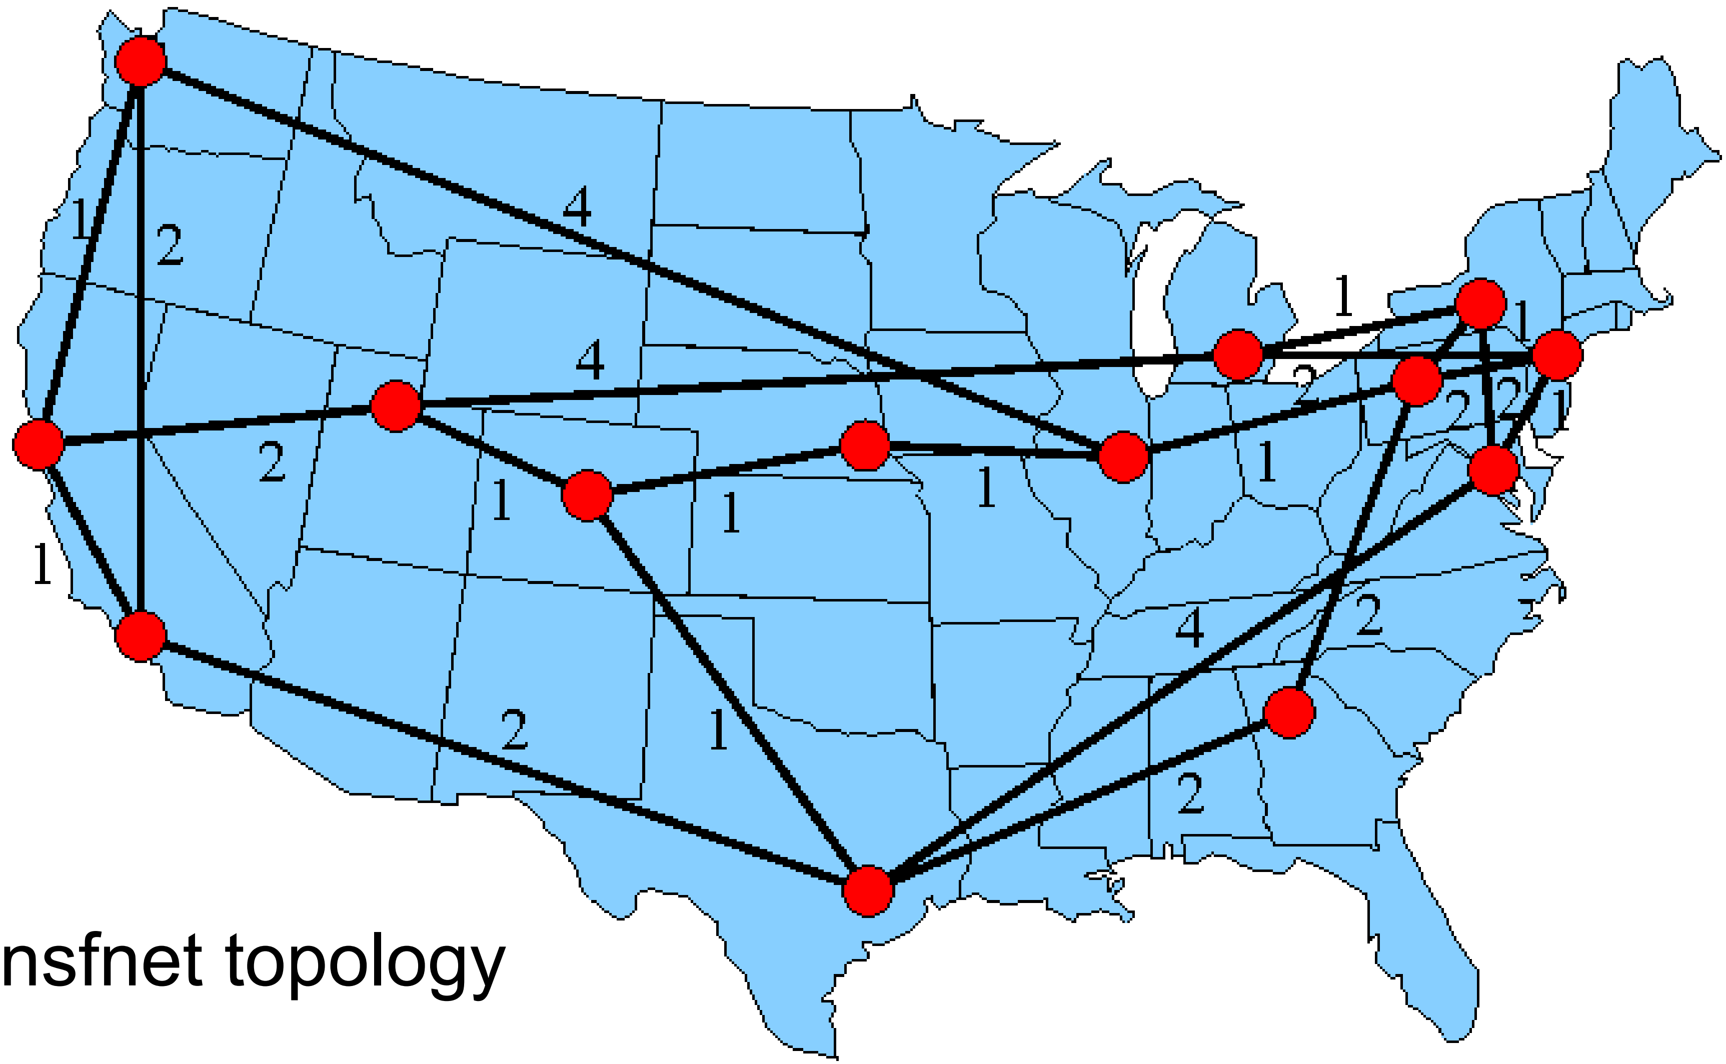

nsfnet topology

14 nodes, 42 links, 8 wavelengths

$\approx$ 60 lightpaths

Empirical Compressibility

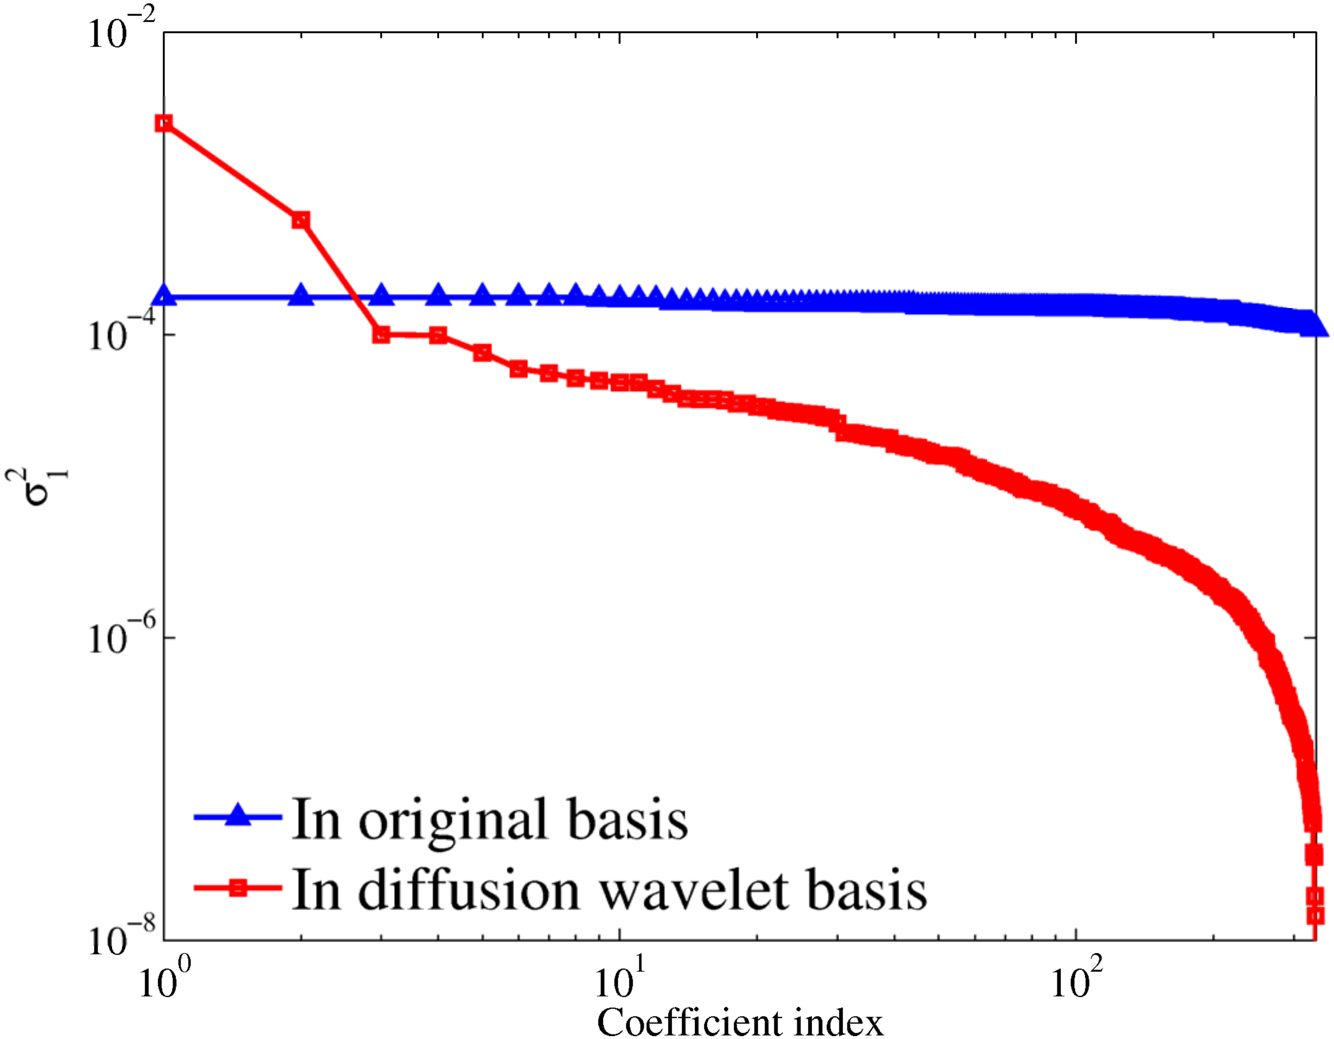

Performance Results

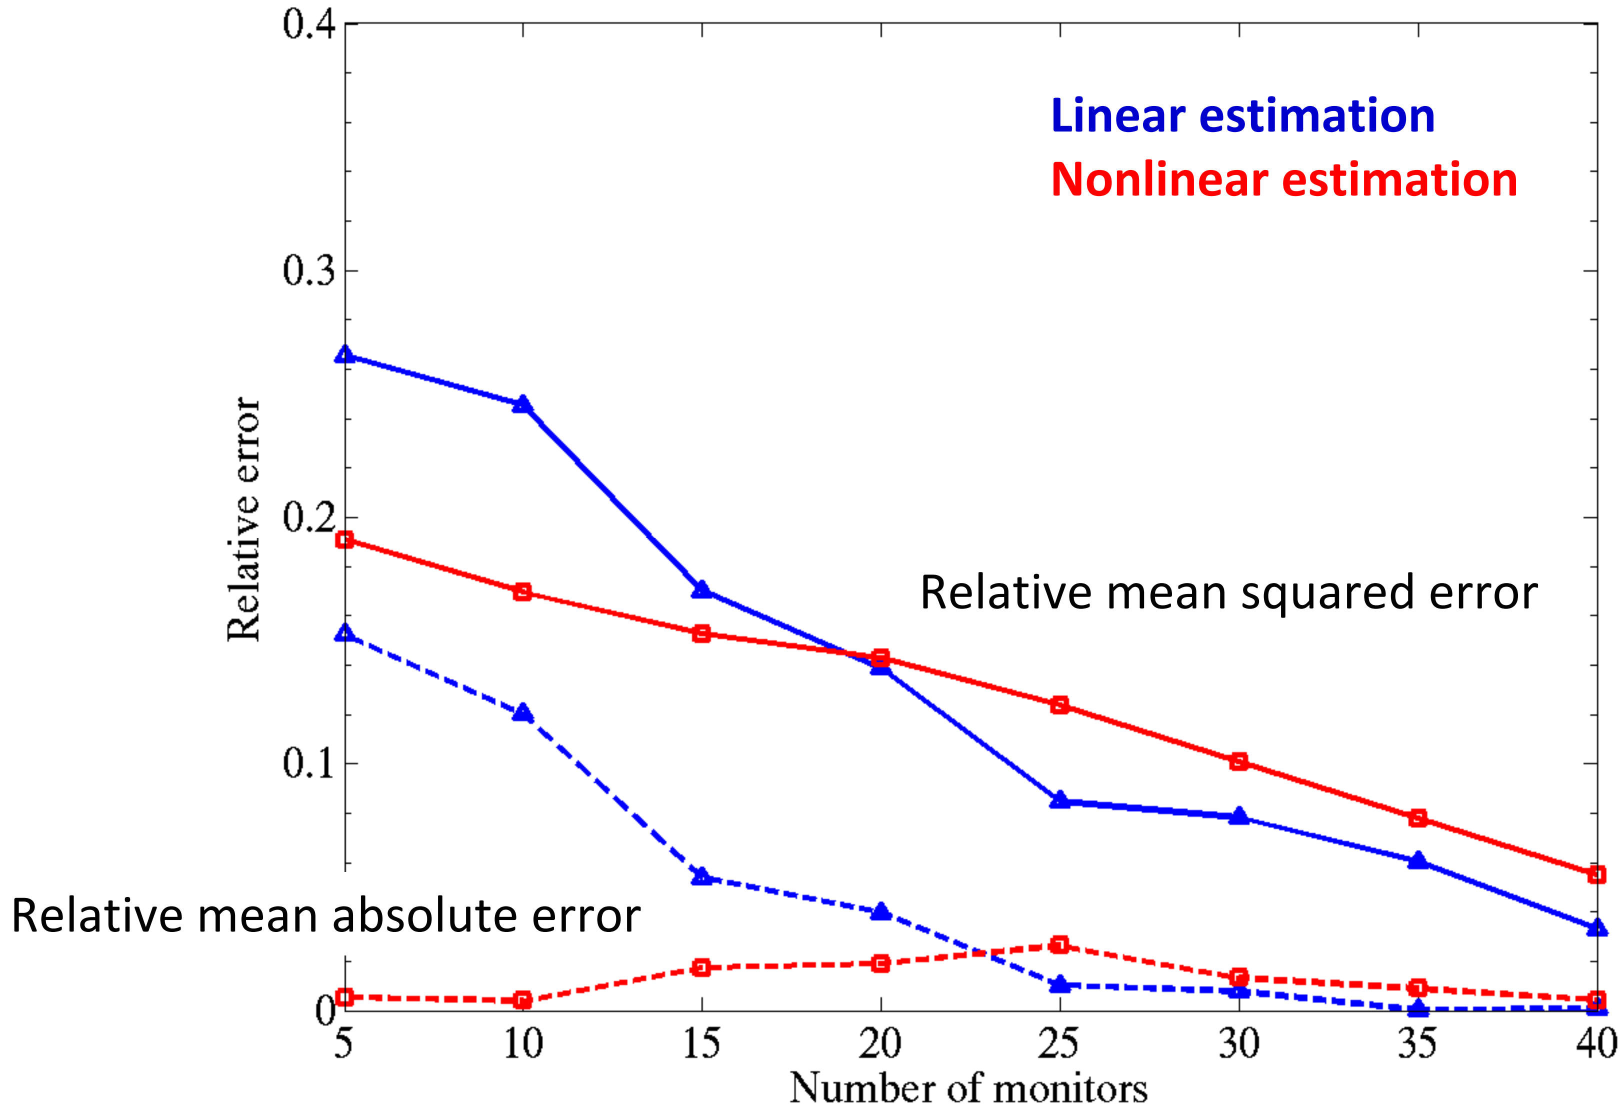

All Optical Extensions

- Monitoring Equipment is expensive
  - Joint monitor placement, path selection
  - Actively lighting augmented paths to reach a monitor

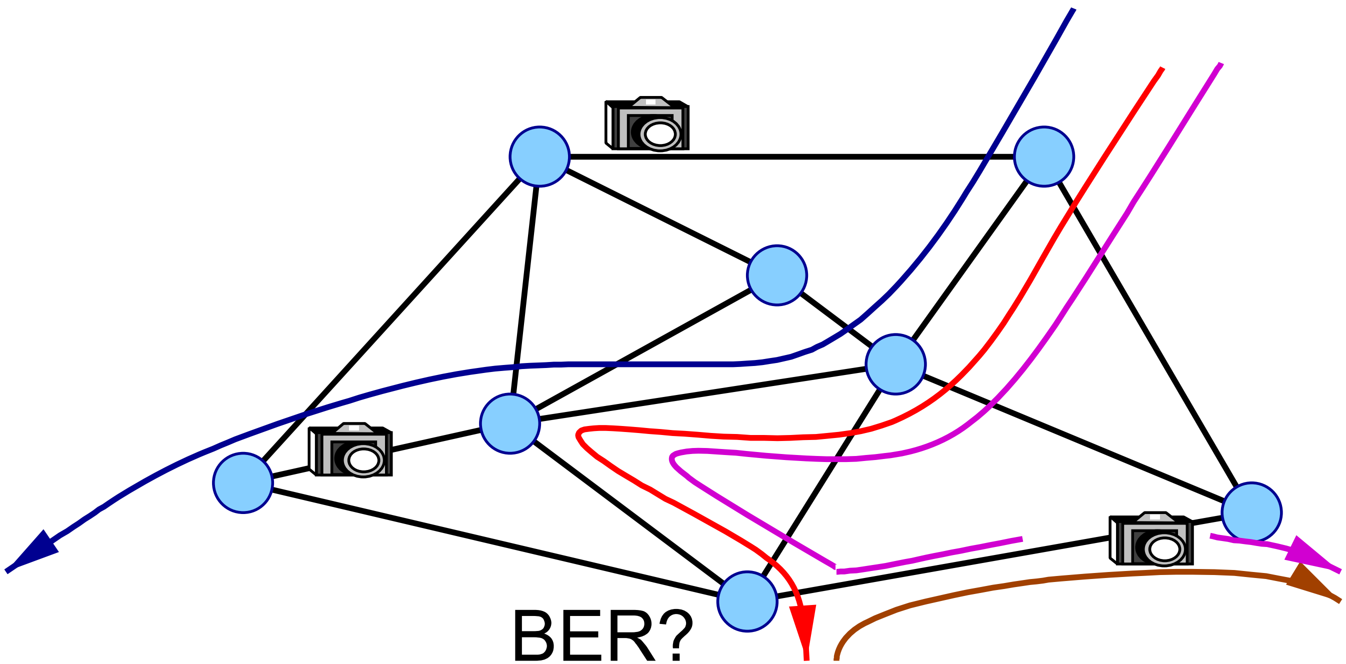

Active Monitoring of All-Optical Nets

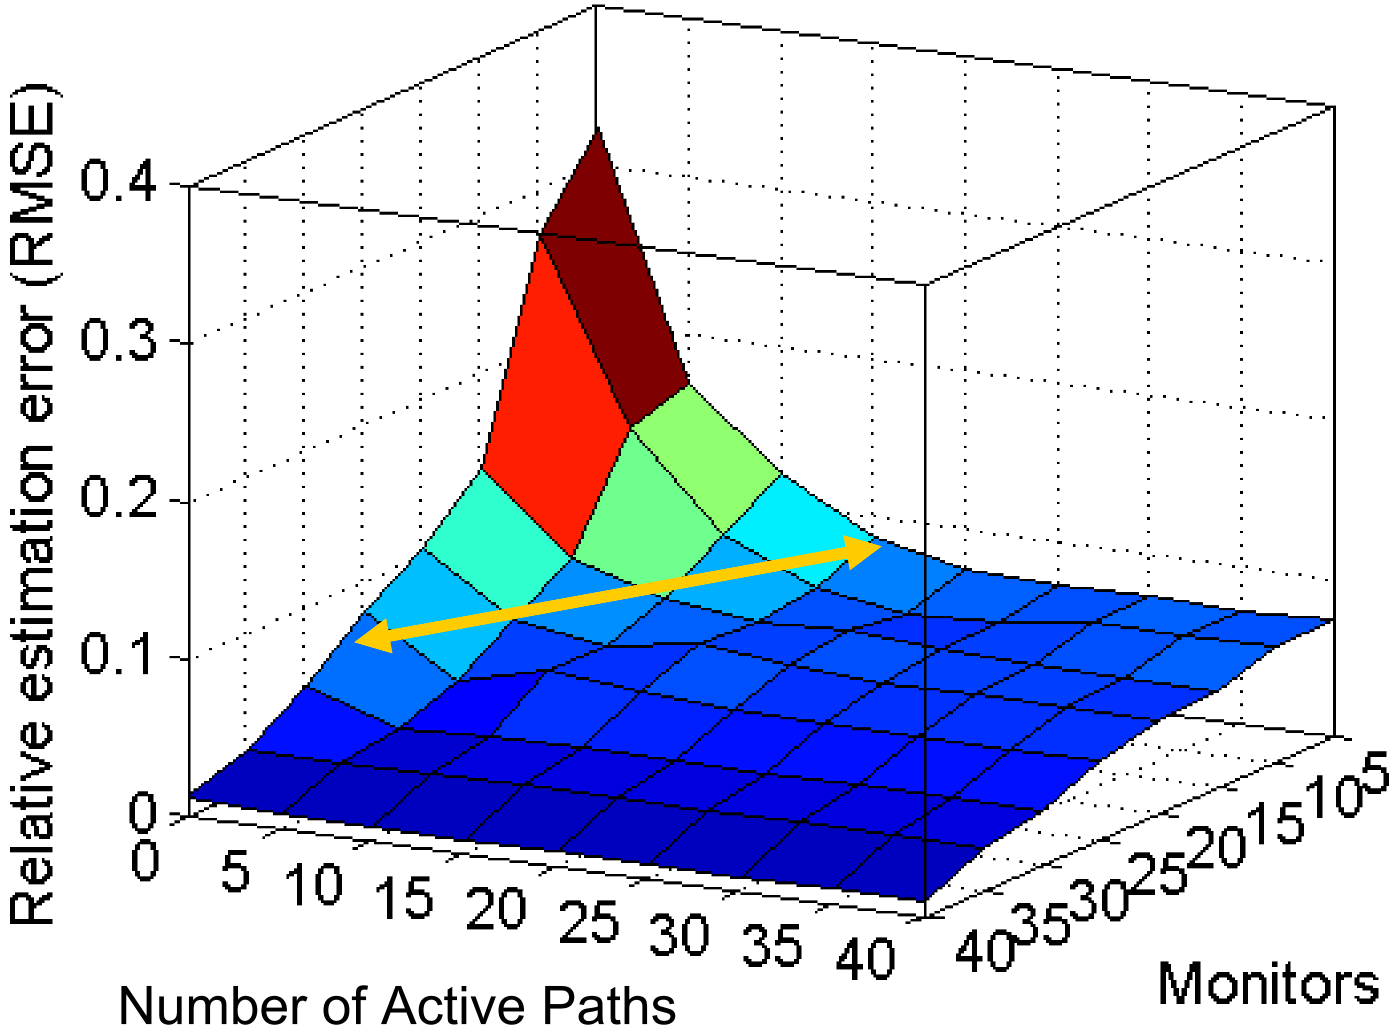

Conclusions

- Multi-scale network monitoring
  - Predict performance on correlated paths by building good representations of the data
  - Representations only as good as our prior knowledge
- Ongoing and Future work
  - Other compressing representations for network data
  - Active, online path selection algorithms
  - Extension to available bandwidth estimation

michael.rabbat@mcgill.ca
http://www.ece.mcgill.ca/~mrabba1

Diffusion Scale 0

T

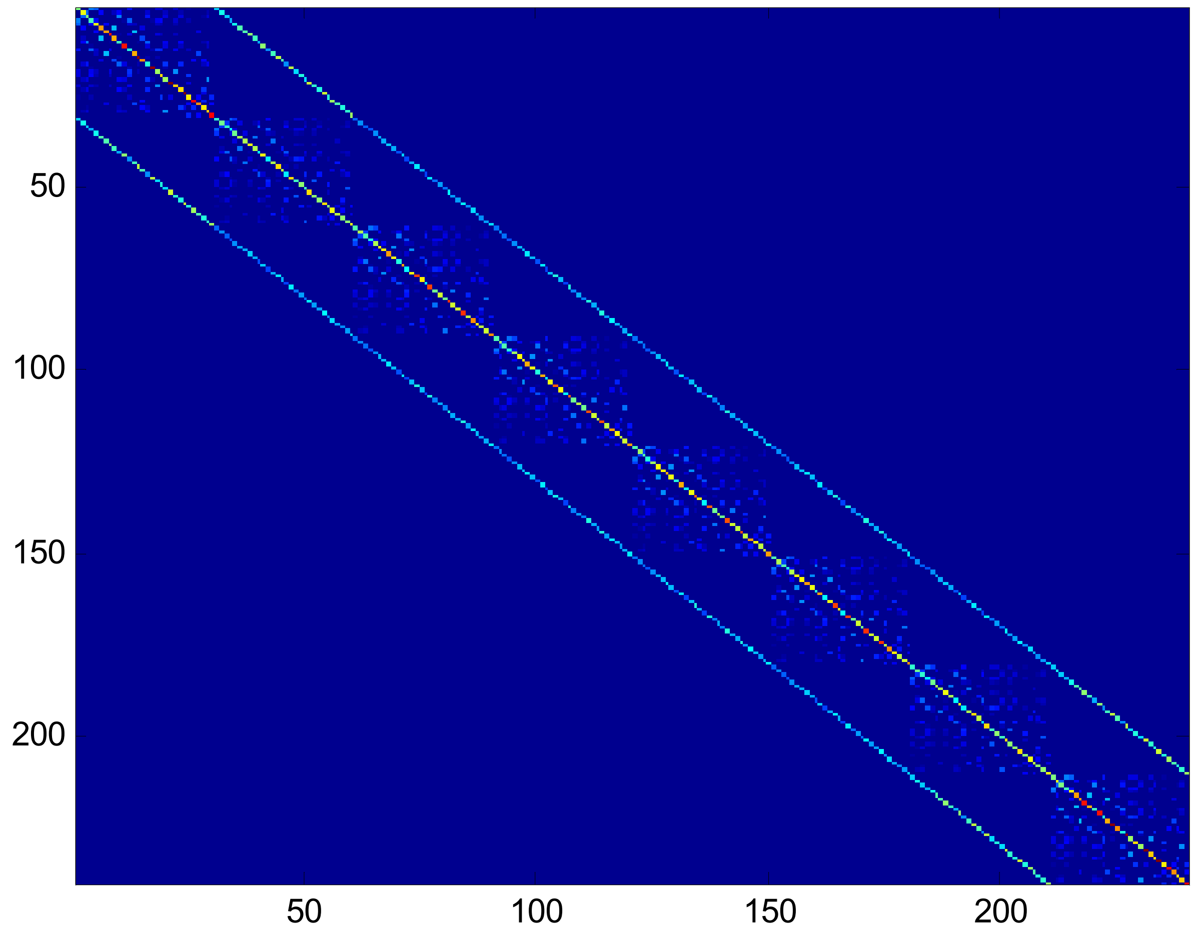

Diffusion Scale 1

T^2

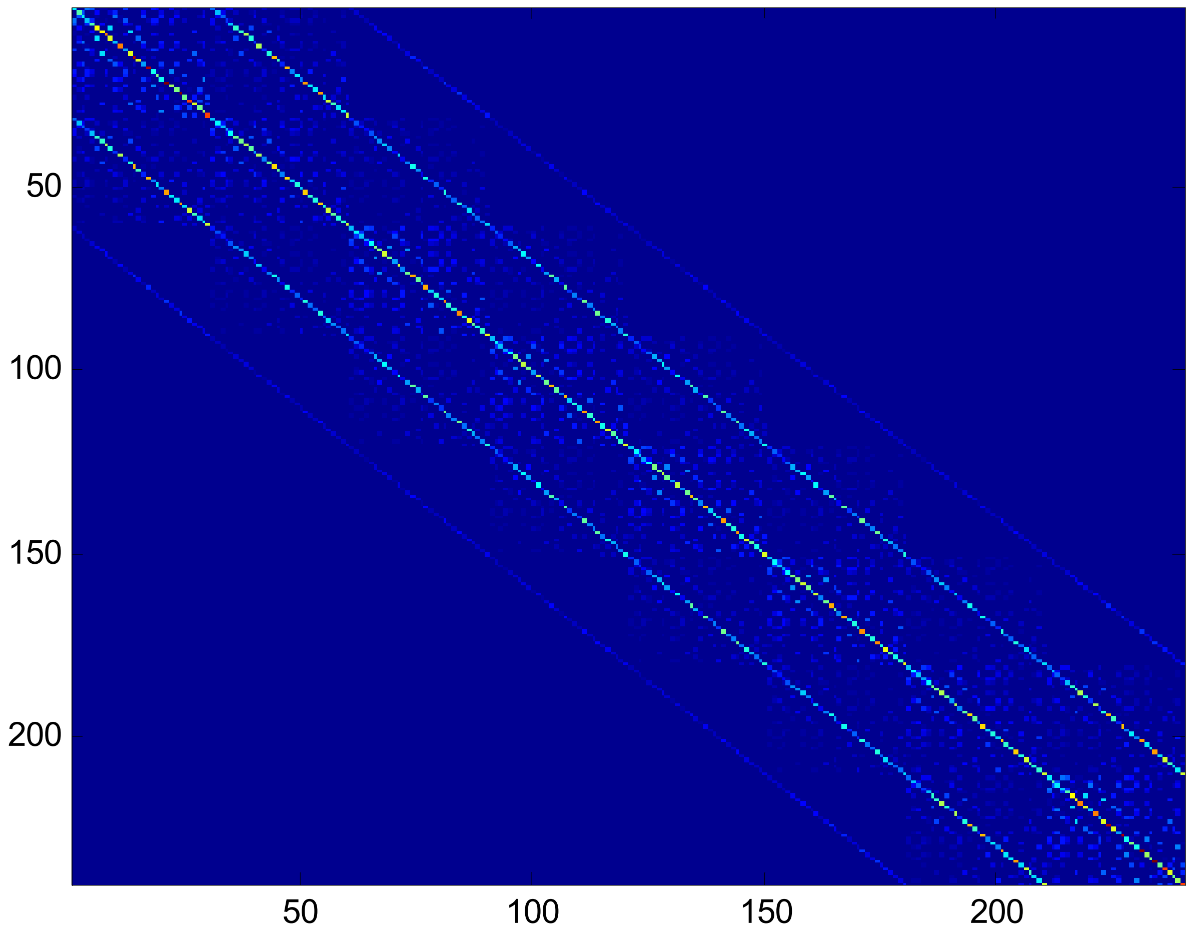

Diffusion Scale 2

T^4