Analysis of Network Flow Data

Many networks serve as conduits – either literally or figuratively – for *flows*.

I.e., they facilitate the movement of something.

Examples include

- transportation networks
  i.e., flows of commodities and/or people
- communication networks
  i.e., flows of data
- networks of international trade relations
  i.e., flows of capital

Statistical Analysis of Network Flow Data

Large literature on flows concerned with network design, provisioning, and routing.

More recently, substantial energy devoted to *statistical analysis of network flow data,* particularly for purposes of *prediction*.

Examples include prediction of

- origin-destination (OD) flow volume
  (Gravity Modeling)
- link volumes
  (OD Traffic Matrix Estimation)
- OD and/or link flow costs
  (Quality of Service (QoS) Monitoring)

Focus for Today

We'll focus today on the prediction of origin-destination flow costs.

Work of Chua, Crovella, and Kolaczyk (2005,2006)[a].

Context is the Internet and prediction of cumulative *delay* for sending packets between origins and destinations.

Mathematical/statistical framework broadly applicable.

---

[a] Also summarized in Ch9.4.2 of the book. Note, however, that the notation here follows that of the papers, and *not* the book.

Collaborators and Support

Collaborators:

- Mark Crovella (BU, Computer Science)
- David Chua (State Street Global Advisors)

Work supported in part by NSF grant CCR-0325701 and ONR award N000140310043.

End-to-End Monitoring

Our interest is in network-wide monitoring of ‘end-to-end’ properties on paths.

Full measurement quickly becomes infeasible, since $\#paths \sim (\#nodes)^2$.

$\Rightarrow$ Motivates the development of methods based on reduced subsets of path measurements.

Some Notation

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a strongly connected digraph, and let $\mathcal{P}$ be the set of all paths on $\mathcal{G}$.

Define $n_v = |\mathcal{V}|$, $n_e = |\mathcal{E}|$, and $n_p = |\mathcal{P}|$.

Finally, let

- y_i, for $i \in \mathcal{P}$, be a path-index quantity, and
- x_j, for $j \in \mathcal{E}$, be the same quantity indexed on links.

Some notation (cont)

Our interest will focus on problems in which $y = Gx$, where

$$G_{i,j} = \begin{cases} 1 & \text{if path } i \text{ traverses link } j \\ 0 & \text{otherwise} \end{cases} .$$

is a routing matrix.

NOTE: Typically $n_p \gg n_e$.

Illustration: Abilene

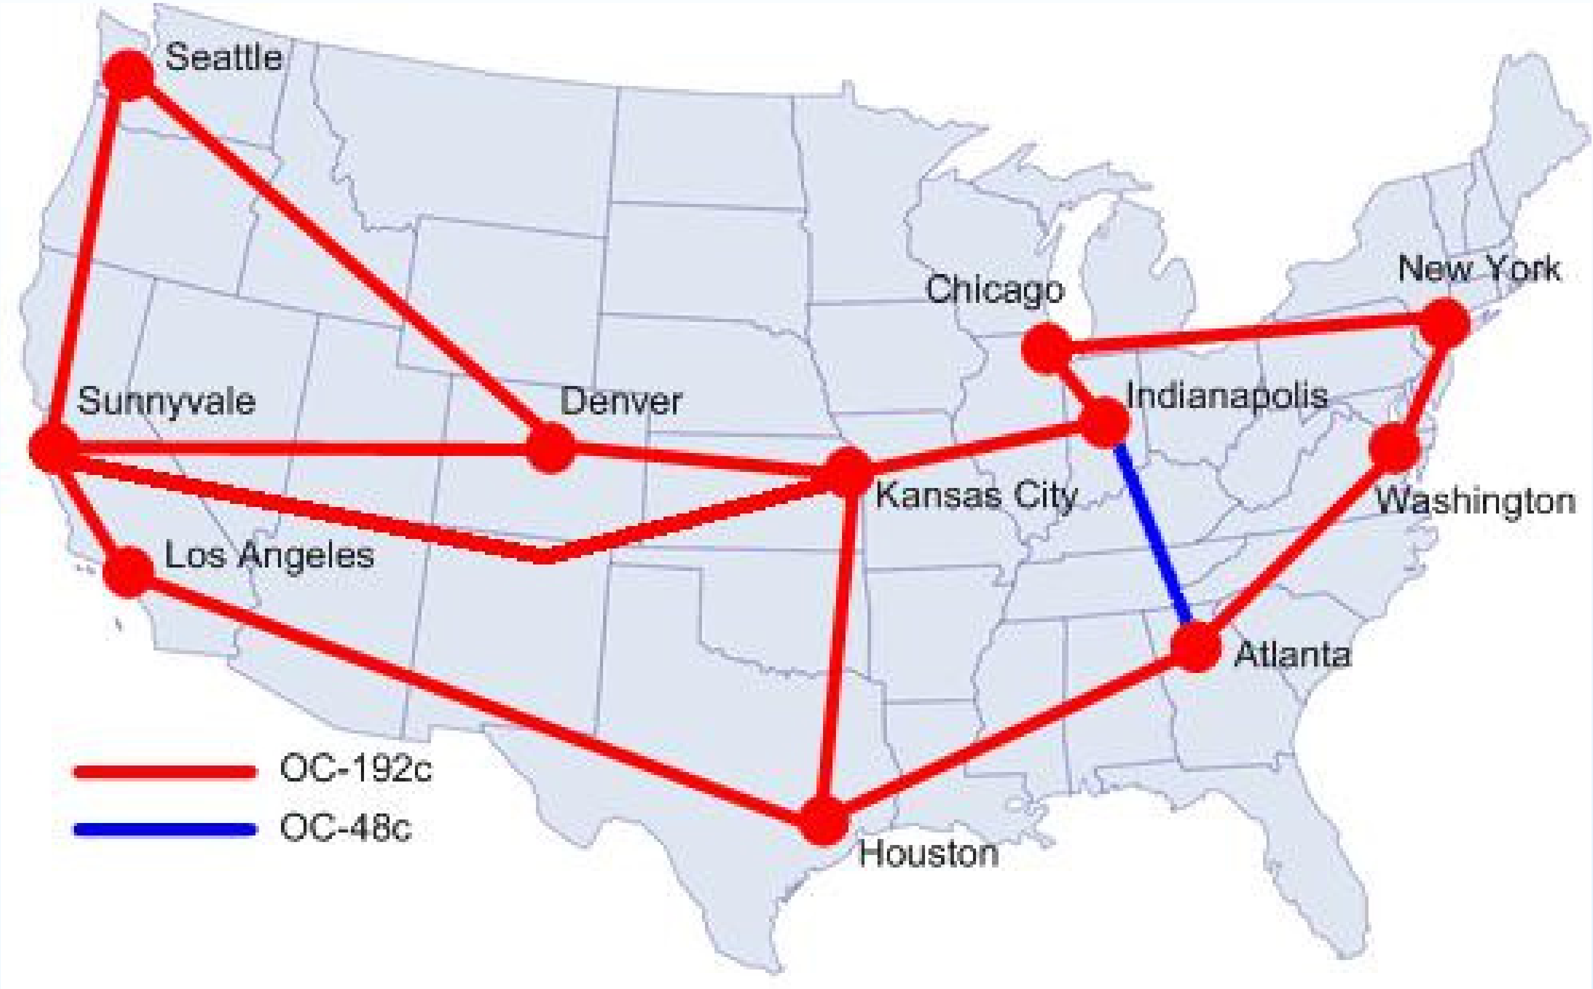

Previous Work

Following seminal work of Shavitt *et al* '04, Chen *et al.* '03 & '04 have shown that only a number

$$\text{Rank}(G) = O(n_e)$$

of paths need be measured for *exact* recovery of y.

Proof follows from basic linear algebraic arguments, and requires simply that a spanning subset of rows of G be chosen and y_i measured on the corresponding paths.

Subset selection algorithm, from computational linear algebra, is used to find a suitable subset.

Our Work

Our contribution to this area is a framework in which substantial additional savings can be achieved by allowing for the *approximate* monitoring of path metrics.

Based on two key components:

- Exploitation of low *effective* rank of G in real networks;
- adaptation of linear predictive statistical methods to the problem.

Illustration of Savings

Network	n_p	$\mathrm{Rank}(G)$	Rel-RMSPE $< 10\%$
Abilene	110	30	8
AS 1221	10716	306	65
AS 1755	7482	318	61
AS 3967	6162	280	66

Remainder of this Talk

1. Network Kriging Framework
   - Path redundancy
   - Statistical prediction

2. Evaluation
   - Theoretical evaluation
   - Empirical evaluation using Abilene delay data

3. Discussion

Quantifying Path Redundancy: SVD

- Higher degree of path redundancy implies less unique paths to measure.
- Similar paths have ‘more similar’ rows in G.
- Similarity of rows suggests reduced dimensionality for G.
- Dimensionality can be assessed through the SVD of G.

Interpretation of SVD of G

- SVD can be seen as deriving from eigen-analysis of $B = G^T G$.
- Largest eigenvalue is that 'direction' in space of link measurements x that maximizes path volume y i.e.,

$$
\begin{aligned}
v_1 &= \arg \max_{x:||x||=1} x^T B x \\
&= \arg \max_{x:||x||=1} (Gx)^T (Gx) \\
&= \arg\max y^T(x) \cdot y(x)
\end{aligned}
$$

Illustration: Eigen-analysis for Abilene

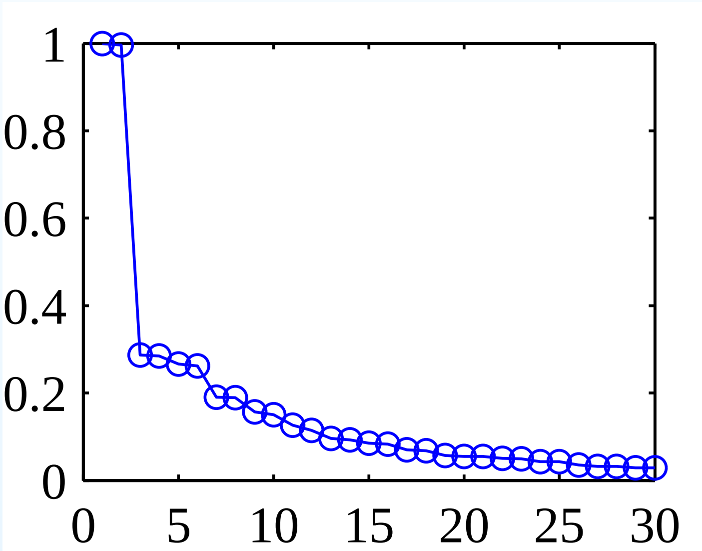

Eigenvalues of $G^T G$

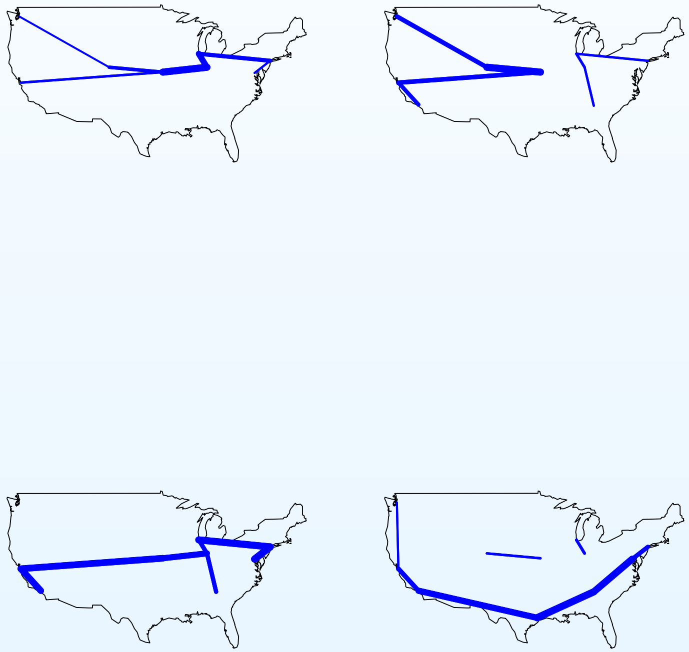

Visual display of energy in eigen-vectors 1, 3, 5, and 7.

Confirmation on Rocketfuel Topologies

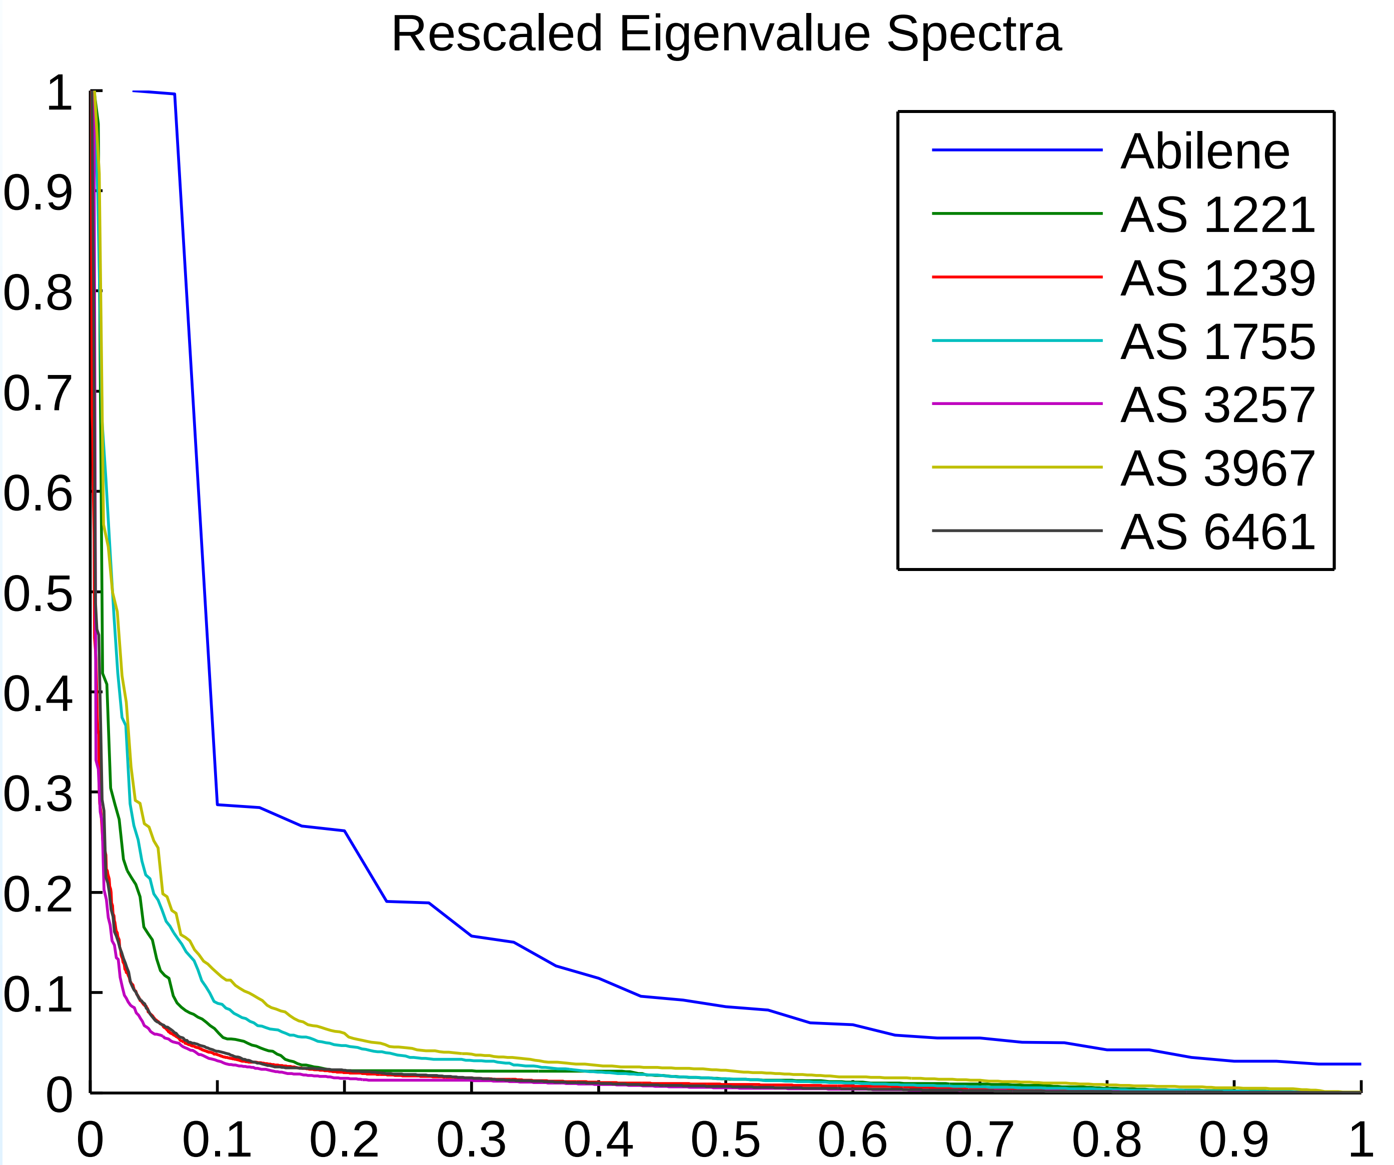

A Partial Characterization

- Let $B = G^T G$, sorted to have non-increasing diagonal i.e., $b_{11} \geq b_{22} \geq \ldots \geq b_{n_e,n_e}$.
- Denote the eigenvalues of B by $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_{n_e}$.
- THEOREM: For $k > 1$

$$\frac{\lambda_k}{\lambda_1} \leq \frac{b_{kk}}{b_{11}} \text{ diam}(\mathcal{G}),$$

  where $\mathcal{G}$ is the underlying network graph.

$\Rightarrow$ Spectral decay parallels that of the 'edge betweenness' of $\mathcal{G}$.

Network Monitoring as Statistical Prediction

Goal: Accurate prediction of a linear summary $l^T y$ of path conditions y, based on measurements from a subset of paths.

Examples:

- Networks Average: $l_i \equiv 1/n_p$
- Difference of Subnetwork Averages:

$$l_i = \begin{cases} 1/|\mathcal{P}_1| & \text{if } i \in \mathcal{P}_1 \\ -1/|\mathcal{P}_2| & \text{if } i \in \mathcal{P}_2 \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Notation

- $\mu = E[x]$ and $\Sigma =$ Cov(x)
  $\Rightarrow \nu = G\mu$ and $V = G\Sigma G^T$
- $y_s = (y_{i_1}, \ldots, y_{i_k})'$ is the value of our metrics on the k paths measured i.e., $i_1, \ldots, i_k \in \mathcal{P}$
- y_r is the value on the remaining $n_p - k$ paths

So rewrite $y = (y_s, y_r)'$ and $G = [G_s, G_r]'$.

Notation (cont)

Similarly

$$\nu = \begin{bmatrix} \nu_s \\ \nu_r \end{bmatrix} = \begin{bmatrix} G_s \mu \\ G_r \mu \end{bmatrix} \qquad (2)$$

and

$$V = \begin{bmatrix} V_{ss} & V_{sr} \\ V_{rs} & V_{rr} \end{bmatrix} = \begin{bmatrix} G_s \Sigma G_s^T & G_s \Sigma G_r^T \\ G_r \Sigma G_s^T & G_r \Sigma G_r^T \end{bmatrix}. \qquad (3)$$

Optimal Prediction

Gauge the quality of a predictor $p(y_s)$ by its mean squared prediction error (MSPE) i.e., $E[(l^T y - p(y_s))^2]$.

Optimal predictor given by

$$E[l^T y | y_s] = l_s^T y_s + E[l_r^T y_r | y_s], \tag{4}$$

where l_s and l_r are defined similarly to y_s and y_r, respectively.

NOTE: Requires information on joint (y_s, y_r) distribution.

Optimal Linear Prediction

Restrict attention to linear predictors of the form $b^T y_s$, in which case the best linear predictor (BLP) is

$$a^T y_s = l_s^T y_s + l_r^T G_r \mu + l_r^T c_* (y_s - G_s \mu), \quad (5)$$

where c_* is any solution to $c_* V_{ss} = V_{rs}$.

NOTE: Requires knowledge of μ.

Our Predictor

We estimate μ from y_s using generalized least squares

$$\hat{\mu} = [G_s^T V_{ss}^{-1} G_s]^- G_s^T V_{ss}^{-1} y_s. \tag{6}$$

Substitution yields

$$\begin{aligned} \hat{a}^T y_s &= l_s^T y_s + l_r^T G_r [G_s^T V_{ss}^{-1} G_s]^- G_s^T V_{ss}^{-1} y_s \\ &= l_s^T y_s + l_r^T V_{rs} V_{ss}^{-1} y_s. \end{aligned} \tag{7}$$

[NOTE: $\hat{a}^T y_s$ is *not* a BLUP, due to deficiency of rank.]

Path Selection

Question: How do we choose the k paths $i_1, \ldots, i_k \in \mathcal{P}$?

A natural answer is to choose a set of k paths for whom the $\mathrm{MSPE}(\hat{a}^T y_s)$ is minimized.

Can write

$$\mathrm{MSPE}(\hat{a}^T y_s) = \mathrm{MSPE}(a^T y_s) + (\mathrm{Bias}\, \hat{a}^T y_s)^2 \tag{8}$$

Path Selection (cont.)

$$\mathrm{MSPE}(a^T y_s) = l_r^T (G_r C)(I - B_s)(G_r C)^T l_r, \tag{9}$$

where

$$\begin{aligned} B_s &= (G_s C)^T [(G_s C)(G_s C)^T]^{-1} (G_s C) \\ &= C^T G_s^T V_{ss}^{-1} G_s C \end{aligned} \tag{10}$$

is the orthog projection onto $\mathrm{Row}(G_s C)$ and $\Sigma = CC^T$.

INTERPRETATION: Select a set of k paths for whom the span of $G_s C$ is as representative of GC as possible.

Path Selection Algorithm

Choosing paths to minimize $\mathrm{MSPE}(a^T y_s)$ is equivalent to a version of the 'subset selection' problem in computational linear algebra.

Exact solution is NP-complete, algorithms for numerous approximate solutions exist.

IDEA: Choose k paths whose rows in G span the first k singular dimensions of G to the greatest degree possible.

MSPE Properties: Theoretical

Theorem 1 *Denoting the* i^{th} *row of* G *as* $G_{(i)}$ *, let* $p_i = \left\|G_{(i)}\right\|_2^2 / \|G\|_F^2$ *, where* $\|\cdot\|_2$ *and* $\|\cdot\|_F$ *are the Euclidean and Frobenius matrix norms, respectively. Let* $\tilde{G}_s$ *be a rescaled version of* G_s *, under the operation* $G_{(i)} \to G_{(i)}/\sqrt{kp_i}$ *for each of the* k *rows in* G_s*. Then if*

$$\left\|G^T G - \tilde{G}_s^T \tilde{G}_s\right\|_F \leq f(k) \|G\|_F^2 \quad , \tag{11}$$

for some $f(k)$*, the MSPE can be bounded as*

$$\mathrm{MSPE}(\hat{a}^T y_s) \leq (\|\mu\|_2^2 + 1) \left(\lambda_{k+1} + 2f(k) \|G\|_F^2\right) \|l\|_2^2 . \tag{12}$$

MSPE Properties: Numerical

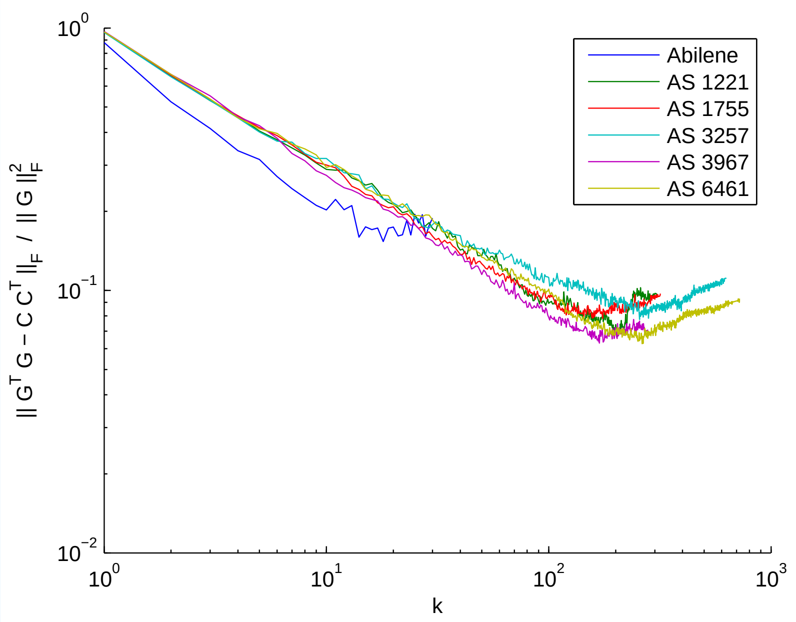

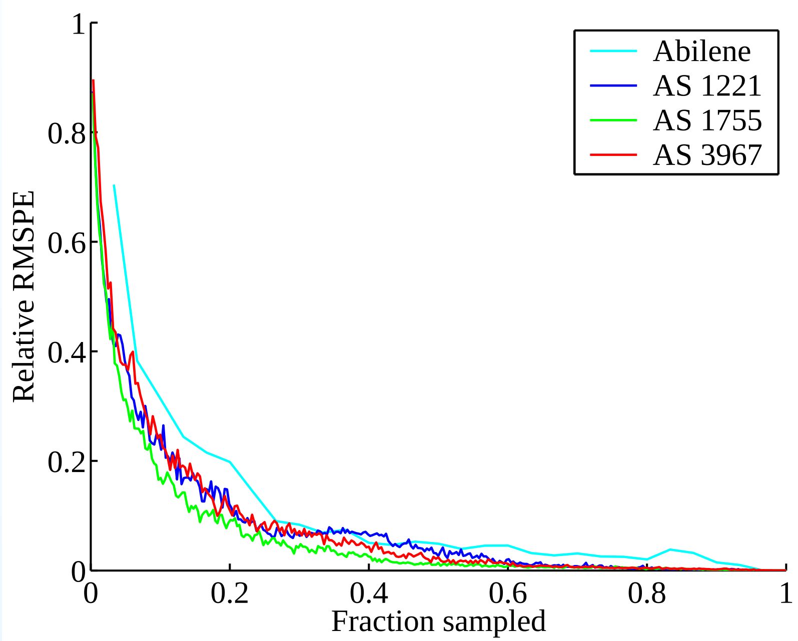

Left: Empirical calculations showing $f(k) \sim k^{-1/2}$.

Right: Computed MSPE for simple packet-loss model.

MSPE Properties: Numerical (cont)

Table 1: Comparison of Paths Measured

Network	n_p	$\mathrm{Rank}(G)$	Rel-RMSPE $<$ 10%
Abilene	110	30	8
AS 1221	10716	306	65
AS 1755	7482	318	61
AS 3967	6162	280	66

An Extreme Case: Monitoring of Individual Paths

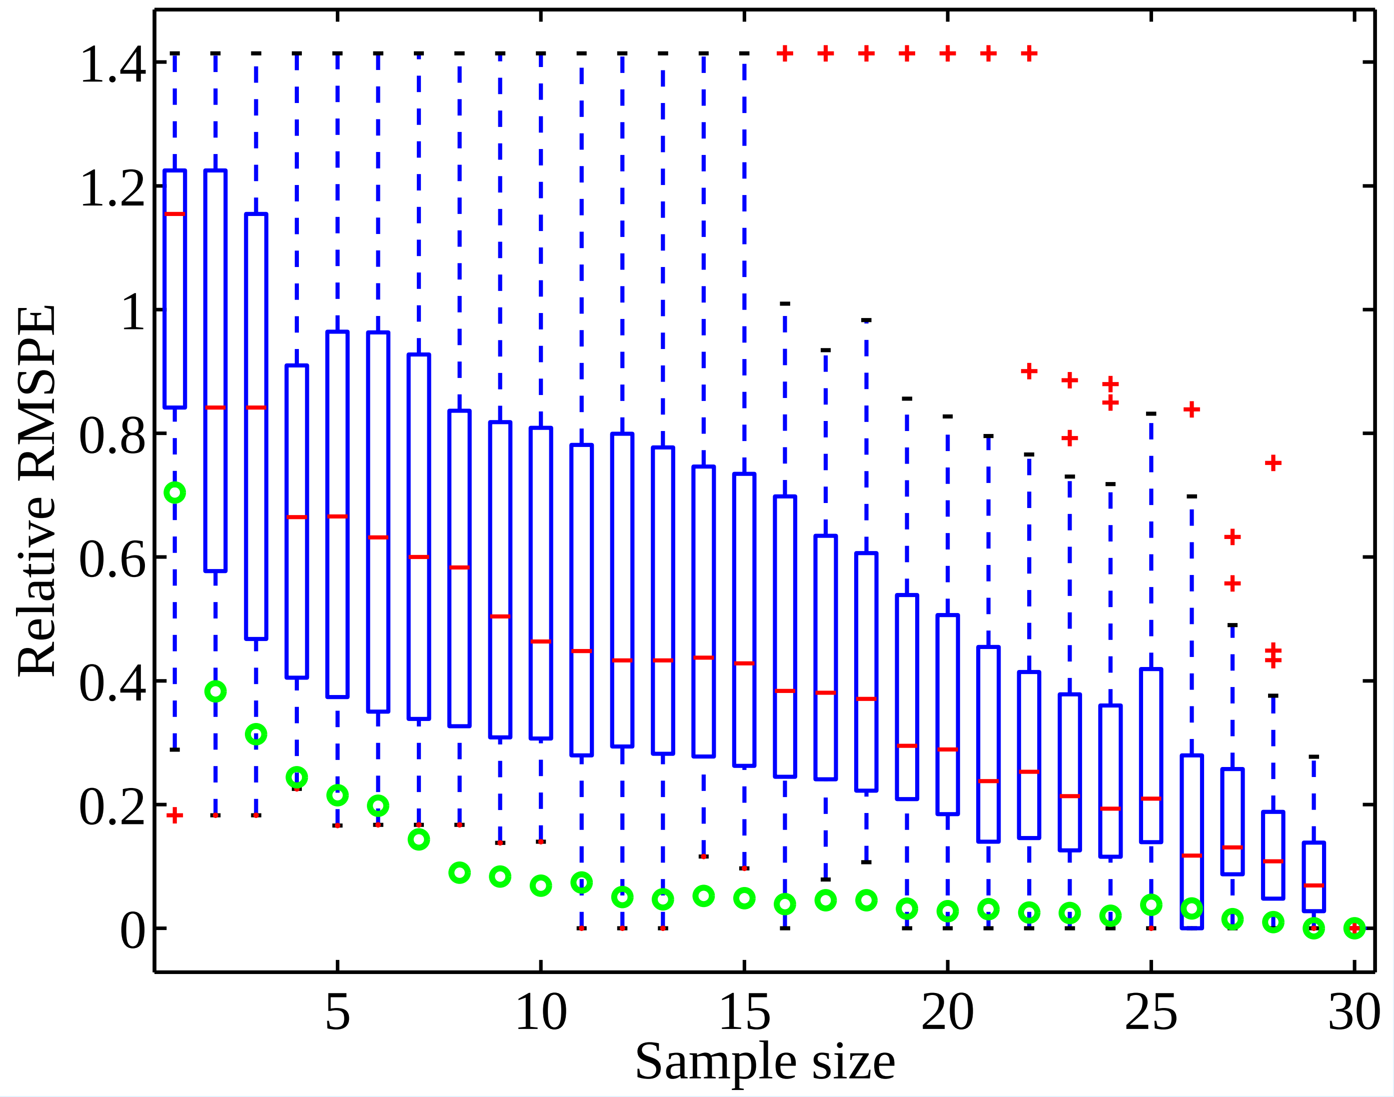

A Theoretical Aside: Random Path Selection

Theorem 2 *Let $\tilde{G}_s$ be a matrix constructed as before, but now from c paths randomly selected with respect to the probabilities $\{p_i\}$.*

If the estimator $\hat{a}^T y_s$ is constructed based on the first (at most) $k \leq c$ singular dimensions of $\tilde{G}_s$, then for any $c \leq n_p$ and $\delta > 0$, the MSPE is bounded by

$$(\|\mu\|_2^2 + 1)(\lambda_{k+1} + 2(1 + \sqrt{\ln(2/\delta)})c^{-\frac{1}{2}} \, \|G\|_F^2) \, \|l\|_2^2$$

with probability at least $1 - \delta$.

Empirical Evaluation: Abilene Delay Data

- Traceroute data from NLANR Active Measurement Project
- 10 minute intervals over 3 days in 2003, yielding 432 time points
- 14917 OD pairs reduced to ingress/egress points on Abilene
- Edge delays ranged from 2 to 36 milliseconds, with standard deviations 0.16 to 0.94.
- Data suggest a diagonal covariance; entries estimated from one day of data.

Effect of Link Variance: Spectral Decay

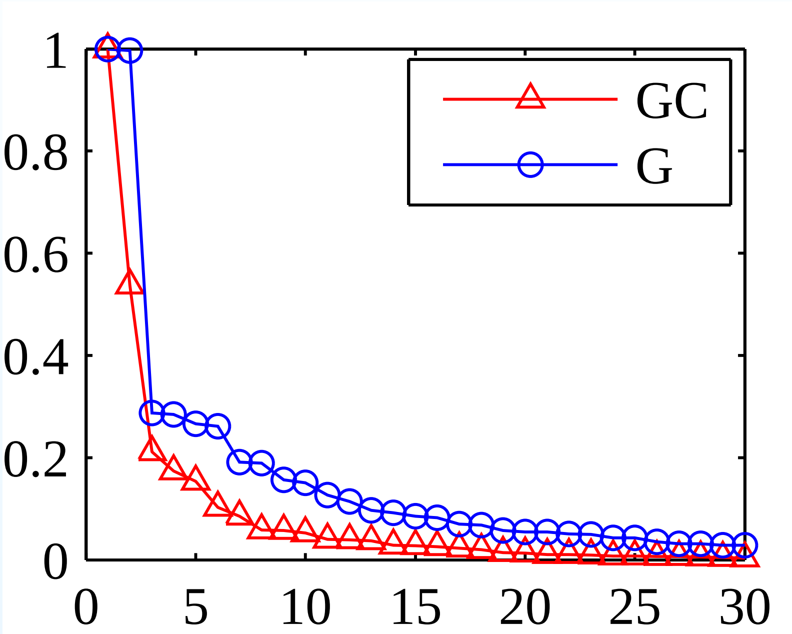

- Spectra of G versus GC shows that reduced effective rank persists.

Effect of Link Variance: Path Selection

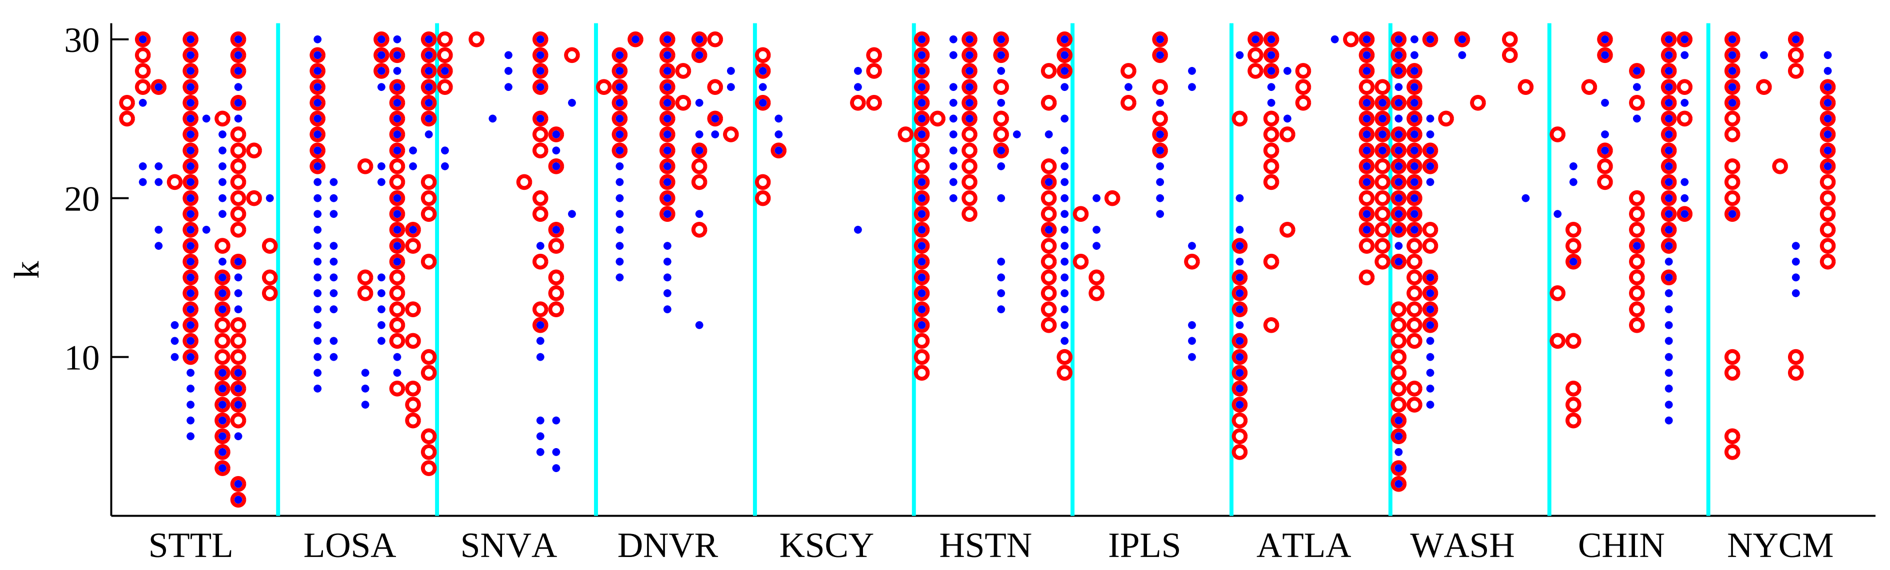

- Selection of paths under G (red) and GC (blue).
- Grouped by egress point, from west to east, and ordered similarly within groups.
- Shows a combination of stability and flexibility

Robustness to Link Failure

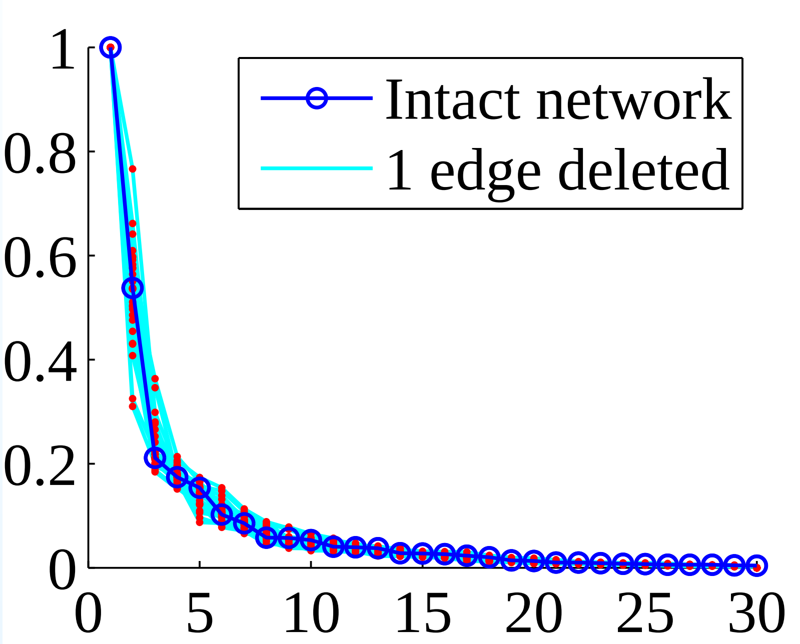

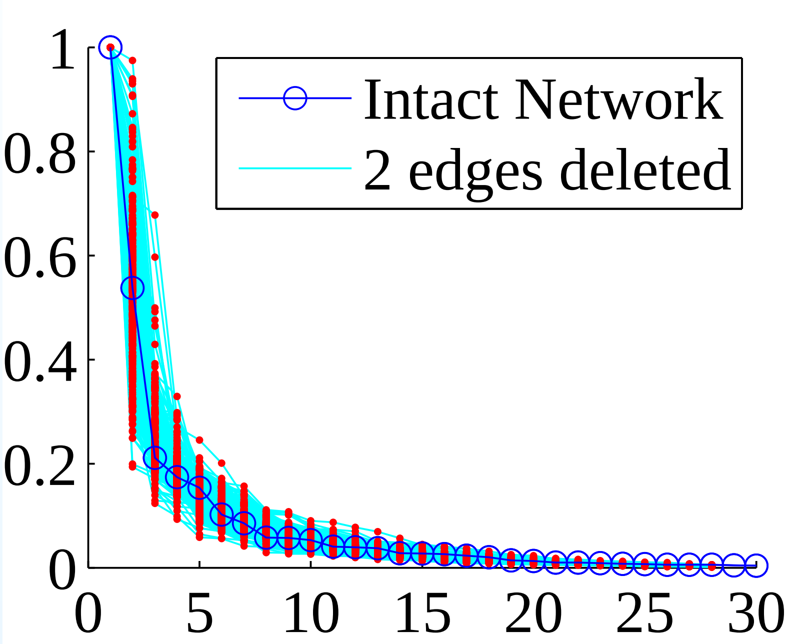

- Spectral decay remains robust to deletion of 1, and even 2, links out of 30.

Applications

Three applications of interest to network operators and users:

- Network 'health': Monitoring of network-average delay.
- Anomaly Detection: Flagging spikes in delay.
- Subnetwork Comparison: Monitoring delay differences from two points of entry into the network.

Application 1: Monitoring Global Delay

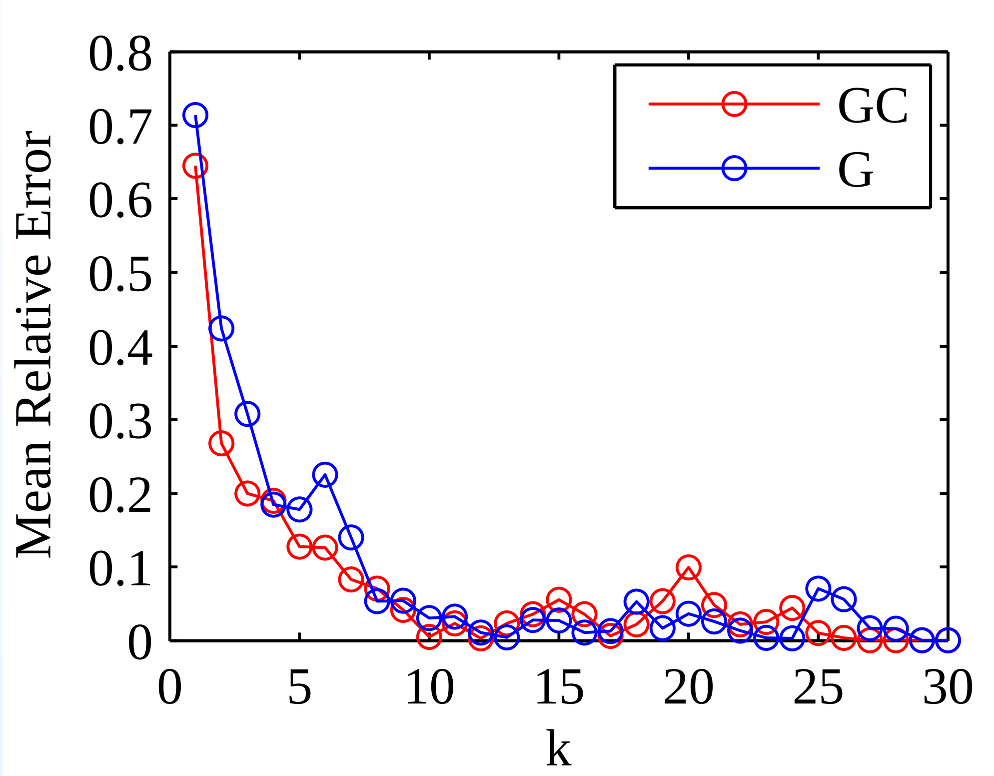

- Characteristics of (time) averaged relative error mimic those of the underlying spectra.

Bias Correction

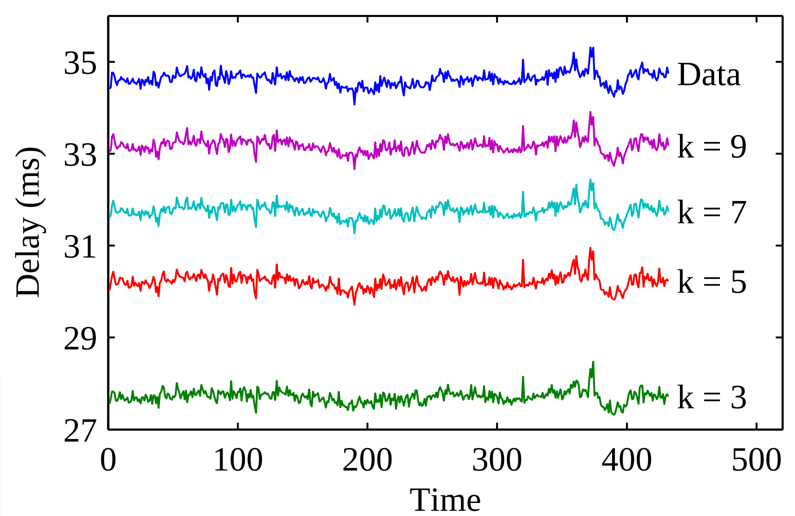

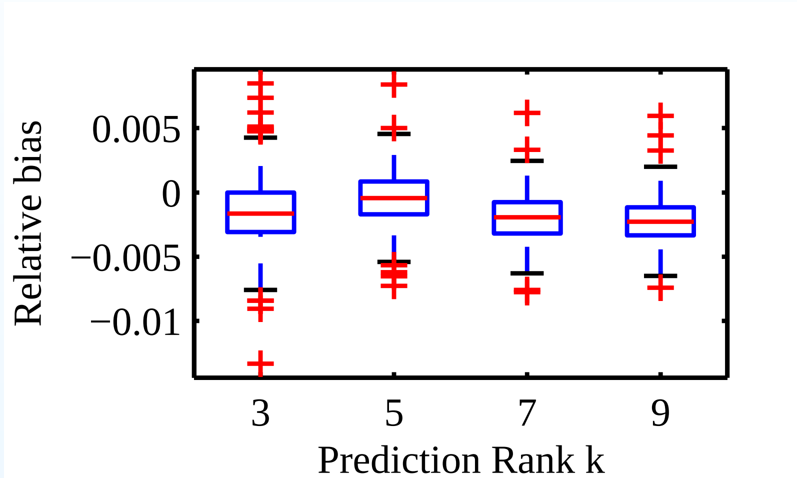

- Predictions show increasing bias with increasing k.
- Due to lack of information on unseen links.
- Correction can be made through use of link-level measurements at $t = 1$.
- Relative errors now on the order of 1%, even for $k = 3$!

Application 2: Anomaly Detection

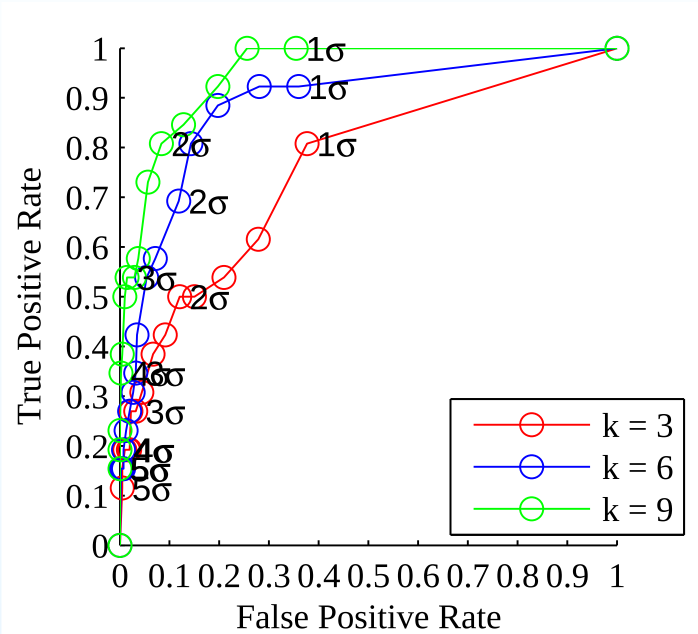

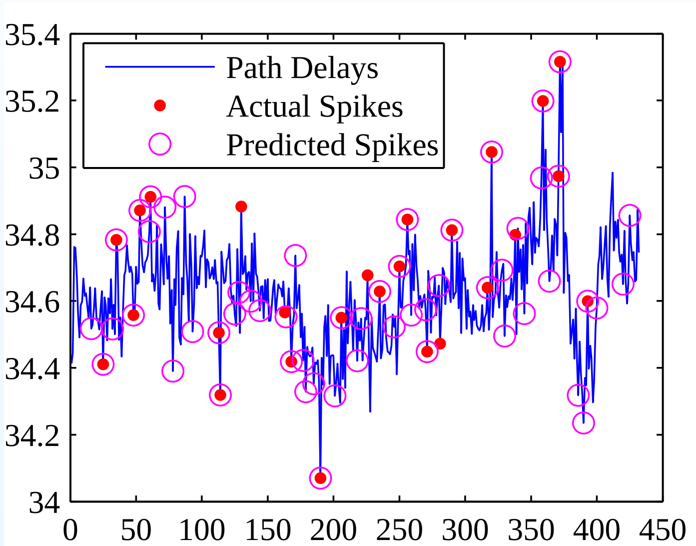

- Potential ‘anomaly’ defined in true network average if value at time t exceeds 3 standard deviations of the past 6 epochs.
- ROC curve considers effect of both choice of k and choice of σ used in flagging anomalies in predicted time series.
- True Pos Rate = 0.81 and False Pos Rate = 0.08

Application 3: Multi-Homing

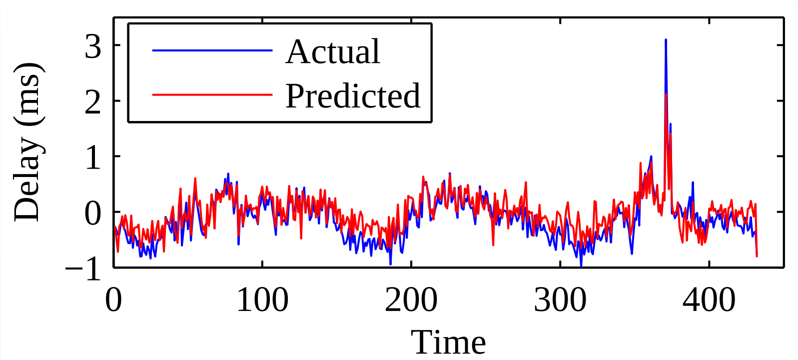

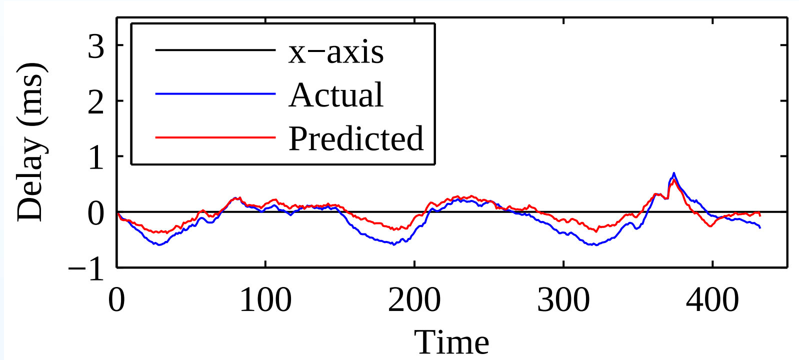

- Comparison of subnetworks originating from Chicago and Atlanta.
- Correlation of 0.866 between predicted and actual differences in average subnetwork path delay.
- Accuracy of 88% in predicting subnetwork with shorter mean delay.

Wrap Up

- Network-wide monitoring from a small collection of measurements a fundamental goal.
- Our method combines methods from
  statistical sampling and subset selection
  to exploit inherent network redundancies.
- Theoretical analyses cast some light on both redundancies and performance of the method.
- Empirical illustrations demonstrate applicability.

See Coates, Pointurier, and Rabbat (2007) for a 'spatio'-temporal extension of this work.

References

- Coates, M., Pointurier, Y., and Rabbat, M. (2007). Compressed Network Monitoring for IP and All-Optical Networks. *Proceedings of the ACM Internet Measurement Conference,* San Diego, California, October, 2007.
- Chua, D.B., Kolaczyk, E.D., and Crovella, M. Efficient monitoring of end-to-end network properties. Proceedings of the IEEE Infocom 2005.
- Chua, D.B., Kolaczyk, E.D., and Crovella, M. (2006). Network Kriging. *IEEE Journal on Selected Areas of Communication (Special Issue on Sampling the Internet),* **24:12**, 2263-2272.

See *http://math.bu.edu/people/kolaczyk/pubs.html*