Analysis of Network Flow Data

Gonzalo Mateos
Dept. of ECE and Goergen Institute for Data Science
University of Rochester
gmateosb@ece.rochester.edu
http://www.ece.rochester.edu/~gmateosb/

April 26, 2016

Network flows

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Traffic flows

- Networks often serve as conduits for traffic flows

Example

- Commodities and people flow over transportation networks;
- Data flows over communication networks; and
- Capital flows over networks of trade relations

- Flow-related questions on network design, provisioning and routing
  $\Rightarrow$ Solutions involve tools in optimization and algorithms
- **Our focus:** statistical analysis and modeling of network flow data
  $\Rightarrow$ Regression-based prediction of unknown flow characteristics

Routing matrix

- Let $G(V, E)$ be a digraph. Flows are directed: origin $\rightarrow$ destination
  - $\Rightarrow$ Directed edges (arcs) here referred to as links
  - $\Rightarrow$ Number of flows is N_f, typically have $N_f = O(N_v^2)$
  - $\Rightarrow$ Flows traverse multiple links en route to their destinations
- Routing matrix $\mathbf{R} \in \{0, 1\}^{N_e \times N_f}$ states incidence of routes with links

$$r_{e,f} = \begin{cases} 1, & \text{if flow } f \text{ routed via link } e, \\ 0, & \text{otherwise} \end{cases}$$

- Assumed flows follow a single route from origin to destination

Example: Routing of two flows

Ex: Consider a digraph with $N_e = 7$ links and $N_f = 2$ active flows

$$\mathbf{R} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \end{pmatrix}$$

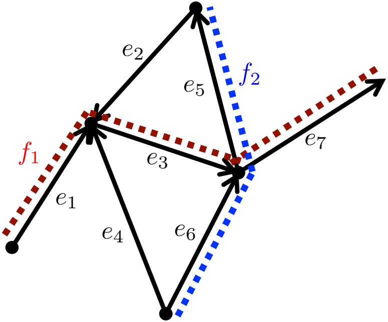

- Strongly connected digraph: flows can be as many as $N_v(N_v - 1)$

Traffic matrix

- Central to study of network flows is the traffic matrix $\mathbf{Z} \in \mathbb{R}^{N_v \times N_v}$
  - Entry z_{ij} is total volume of flow from origin vertex i to destination j
- Ex: net out-flow from i and net in-flow to j given by

$$z_{i+} = \sum_j z_{ij} \quad \text{and} \quad z_{+j} = \sum_i z_{ij}$$

- Link-level aggregate traffic vector $\mathbf{x} := [x_1, \ldots, x_{N_e}]^T$ related to $\mathbf{Z}$ as

$$\mathbf{x} = \mathbf{R}\mathbf{z}, \qquad \text{where } \mathbf{z} := \text{vec}(\mathbf{Z})$$

$\Rightarrow$ Link counts x_e equal the sum of flow volumes routed through e

Flow costs and time dependencies

- Notion of cost c associated with paths or links also important

  Ex: generalized socioeconomic cost for transportation analysis

  $\Rightarrow$ Study choices made by consumers of transportation resources

  Ex: quality of service (QoS) in network traffic analysis

  $\Rightarrow$ Monitor delays to unveil congestion or anomalies

- Implicitly assumed a static snapshot taken of the network flows

  $\Rightarrow$ Flows dynamic in nature. Time-varying models more realistic

  $\Rightarrow$ When appropriate will denote $\mathbf{x}(t)$, $\mathbf{Z}(t)$ or $\mathbf{R}(t)$

- Common assumption to treat routing matrix $\mathbf{R}$ as being fixed

  $\Rightarrow$ Routing changes at slower time scale than flow dynamics

Example: Internet2 traffic matrix

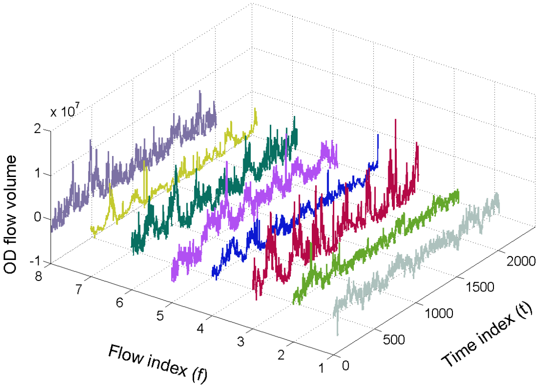

- Internet2 backbone: $N_f = 110$ flows (8 shown) over a week

  $\Rightarrow$ Temporal periodicity and "spatial" correlation apparent

Roadmap

- Roadmap dictated by types of measurement and analysis goal

- Measure: origin-destination (OD) flow volumes z_{ij} in full
- Goal: model flows to understand and predict future traffic

  $\Rightarrow$ Gravity models

- Measure: link counts x_e, flow volumes unavailable
- Goal: traffic matrix estimation, i.e., predict unobserved OD flows z_{ij}

  $\Rightarrow$ Gaussian and Poisson models, entropy minimization

- Measure: OD costs c_{ij} for a subset of paths
- Goal: predict unobserved OD and link costs

  $\Rightarrow$ Active network tomography and network kriging

Gravity models

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Gravity models

- Gravity models originate in the social sciences [Stewart '41]
  - $\Rightarrow$ Describe aggregate level of interactions among populations
- Ex: geography, economics, sociology, hydrology, computer networks
- Newton's law of gravitation for masses m_1, m_2 separated by d_{12}

$$F_{12} = G\frac{m_1 m_2}{d_{12}^2}$$

- Gravity models specify interactions among populations vary:
  - $\Rightarrow$ In direct proportion to the population's sizes; and
  - $\Rightarrow$ Inversely with some measure of their separation
- Intuition: OD flows as "population interactions", makes sense!

Model specification

- Sets of origins $\mathcal{I}$ and destinations $\mathcal{J}$. Flows Z_{ij} from $i \in \mathcal{I}$ to $j \in \mathcal{J}$
- Gravity models state Z_{ij} are independent, Poisson, with mean

$$\mathbb{E}\left[Z_{ij}\right] = h_O(i) h_D(j) h_S(\mathbf{c}_{ij})$$

$\Rightarrow$ Origin $h_O(\cdot)$, destination $h_D(\cdot)$, and separation function $h_S(\cdot)$

$\Rightarrow$ "Distance" between i, j captured by separation attributes $\mathbf{c}_{ij}$

- Ex: Stewart's theory of demographic gravitation specifies

$$\mathbb{E}\left[Z_{ij}\right] = \gamma \pi_{O,i} \pi_{D,j} d_{ij}^{-2}$$

$\Rightarrow$ Population sizes measured by $\pi_{O,i}$ and $\pi_{D,j}$, distance by d_{ij}

$\Rightarrow$ Demographic gravitational constant γ

- Unlike Netwon's law, no empirical or theoretical support here

Origin, destination and separation functions

- Multiple origin, destination and separation functions proposed

  $\Rightarrow$ Motivated from sociophysics and economic utility theory

- Ex: power functions for $h_O(i)$ and $h_D(j)$, where for $\alpha, \beta \geq 0$

$$h_O(i) = (\pi_{O,i})^{\alpha} \quad \text{and} \quad h_D(j) = (\pi_{D,j})^{\beta}$$

- Ex: power function $h_S(c_{ij}) = c_{ij}^{-\theta}$, $\theta \geq 0$. General exponential form

$$h_S(\mathbf{c}_{ij}) = \exp(\boldsymbol{\theta}^T \mathbf{c}_{ij}), \quad \boldsymbol{\theta}, \mathbf{c}_{ij} \in \mathbb{R}^K$$

- Convenient for inference of model parameters, since

$$\log \mathbb{E}\left[Z_{ij}\right] = \log \gamma + \alpha \log \pi_{O,i} + \beta \log \pi_{D,j} + \boldsymbol{\theta}^T \mathbf{c}_{ij}$$

  $\Rightarrow$ Log-linear form facilitates standard regression software

Example: Austrian phone-call data

- Q: Structure of telecommunication interactions among populations?
  - $\Rightarrow$ Planning for government (de)regulation of the sector
  - $\Rightarrow$ Predict influence of technologies in regional development
- Gravity models to model telecommunication patterns as flows
- Data for phone-call traffic among 32 Austrian districts in 1991
  - $\Rightarrow$ $32 \times 31 = 992$ flow measurements z_{ij}, $i \neq j = 1, \ldots, 32$
  - $\Rightarrow$ Gross regional product (GRP) per region $\rightarrow$ Size proxy
  - $\Rightarrow$ Road-based distance among regions $\rightarrow$ Separation proxy

Phone-call data scatterplots

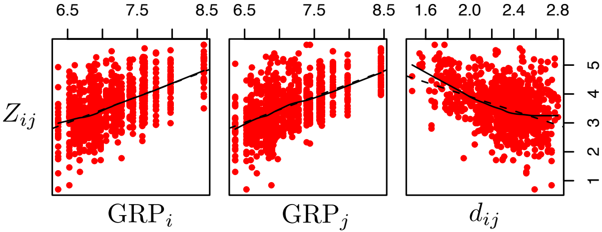

- Data (in $\log_{10}$ scale) suggest a gravity model of the form

$$\mathbb{E}[Z_{ij}] = \gamma(\pi_{O,i})^{\alpha}(\pi_{D,j})^{\beta}(c_{ij})^{-\theta}$$

$\Rightarrow \pi_{O,i} = \mathrm{GRP}_i$, $\pi_{D,j} = \mathrm{GRP}_j$, $c_{ij} = d_{ij}$ i-j's road-based distance

- Typical that flow volumes vary widely in scale

Inference for gravity models

- Specified Z_{ij} as independent Poisson RVs, with means $\mu_{ij} = \mathbb{E}[Z_{ij}]$

  $\Rightarrow$ ML for statistical inference in the general gravity model

- Let $\alpha_i = \log h_O(i)$, $\beta_i = \log h_D(j)$ and $\boldsymbol{\theta} \in \mathbb{R}^K$. Will focus on

$$\log \mu_{ij} = \alpha_i + \beta_j + \boldsymbol{\theta}^T \mathbf{c}_{ij}$$

  $\Rightarrow$ Log-linear model $\in$ class of generalized linear models

- P. McCullagh and J. Nedler, *Generalized Linear Models*. CRC, 1989

- Given flow observations $\mathbf{Z} = \mathbf{z}$, the Poisson log-likelihood for $\boldsymbol{\mu}$ is

$$\ell(\boldsymbol{\mu}) = \sum_{i,j \in \mathcal{I} \times \mathcal{J}} z_{ij} \log \mu_{ij} - \mu_{ij}$$

  $\Rightarrow$ Substitute the gravity model and maximize $\ell(\boldsymbol{\mu})$ for MLE

ML parameter estimates

- MLEs $\hat{\boldsymbol{\alpha}} := \{\hat{\alpha}_i\}_{i\in\mathcal{I}}$, $\hat{\boldsymbol{\beta}} := \{\hat{\beta}_j\}_{j\in\mathcal{J}}$ and $\hat{\boldsymbol{\theta}}$ satisfy

$$\log \hat{\mu}_{ij} = \hat{\alpha}_i + \hat{\beta}_j + \hat{\boldsymbol{\theta}}^T \mathbf{c}_{ij},\ i,j \in \mathcal{I}\times\mathcal{J} \ \Rightarrow \log \hat{\boldsymbol{\mu}} = \mathbf{M}\hat{\gamma}$$

- Defined $\hat{\gamma} := \left[\hat{\boldsymbol{\alpha}}^T\ \hat{\boldsymbol{\beta}}^T\ \hat{\boldsymbol{\theta}}^T\right]^T$, mean flow estimates $\hat{\mu}_{ij}$ solve

$$\sum_j \hat{\mu}_{ij} = z_{i+},\ i\in\mathcal{I} \quad \text{and} \quad \sum_i \hat{\mu}_{ij} = z_{+j},\ j\in\mathcal{J}$$

$$\sum_{i,j} \mathbf{c}_{ij}(k)\hat{\mu}_{ij} = \sum_{i,j} \mathbf{c}_{ij}(k) z_{ij}, \quad k = 1,\ldots,K$$

- Unique MLE $\hat{\boldsymbol{\theta}}$ under mild conditions, e.g., $\text{rank}(\mathbf{M}) = I + J + K - 1$

  $\Rightarrow$ Values $\hat{\alpha}_i$, $\hat{\beta}_j$ unique only up to a constant

- A. Sen, "Maximum likelihood estimation of gravity model parameters," *J. Regional Science*, vol. 26, pp. 461-474, 1986

LS parameter estimates

- LS procedures the norm early on, based on models

$$\log Z_{ij} \approx \alpha_i + \beta_j + \boldsymbol{\theta}^T \mathbf{c}_{ij} + \epsilon_{ij}, \quad i,j \in \mathcal{I} \times \mathcal{J}$$

- Beware: ordinary LS estimation doomed to yield poor results
  - $\Rightarrow$ Biased estimates, $\mathbb{E}[\log Z_{ij}] \leq \log \mu_{ij}$ by Jensen's inequality
  - $\Rightarrow$ Variance not constant, $\text{var}[\log Z_{ij}]$ depends on μ_{ij}
- Corrective measures: replace $\log Z_{ij} \leftrightarrow \tilde{Z}_{ij} := \log(Z_{ij} + 1/2)$
  - $\Rightarrow \mathbb{E}\left[\tilde{Z}_{ij}\right] = \log \mu_{ij}$ and $\text{var}\left[\tilde{Z}_{ij}\right] = \mu_{ij}^{-1}$ up to $O(\mu_{ij}^{-2})$ terms
  - $\Rightarrow$ Use weighted LS with $w_{ij} \propto \mu_{ij}^{1/2}$ (start with $z_{ij}^{1/2}$, then $\hat{\mu}_{ij}^{1/2}$)
- LS is simple, but all things being equal ML is preferable

Example: Analysis of Austrian phone-call data

- Given phone-call data, form MLEs of parameters in two models

$$\text{Standard gravity model: } \mu_{ij} = \gamma(\pi_{O,i})^{\alpha}(\pi_{D,j})^{\beta}(c_{ij})^{-\theta}$$

$$\text{General gravity model: } \log \mu_{ij} = \alpha_i + \beta_j - \theta c_{ij}$$

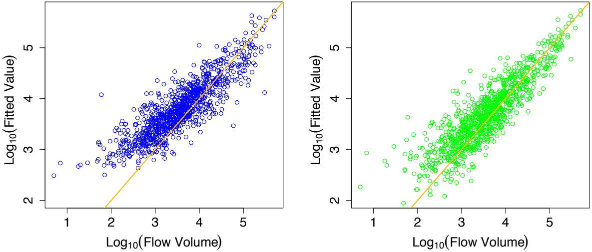

- Prediction of traffic flows. Plot $\hat{\mu}_{ij}$ vs z_{ij} in log-log scale
  - $\Rightarrow$ Fairly linear trend for both gravity models
  - $\Rightarrow$ Standard model tends to over-estimate low-volume flows

Relative prediction error

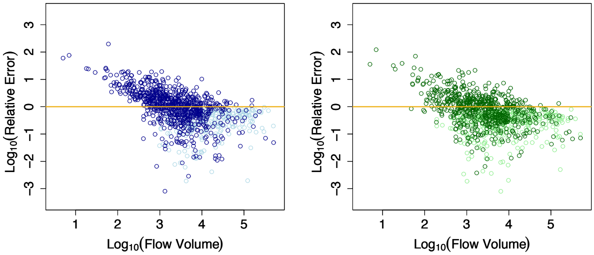

- Relative prection error. Plot $(z_{ij} - \hat{\mu}_{ij})/z_{ij}$ vs z_{ij} in log-log scale

  $\Rightarrow$ For both models error varies widely in magnitude

  $\Rightarrow$ Roughly, error decreases with flow volume

  $\Rightarrow$ Tendency to over- (under)-estimate low (high) volumes

Model accuracy comparison

- Plot empirical CDF of models' relative prediction errors

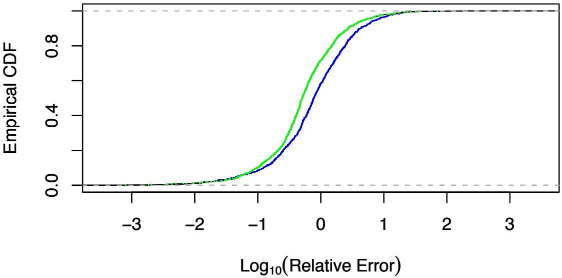

- General model's CDF lies to the left of that for the standard model

  $\Rightarrow$ The general model dominates in terms of accuracy

- Ex: Standard model errors $\leq z_{ij}$ for 58% of the OD pairs

  $\Rightarrow$ Compare with 72% under the general model

Estimating traffic matrices

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Monitoring flows

- Monitoring OD flow volumes Z_{ij} fundamental to:
  - $\Rightarrow$ Traffic management
  - $\Rightarrow$ Network provisioning
  - $\Rightarrow$ Planning for network growth
- Often difficult (even impossible) to measure the Z_{ij}...

  Ex: large-scale surveys prohibitive in transportation networks

  Ex: flow sampling, storing, transmission affects Internet user QoS
- ...but relatively easy to acquire link counts X_e

  Ex: highway networks, place sensors in on- and off-ramps

  Ex: routers monitor data on incident links (e.g., SNMP)

Traffic matrix estimation

Traffic matrix estimation

Given $\mathbf{R}$ and link counts $\{X_e\}_{e \in E}$, predict flows Z_{ij} (or estimate μ_{ij})

- Highly underdetermined inverse problem. "Invert" known fat $\mathbf{R}$ in

$$\mathbf{X} = \mathbf{RZ}, \quad \text{where } \mathbf{R} \in \{0,1\}^{N_e \times N_f} \text{ and } N_e \ll N_f = O(N_v^2)$$

  $\Rightarrow$ Leverage side information to constrain the solution set

- Also dubbed network tomography. Taxonomy of methods:
  - $\Rightarrow$ Static: estimate $\mathbf{Z}$ for a single time period
  - $\Rightarrow$ Dynamic: estimate $\mathbf{Z}$ successively over multiple time periods

- Y. Vardi, "Network tomography: Estimating source-destination traffic intensities from traffic counts," *JASA*, vol. 91, pp. 365-377, 1996

Gaussian models and LSE

- Traffic often has units of "counts" e.g., cars per hour or Mbps

  $\Rightarrow$ Still, early approaches based on LS and Gaussian models

- Simple linear model for observed link counts $\mathbf{X} = \{X_e\}_{e \in E}$

$$\mathbf{X} = \mathbf{R}\boldsymbol{\mu} + \boldsymbol{\varepsilon}$$

  - $\mathbf{R} \in \{0,1\}^{N_e \times N_f}$ is the known routing matrix
  - $\boldsymbol{\mu} \in \mathbb{R}_+^{N_f}$ is vector of expected OD flow volumes
  - $\boldsymbol{\varepsilon}$ is a $N_e \times 1$ vector of i.i.d. zero-mean errors, with variance σ^2

- Formulation suggests estimating $\boldsymbol{\mu}$ via ordinary LS

  $\Rightarrow$ Gaussian $\boldsymbol{\varepsilon}$ reasonable in high-count settings (LS $\Leftrightarrow$ ML)

  $\Rightarrow$ However, typically $N_e \ll N_f$ and LS is poorly posed

Example: Toy network

- Graph $G(V, E)$ with $N_v = 5$ and $N_e = 4$, OD pairs $\{ac, ad, bc, bd\}$

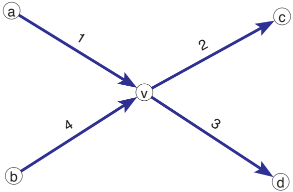

$$\begin{bmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} \mu_{ac} \\ \mu_{ad} \\ \mu_{bc} \\ \mu_{bd} \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \end{bmatrix}$$

- Although $N_e = N_f = 4$, $\text{rank}(\mathbf{R}) = 3$ and $\mathbf{R}^T\mathbf{R}$ not invertible

  $\Rightarrow$ For link counts $\mathbf{X} = \mathbf{x}$, there are infinite solutions $\hat{\boldsymbol{\mu}}$ to

$$\min_{\boldsymbol{\mu}} \|\mathbf{x} - \mathbf{R}\boldsymbol{\mu}\|^2$$

Leveraging historical data

- Suppose we have initial OD flow volume measurements $\mathbf{Z}_0 = \mathbf{z}_0$

  $\Rightarrow$ Historical data, maybe even rough and innacurate

- Use $\mathbf{z}_0$ to constrain the LS problem. Consider the model

$$\begin{bmatrix} \mathbf{Z}_0 \\ \mathbf{X} \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{R} \end{bmatrix} \mu + \begin{bmatrix} \xi \\ \varepsilon \end{bmatrix}$$

  - Independent errors ξ and ε have covariance matrices $\mathbf{\Psi}$ and $\mathbf{\Sigma}$

- Generalized LS estimator

$$\min_{\mu} \begin{bmatrix} \mathbf{z}_0 - \mu \\ \mathbf{x} - \mathbf{R}\mu \end{bmatrix}^T \begin{bmatrix} \mathbf{\Psi}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{z}_0 - \mu \\ \mathbf{x} - \mathbf{R}\mu \end{bmatrix}$$

  $\Rightarrow$ From likelihood-based perspective a Gaussian model implicit

Generalized LS solution

- Generalized LSE is a linear combination of $\mathbf{z}_0$ and $\mathbf{x}$, namely

$$\hat{\boldsymbol{\mu}} = \left(\boldsymbol{\Psi}^{-1} + \mathbf{R}^T\boldsymbol{\Sigma}^{-1}\mathbf{R}\right)^{-1}\left(\boldsymbol{\Psi}^{-1}\mathbf{z}_0 + \mathbf{R}^T\boldsymbol{\Sigma}^{-1}\mathbf{x}\right)$$

- Model is linear so $\hat{\boldsymbol{\mu}}$ is unbiased and a MVUE, with

$$\text{var}\left[\hat{\boldsymbol{\mu}}\right] = \left(\boldsymbol{\Psi}^{-1} + \mathbf{R}^T\boldsymbol{\Sigma}^{-1}\mathbf{R}\right)^{-1}$$

- Typically $\boldsymbol{\Sigma}$ is diagonal and $\boldsymbol{\Psi}$ depends on sampling of $\mathbf{z}_0$

  $\Rightarrow$ Estimate from historical data $\{\mathbf{z}_0\}$ or previous estimates $\hat{\boldsymbol{\mu}}$

- Likely to obtain negative $\hat{\mu}_{ij}$ if link counts are low. Constrain $\mu_{ij} \geq 0$

- M. Bell, "The estimation of OD matrices by constrained generalized least squares," *Transportation Research*, vol. 25B, pp. 13-22, 1991

Bayesian approach

- Instead of historical data, regularize with prior $\boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\mu}_0, \tau^2 \mathbf{I})$
- Suppose $\mathbf{X} = \mathbf{R}\boldsymbol{\mu} + \boldsymbol{\varepsilon}$, with $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. MAP estimator

$$\hat{\boldsymbol{\mu}} := \mathbb{E}\left[\boldsymbol{\mu} \,\middle|\, \mathbf{X} = \mathbf{x}\right] = \boldsymbol{\mu}_0 + \mathbf{R}^T(\mathbf{R}\mathbf{R}^T + \lambda \mathbf{I})^{-1}(\mathbf{x} - \mathbf{R}\boldsymbol{\mu}_0)$$

$\Rightarrow$ Correction of $\boldsymbol{\mu}_0$ driven by error in predicting $\mathbf{x}$ as $\mathbf{R}\boldsymbol{\mu}_0$

- Uncertainty in the estimate assessed via the covariance matrix

$$\text{var}\left[\boldsymbol{\mu} \,\middle|\, \mathbf{X} = \mathbf{x}\right] = \tau^2 \left[\mathbf{I} - \mathbf{R}^T(\mathbf{R}\mathbf{R}^T + \lambda \mathbf{I})^{-1}\mathbf{R}\right]$$

- Smoothing parameter $\lambda := \sigma^2/\tau^2$. Limiting cases:

$\Rightarrow$ As $\lambda \to 0$ enforce $\mathbf{x} = \mathbf{R}\hat{\boldsymbol{\mu}}$

$\Rightarrow$ As $\lambda \to \infty$ then $\hat{\boldsymbol{\mu}} \to \boldsymbol{\mu}_0$

Poisson models and MLE

- Gaussian model inappropriate even if few $\{\mu_{ij}\}$ are small
- Independent, Poisson OD flows modeled as

$$P(\mathbf{Z} = \mathbf{z}; \boldsymbol{\mu}) = \prod_{ij} P(Z_{ij} = z_{ij}; \mu_{ij}) = \prod_{ij} \frac{e^{-\mu_{ij}} \mu_{ij}^{z_{ij}}}{z_{ij}!}$$

- Consider error-free observations $\mathbf{X} = \mathbf{RZ}$

  $\Rightarrow$ Distribution of $\mathbf{X}$ induced by that of $\mathbf{Z}$ above

  $\Rightarrow$ Elements of $\mathbf{X}$ not independent in general

  $\Rightarrow$ Multiple $\mathbf{z}$ solve $\mathbf{x} = \mathbf{Rz}$, for observed $\mathbf{X} = \mathbf{x}$

- Still $\boldsymbol{\mu}$ identifiable if columns of $\mathbf{R}$ all distinct and nonzero [Vardi '96]

$$P(\mathbf{X}; \boldsymbol{\mu}) = P(\mathbf{X}; \tilde{\boldsymbol{\mu}}) \Rightarrow \boldsymbol{\mu} = a\tilde{\boldsymbol{\mu}}$$

Example: Toy network (encore)

- Subgraph induced by $V' = \{a, v, c\}$, OD pairs $\{av, vc, ac\}$

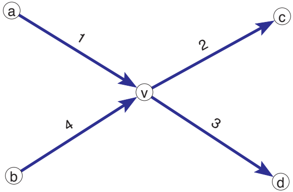

$$\mathbf{R} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

- Observe link counts $\mathbf{x} = [1, 2]^T$
- Two consistent flow sets

  $\mathbf{z}_1 = [0, 1, 1]^T$ and $\mathbf{z}_2 = [1, 2, 0]^T$

- Data likelihood $\mathcal{L}(\boldsymbol{\mu}; \mathbf{x}) = \mathrm{P}\left(\mathbf{X} = [1, 2]^T; \boldsymbol{\mu}\right)$ is

$$\begin{aligned}\mathcal{L}(\boldsymbol{\mu}; \mathbf{x}) &= \mathrm{P}\left(\mathbf{Z} = [0, 1, 1]^T; \boldsymbol{\mu}\right) + \mathrm{P}\left(\mathbf{Z} = [1, 2, 0]^T; \boldsymbol{\mu}\right) \\ &= (\mu_{ac}\mu_{vc} + \mu_{av}\mu_{vc}^2/2)\exp(-\mu_{ac} - \mu_{av} - \mu_{vc})\end{aligned}$$

Issues with Poisson MLE

- Q: What is the MLE $\hat{\boldsymbol{\mu}} = \arg\max_{\boldsymbol{\mu}\succeq\mathbf{0}} \mathcal{L}(\boldsymbol{\mu}; \mathbf{x})$?

  Solve $\max_{\boldsymbol{\mu}\succeq\mathbf{0}} (\mu_{ac}\mu_{vc} + \mu_{av}\mu_{vc}^2/2)\exp(-\mu_{ac} - \mu_{av} - \mu_{vc})$

  $\Rightarrow \nabla_{\boldsymbol{\mu}}\mathcal{L}(\boldsymbol{\mu}^*; \mathbf{x}) = \mathbf{0}$ for $\boldsymbol{\mu}^* = [1, 2, 0]^T$, but $\hat{\boldsymbol{\mu}} = [0, 1, 1]^T$

- Paradox? No, solution in the boundary of the feasible set

- For Poisson models $\mathcal{L}(\boldsymbol{\mu}; \mathbf{x})$ not concave in general [Vardi '96]

  $\Rightarrow$ Asymptotically concave for i.i.d. $\mathbf{x}_1, \ldots, \mathbf{x}_n$ if $\boldsymbol{\mu} \succ \mathbf{0}$

- EM-based MLE solver impractical ($\mathbb{E}\left[\mathbf{Z} \,\middle|\, \mathbf{X}, \boldsymbol{\mu}\right]$ tricky)

  $\Rightarrow$ Workaround: approximate $\mathbf{X} \sim \mathcal{N}(\mathbf{R}\boldsymbol{\mu}, \mathbf{R}\text{diag}(\boldsymbol{\mu})\mathbf{R}^T)$

  $\Rightarrow$ Resort to a method-of-moments estimator

Bayesian approach

- Goal: inference based on the posterior $\mathrm{P}\left(\mathbf{Z} \,\middle|\, \mathbf{X}\right)$

  $\Rightarrow$ Requires a prior $\mathrm{P}\left(\mathbf{Z}\right)$ and the model $\mathbf{X} = \mathbf{RZ}$

- Prior specification: $\mathbf{Z}$ independent, Poisson($\boldsymbol{\mu}$); along prior $\mathrm{P}\left(\boldsymbol{\mu}\right)$

$$\mathrm{P}\left(\mathbf{Z}, \boldsymbol{\mu}\right) = \mathrm{P}\left(\boldsymbol{\mu}\right) \prod_{ij} \mathrm{P}\left(Z_{ij} \,\middle|\, \mu_{ij}\right) = \mathrm{P}\left(\boldsymbol{\mu}\right) \prod_{ij} \frac{e^{-\mu_{ij}} \mu_{ij}^{z_{ij}}}{z_{ij}!}$$

- Observe link counts $\mathbf{X}$, conduct inference based on $\mathrm{P}\left(\mathbf{Z}, \boldsymbol{\mu} \,\middle|\, \mathbf{X}\right)$

  $\Rightarrow$ Simulate from the posterior via Gibbs sampler

  $\Rightarrow$ Iteratively resample from $\mathrm{P}\left(\mathbf{Z} \,\middle|\, \boldsymbol{\mu}, \mathbf{X}\right)$ and $\mathrm{P}\left(\boldsymbol{\mu} \,\middle|\, \mathbf{X}, \mathbf{Z}\right)$

- C. Tebaldi and M. West, "Bayesian inference on network traffic using link count data," *JASA*, vol. 93, pp. 557-573, 1998

Conditional posterior distributions

- $\mathrm{P}\left(\boldsymbol{\mu} \,\middle|\, \mathbf{X}, \mathbf{Z}\right)$: Independent μ_{ij} priors, i.e., $\mathrm{P}(\boldsymbol{\mu}) = \prod_{ij} \mathrm{P}(\mu_{ij})$, yields

$$\mathrm{P}\left(\boldsymbol{\mu} \,\middle|\, \mathbf{X}, \mathbf{Z}\right) \equiv \mathrm{P}\left(\boldsymbol{\mu} \,\middle|\, \mathbf{Z}\right) = \prod_{ij} \mathrm{P}\left(\mu_{ij} \,\middle|\, Z_{ij}\right) \propto \prod_{ij} \frac{e^{-\mu_{ij}} \mu_{ij}^{z_{ij}}}{z_{ij}!} \mathrm{P}(\mu_{ij})$$

  $\Rightarrow$ Given $\mathbf{Z}$, easy to simulate $\{\mu_{ij}\}$ from univariate posteriors

  $\Rightarrow$ Ex: If $\mathrm{P}(\mu_{ij})$ uniform or Gamma $\rightarrow \mathrm{P}\left(\mu_{ij} \,\middle|\, Z_{ij}\right)$ also Gamma

- $\mathrm{P}\left(\mathbf{Z} \,\middle|\, \boldsymbol{\mu}, \mathbf{X}\right)$: Model $\mathbf{X} = \mathbf{RZ}$ constrains $\mathbf{Z}$ given $\mathbf{X} = \mathbf{x}$

  $\Rightarrow$ Condition algebraically, rather than using Bayes' rule

- Illustrate through an example, then give general form of $\mathrm{P}\left(\mathbf{Z} \,\middle|\, \boldsymbol{\mu}, \mathbf{X}\right)$

Example: Toy network (second encore)

- Subgraph induced by $V' = \{a, v, c\}$, OD pairs $\{av, vc, ac\}$

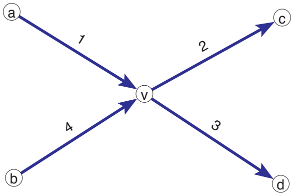

$$\mathbf{R} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

- Given $\mathbf{X} = \mathbf{x}$ and Z_{ac}

  $\Rightarrow$ Know Z_{av} and Z_{vc} since

  $Z_{av} = X_1 - Z_{ac}$ and $Z_{vc} = X_2 - Z_{ac}$

- Simulate from the full joint conditional posterior $\mathrm{P}\left(\mathbf{Z} \,\middle|\, \boldsymbol{\mu}, \mathbf{X}\right)$ by:
  (i) Drawing z_{ac} from the marginal posterior

$$\mathrm{P}\left(Z_{ac=}z_{ac} \,\middle|\, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right) \propto \frac{\mu_{ac}^{z_{ac}}}{z_{ac}!} \frac{\mu_{av}^{x_1 - z_{ac}}}{(x_1 - z_{ac})!} \frac{\mu_{vc}^{x_2 - z_{ac}}}{(x_2 - z_{ac})!}$$

  (ii) Evaluating $z_{av} = x_1 - z_{ac}$ and $z_{vc} = x_2 - z_{ac}$

General form of the OD flow posterior

- If $\text{rank}(\mathbf{R}) = N_e$, write $\mathbf{R} = [\mathbf{R}_1 \ \mathbf{R}_2]$ with $\mathbf{R}_1 \in \{0,1\}^{N_e \times N_e}$ invertible

  $\Rightarrow$ Can split flows $\mathbf{Z}^T = [\mathbf{Z}_1^T, \mathbf{Z}_2^T]^T$, where $\mathbf{Z}_1 = \mathbf{R}_1^{-1}(\mathbf{X} - \mathbf{R}_2\mathbf{Z}_2)$

- The sought conditional posterior has the form

$$\text{P}\left(\mathbf{Z} = \mathbf{z} \,\middle|\, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right) = \text{P}\left(\mathbf{Z}_1 = \mathbf{z}_1 \,\middle|\, \mathbf{Z}_2 = \mathbf{z}_2, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right) \text{P}\left(\mathbf{Z}_2 = \mathbf{z}_2 \,\middle|\, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right)$$

  $\Rightarrow \text{P}\left(\mathbf{Z}_1 = \mathbf{z}_1 \,\middle|\, \mathbf{Z}_2 = \mathbf{z}_2, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right) = \mathbb{I}\left\{\mathbf{z}_1 = \mathbf{R}_1^{-1}(\mathbf{x} - \mathbf{R}_2\mathbf{z}_2)\right\}$

  $\Rightarrow$ The "independent flows" $\mathbf{Z}_2$ have distribution

$$\text{P}\left(\mathbf{Z}_2 = \mathbf{z}_2 \,\middle|\, \boldsymbol{\mu}, \mathbf{X} = \mathbf{x}\right) \propto \prod_{ij} \frac{\mu_{ij}^{z_{ij}}}{z_{ij}!}$$

- Amenable to drawing entries of $\mathbf{Z}_2$ via a Gibbs sampler

Example: North Carolina road network

- Monroe, NC road network: $N_e = 20$ links and $N_f = 64$ flows

  $\Rightarrow$ Studied by transportation engineers at NC State University

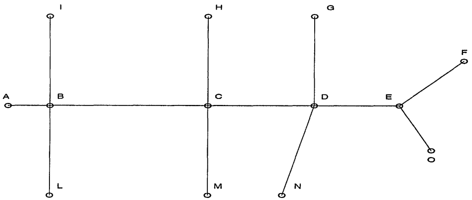

- Network fed to the traffic simulator *Integration* [Van Aerde et al '96]

  $\Rightarrow$ Modeled delays: congestion, traffic lights, turns, lanes merging

- Data (OD flows and link counts) for 2-hour morning period

Flow marginal posterior distributions

- Estimated marginal posteriors for 8 of the 64 OD flows (▲ = true)
  - ⇒ Uniform priors (top), and "informed" Gamma priors (bottom)

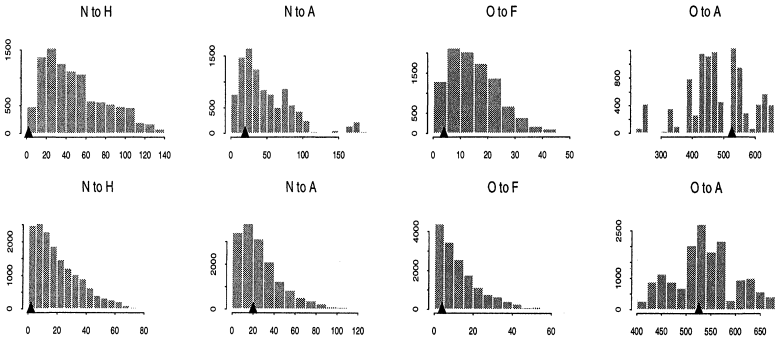

- Tend to overestimate smaller flows with a uniform prior
  - ⇒ Gamma priors based on recent data remove ambiguities

Relative entropy

- Consider a prior guess $\boldsymbol{\mu}^{(0)}$ of $\boldsymbol{\mu}$, normalized such that

$$\sum_{ij} \mu_{ij}^{(0)} = \sum_{ij} \mu_{ij} =: \mu_{++}$$

- Relative entropy "distance" between $\boldsymbol{\mu}$ and $\boldsymbol{\mu}^{(0)}$ given by

$$D(\boldsymbol{\mu} \| \boldsymbol{\mu}^{(0)}) = \sum_{ij} \frac{\mu_{ij}}{\mu_{++}} \log \left(\frac{\mu_{ij}}{\mu_{ij}^{(0)}} \right)$$

- Remarks
  - (i) Also known as Kullback-Liebler (KL) divergence
  - (ii) Dissimilarity between "distributions" $\{\mu_{ij}/\mu_{++}\}$ and $\{\mu_{ij}^{(0)}/\mu_{++}^{(0)}\}$
  - (iii) $D(\boldsymbol{\mu} \| \boldsymbol{\mu}^{(0)}) \geq 0$ always, and $D(\boldsymbol{\mu} \| \boldsymbol{\mu}^{(0)}) = 0 \Leftrightarrow \boldsymbol{\mu} = \boldsymbol{\mu}^{(0)}$

Entropy minimization

- Traffic matrix estimation: minimize $D(\boldsymbol{\mu} \| \boldsymbol{\mu}^{(0)})$ subject to $\mathbf{x} \approx \mathbf{R}\boldsymbol{\mu}$
- Dualize constraints via Lagrange multipliers $\boldsymbol{\lambda} \in \mathbb{R}^{N_e}$, solve

$$\min_{\boldsymbol{\mu}, \boldsymbol{\lambda}} D(\boldsymbol{\mu} \| \boldsymbol{\mu}^{(0)}) + \boldsymbol{\lambda}^T (\mathbf{x} - \mathbf{R}\boldsymbol{\mu})$$

- Given λ, optimality condition yields the estimator $(\mathbf{R} = [\mathbf{r}_{11}, \ldots, \mathbf{r}_{IJ}])$

$$\hat{\mu}_{ij}(\boldsymbol{\lambda}) = \mu_{ij}^{(0)} \exp\left(-1 - \boldsymbol{\lambda}^T \mathbf{r}_{ij}\right)$$

$\Rightarrow$ Multiplicative perturbation of $\boldsymbol{\mu}^{(0)}$, $\boldsymbol{\lambda}$ obtained numerically

$\Rightarrow$ Specify $\boldsymbol{\mu}^{(0)}$ from historical data $\mathbf{z}_0$, or prior estimates $\hat{\boldsymbol{\mu}}$

$\Rightarrow$ Non-negative solution guaranteed if $\boldsymbol{\mu}^{(0)} \succeq \mathbf{0}$

Entropy regularization

- Can view $D(\boldsymbol{\mu}\|\boldsymbol{\mu}^{(0)})$ as regularizer for $\mathbf{x} = \mathbf{Rz}$ → Penalized LS

$$\min_{\boldsymbol{\mu}\succeq\mathbf{0}} \|\mathbf{x} - \mathbf{Rx}\|^2 + \lambda D(\boldsymbol{\mu}\|\boldsymbol{\mu}^{(0)})$$

$\Rightarrow$ Convex problem, λ chosen via cross validation

- Couple interpretations:
  (i) Entropy minimization with relaxed constraint $\|\mathbf{x} - \mathbf{Rx}\|^2 \leq \tau$
  (ii) MAP for Gaussian model and prior $f(\boldsymbol{\mu})$ s. t. $\log f(\boldsymbol{\mu}) \propto D(\boldsymbol{\mu}\|\boldsymbol{\mu}^{(0)})$
  $\Rightarrow$ View as $f(\boldsymbol{\mu}) \approx$ multinomial, with probabilities $\propto \mu_{ij}^{(0)}$

- Ex: simple gravity model prior $\mu_{ij}^{(0)} \propto \mu_{i+}^{(0)}\mu_{+j}^{(0)}$ (more soon)

- Y. Zhang et al, "An information-theoretic approach to traffic matrix estimation," *SIGCOMM*, pp. 301-312, 2003

Dynamic methods

- Q: Traffic matrix estimation over time periods $t = 1, \ldots, \tau$?
- Given: link counts $\mathbf{x}_{1:\tau} := \{\mathbf{x}(t)\}_{t=1}^{\tau}$ and routing $\mathbf{R}_{1:\tau} := \{\mathbf{R}(t)\}_{t=1}^{\tau}$
- Determine: OD flows $\mathbf{z}_{1:\tau} := \{\mathbf{z}(t)\}_{t=1}^{\tau}$, where $\mathbf{x}(t) \approx \mathbf{R}(t)\mathbf{z}(t)$

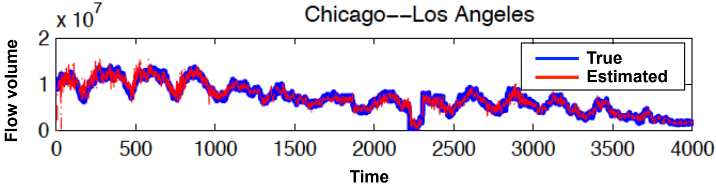

- Dynamic methods categorization: simultaneous or sequential
- A. Soule et al, "Traffic matrices: Balancing measurements, inference and modeling," *SIGMETRICS*, pp. 362-373, 2005

Simultaneous methods

- Simultaneous methods mostly based on the linear model

$$\mathbf{X}(t) = \mathbf{R}(t)\boldsymbol{\mu}(t) + \varepsilon(t), \quad t = 1, \ldots, \tau$$

- Penalized LS criteria employed to form $\hat{\boldsymbol{\mu}}_{1:\tau}$

$$\hat{\boldsymbol{\mu}}_{1:\tau} := \arg\min_{\boldsymbol{\mu}_{1:\tau}} \sum_{t=1}^{\tau} \|\mathbf{x}(t) - \mathbf{R}(t)\boldsymbol{\mu}(t)\|^2 + \lambda J(\boldsymbol{\mu}_{1:\tau})$$

  - Separable penalty $J(\boldsymbol{\mu}_{1:\tau}) = \sum_t J_t(\boldsymbol{\mu}(t))$ not uncommon
  - Ex: $J_t(\cdot)$ based on independent Gaussian or entropy-based priors

- Temporal correlations in $\mathbf{x}_{1:\tau}$ ignored $\rightarrow$ τ decoupled static problems

- Over short spans can assume $\boldsymbol{\mu}(t) = \boldsymbol{\mu}$, treat $\mathbf{x}_{1:\tau}$ as replicates

  $\Rightarrow$ LS ill-posed in general, but Poisson likelihood well behaved

Sequential methods

- Sequential methods leverage time correlations via Kalman filtering
- State $\boldsymbol{\mu}(t)$ and link count (measurement) $\mathbf{X}(t)$ equations

$$\begin{aligned}\boldsymbol{\mu}(t+1) &= \boldsymbol{\Phi}(t)\boldsymbol{\mu}(t) + \boldsymbol{\eta}(t)\\ \mathbf{X}(t) &= \mathbf{R}(t)\boldsymbol{\mu}(t) + \boldsymbol{\varepsilon}(t)\end{aligned}$$

$\Rightarrow$ $\boldsymbol{\eta}(t)$, $\boldsymbol{\varepsilon}(t)$ are zero-mean, white, with covariances $\boldsymbol{\Psi}(t)$, $\boldsymbol{\Sigma}(t)$

- Kalman filter (KF) in a nutshell
  - Prediction step: form prediction $\hat{\boldsymbol{\mu}}_{t+1:t}$ of $\boldsymbol{\mu}(t+1)$ using $\mathbf{x}_{1:t}$
  - Correction step: Update $\hat{\boldsymbol{\mu}}_{t+1:t+1}$ based on $\mathbf{x}(t+1) - \mathbf{R}(t+1)\hat{\boldsymbol{\mu}}_{t+1:t}$
- Also update recursively the error covariance matrix

$$\mathbf{M}_{t:t} := \mathbb{E}\left[(\hat{\boldsymbol{\mu}}_{t:t} - \boldsymbol{\mu}(t))(\hat{\boldsymbol{\mu}}_{t:t} - \boldsymbol{\mu}(t))^T\right]$$

Kalman filter updates

- Initialize $\hat{\boldsymbol{\mu}}_0$, $\mathbf{M}_{0:0}$ and run for $t = 0, \ldots, \tau$
- Prediction step:

$$\hat{\boldsymbol{\mu}}_{t+1:t} = \boldsymbol{\Phi}(t)\hat{\boldsymbol{\mu}}_{t:t}$$
$$\mathbf{M}_{t+1:t} = \boldsymbol{\Phi}(t)\mathbf{M}_{t:t}\boldsymbol{\Phi}^T(t) + \boldsymbol{\Psi}(t)$$

- Kalman gain update:

$$\mathbf{K}_{t+1} = \mathbf{M}_{t+1:t}\mathbf{R}^T(t+1)\left[\mathbf{R}(t+1)\mathbf{M}_{t+1:t}\mathbf{R}^T(t+1) + \boldsymbol{\Sigma}(t+1)\right]^{-1}$$

- Correction step:

$$\hat{\boldsymbol{\mu}}_{t+1:t+1} = \hat{\boldsymbol{\mu}}_{t+1:t} + \mathbf{K}_{t+1}\left[\mathbf{x}(t+1) - \mathbf{R}(t+1)\hat{\boldsymbol{\mu}}_{t+1:t}\right]$$
$$\mathbf{M}_{t+1:t+1} = [\mathbf{I} - \mathbf{K}_{t+1}\mathbf{B}(t+1)]\mathbf{M}_{t+1:t}[\mathbf{I} - \mathbf{K}_{t+1}\mathbf{B}(t+1)]^T + \mathbf{K}_{t+1}\boldsymbol{\Sigma}(t+1)\mathbf{K}_{t+1}^T$$

Practical considerations

- Model matrices $\mathbf{\Phi}(t)$, $\mathbf{\Psi}(t)$ and $\mathbf{\Sigma}(t)$ must be determined
  $\Rightarrow$ Often assumed time-invariant, and estimated from data
- Estimation depends on the model and data available
  $\Rightarrow$ Given $\mathbf{x}_{1:\tau}$, use variant of the EM algorithm
  $\Rightarrow$ Given flows $\mathbf{z}_{1:\tau}$, use AR(1) fitting techniques
- Z. Ghahramani and G. Hinton, "Parameter estimation for linear dynamical systems," *Tech. Rep. CRG-TR-96-2,* U. of Toronto, 1996
- KF should be periodically recalibrated $\rightarrow$ readjust $\mathbf{\Phi}$, $\mathbf{\Psi}$ and $\mathbf{\Sigma}$
  (a) Monitor the error process $\mathbf{x}(t) - \mathbf{R}(t)\hat{\boldsymbol{\mu}}_{t:t}$
  (b) Check if some entry e exceeds e.g., $3\sigma_e$ for few periods
  (c) Obtain σ_e^2 from diagonal of $\mathbf{R}(t)\mathbf{M}_{t:t}\mathbf{R}^T(t) + \mathbf{\Sigma}$

Case study

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Internet traffic monitoring

- Q: Why do ISPs monitor their networks routinely?
  - R1) Identify network (e.g., link) failures, their extent, and reasons
  - R2) Adjust routing $\rightarrow$ control congestion $\rightarrow$ optimize QoS
  - R3) Traffic engineering and management $\rightarrow$ capacity planning
  - R4) Security policies against cyber-attacks (e.g., worms, DoS)
- Availability of traffic matrices $\mathbf{Z}(t)$ key to traffic monitoring
- While possible, rarely measure Internet flows $Z_{ij}(t)$ at ISP level

  $\Rightarrow$ Concern on the volume of data generated

  $\Rightarrow$ Potential to adversely affect end-user QoS
- Limited $\mathbf{z}(t)$ to calibrate Internet traffic matrix estimation methods

Abilene traffic data

- Abilene backbone: $N_v = 11$ PoPs, $N_e = 30$ links, $N_f = 110$ flows

- Measure flows $\mathbf{z}_{1:\tau}$ for $\tau = 12 \times 24 \times 7 = 2,016$ time slots
  - $\Rightarrow$ Router sampling every 5 mins., week of Dec. 22, 2003
- Abilene routing matrix $\mathbf{R} \in \{0,1\}^{30 \times 110}$ given, time invariant
  - $\Rightarrow$ Pseudo-measurements: link counts $\mathbf{x}(t) = \mathbf{R}\mathbf{z}(t)$, $t = 1, \ldots, \tau$

Link counts and OD flow volumes

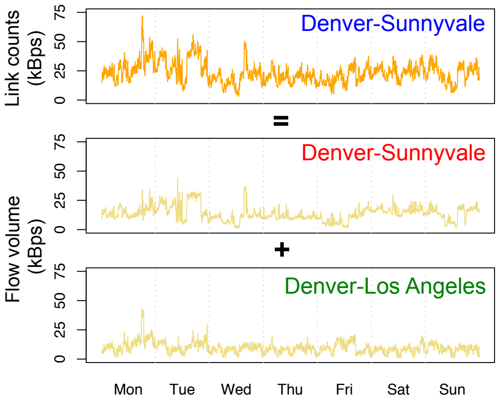

- Few flow patterns discernible in the aggregate (link count) data
  - $\Rightarrow$ OD flow recovery impossible in the absence of side information

Choice of traffic matrix estimation methods

- Compare static and dynamic methods for traffic matrix estimation
- Method 1: entropy-based approach termed tomogravity

$$\min_{\mathbf{z} \succeq \mathbf{0}} \|\mathbf{x} - \mathbf{R}\mathbf{z}\|^2 + \lambda \sum_{ij} \frac{z_{ij}}{z_{++}^{(0)}} \log\left(\frac{z_{ij}}{z_{ij}^{(0)}}\right), \quad \text{where } z_{ij}^{(0)} = z_{i+}^{(0)} z_{+j}^{(0)}$$

$\Rightarrow$ Simple gravity model prior adopted for $\mathbf{z}^{(0)}$, $\lambda = 0.01$

- Method 2: KF with state and measurement equations

$$\mathbf{Z}(t+1) = \mathbf{\Phi}\mathbf{Z}(t) + \boldsymbol{\eta}(t)$$
$$\mathbf{X}(t) = \mathbf{R}\mathbf{z}(t)$$

$\Rightarrow$ No error injected to the pseudo-measurements $\mathbf{x}(t)$

$\Rightarrow$ Matrices $\mathbf{\Phi}$ and $\mathbf{\Psi}$ estimated from $\mathbf{z}_{1:288}$ (Monday's flows)

Relative prediction error versus time

- Relative error averaged over OD pairs, as a function of time
  - $\Rightarrow$ Compare KF, tomogravity and bias-compensated tomogravity

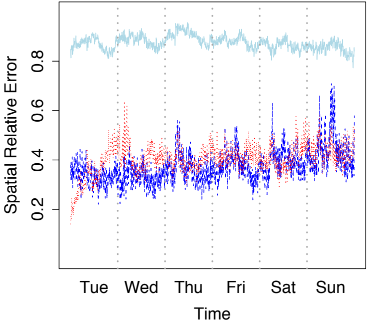

- Tomogravity overestimates, after bias-correction comparable to KF
  - $\Rightarrow$ KF performs better early in the week, then degrades

Relative prediction error versus flows

- Relative error averaged over time, for each OD pair in log-log scale
  - $\Rightarrow$ Symbol area $\propto$ mean volume of the flow
  - $\Rightarrow$ Color code: KF had higher error, tomogravity had higher error

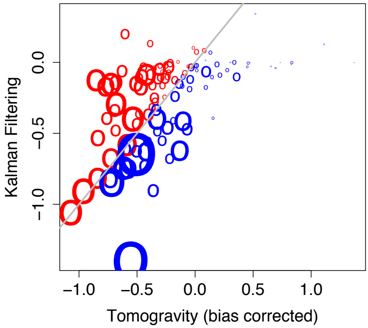

- KF mostly outperforms tomogravity for high- and low-volume flows

Flow volume predictions

- True flows superimposed with tomogravity and KF predictions

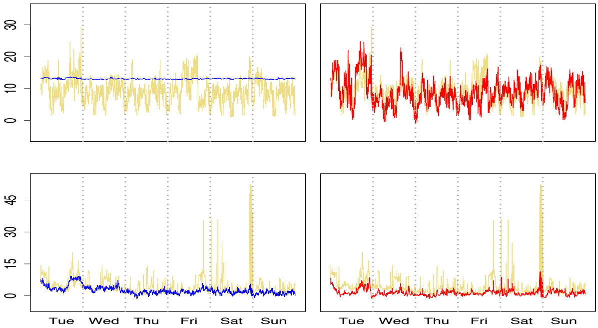

- Tomogravity completely misses the dynamics of the first flow

  $\Rightarrow$ But outperforms KF for the second flow

Network flow costs

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Network flow costs

- Consider a network graph $G(V, E)$. Let P be the set of paths in G
  - $\Rightarrow$ Path i-j has origin vertex $i \in \mathcal{I}$ and destination $j \in \mathcal{J}$
- Network flows costs at two levels of granularity: paths and links
  - $\Rightarrow$ Path costs $\mathbf{c} \in \mathbb{R}^{N_p}$ and link costs $\mathbf{x} \in \mathbb{R}^{N_e}$ related via

$$\mathbf{c} = \mathbf{R}^T \mathbf{x}$$

- Cost associated to path $=$ sum of the costs of the links traversed
- Ex: end-to-end delay is the sum of the delays in intermediate links
- **Our focus:** a particular class of problems involving inference of costs
  - $\Rightarrow$ Given data are limited (path) end-to-end measurements

Link costs from end-to-end measurements

Active network tomography

Given $\mathbf{c}^{obs}$ in paths $P^{obs} \subset P$, infer some characteristic of $\mathbf{x}$

- Actively inject traffic to measure $\mathbf{c}^{obs}$, e.g., multicast probing
  $\Rightarrow$ Traffic matrix estimation $\rightarrow$ observe link counts passively
- Tomography: unveil "internal" network characteristics
  $\Rightarrow$ Infer summands $\{x_e\}_{e \in P_{ij}}$ from aggregate c_{ij}
- Ex: determine link loss rates from packet loss measurements
- M. Coates et al, "Internet tomography," *IEEE Signal Processing Magazine*, vol. 19, pp. 47-65, 2002

Path costs from end-to-end measurements

Network kriging

Given $\mathbf{c}^{obs}$ in paths $P^{obs} \subset P$, predict $\mathbf{c}^{miss}$ in $P^{miss} = P \setminus P^{obs}$

- Kriging coined in geosciences for spatial interpolation or smoothing
- Key: exploit redundancies among links used by various paths

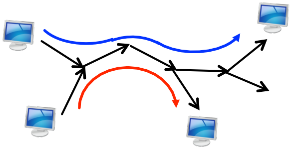

- D. Chua et al, "Network kriging," *IEEE J. Selected Areas in Communications*, vol. 24, pp. 2263-2276, 2006

Interpolation of path costs

- Number of paths N_p is much larger than N_e. Interpolation idea:
  (i) Select only N_e paths P^{obs} to monitor
  (ii) Use $\mathbf{c}^{obs} \in \mathbb{R}^{N_e}$ to determine link costs $\mathbf{x}$
  (iii) Since $\mathbf{R} = [\mathbf{R}_o \; \mathbf{R}_m]$, recover $\mathbf{c}^{miss} = \mathbf{R}_m^T \mathbf{x}$

- But in general $r := \text{rank}(\mathbf{R}) < N_e$, so $\mathbf{x}$ not identifiable

  $\Rightarrow$ Cannot find $x_{\mathcal{N}(\mathbf{R}^T)} \in \text{null}(\mathbf{R}^T)$ from $\mathbf{c} = \mathbf{R}^T\mathbf{x}$

  $\Rightarrow$ Only vectors $x_{\mathcal{R}(\mathbf{R}^T)} \in \text{range}(\mathbf{R}^T)$ can be identified in (ii)

- Of course do not need $\mathbf{x}$ to recover $\mathbf{c}^{miss} \Rightarrow x_{\mathcal{R}(\mathbf{R}^T)}$ suffices

- Y. Chen et al, "An algebraic approach to practical and scalable overlay network monitoring," *SIGCOMM*, vol. 34, pp. 55-66, 2004

Example: Unidentifiable link costs

- Graph $G(V, E)$ with $N_v = 4$ and $N_e = 3$, paths $\{AB, AC, BC\}$

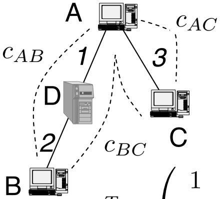

$$\mathbf{R}^T = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

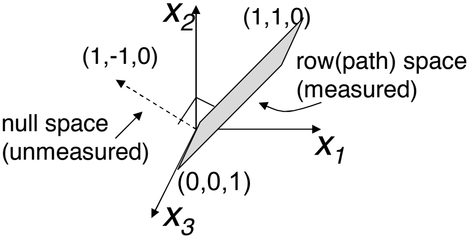

- Cannot identify x_1 and x_2 $\rightarrow$ Always show up summed in paths

Interpolation algorithm

- Key: monitor $r = \text{rank}(\mathbf{R})$ independent paths to recover $x_{\mathcal{R}(\mathbf{R}^T)}$
  $\Rightarrow$ Choose paths via QR decomposition of $\mathbf{R}$ with column pivoting

Interpolation algorithm:

(1) Select $r = \text{rank}(\mathbf{R}) < N_e$ independent paths to monitor
(2) Use $\mathbf{c}^{obs} \in \mathbb{R}^r$ to solve for $x_{\mathcal{R}(\mathbf{R}^T)}$ from $\mathbf{c}^{obs} = \mathbf{R}_o^T x_{\mathcal{R}(\mathbf{R}^T)}$

Least norm solution: $x_{\mathcal{R}(\mathbf{R}^T)} = \left(\mathbf{R}_o^T\right)^{\dagger}\mathbf{c}^{obs} = \mathbf{R}_o\left(\mathbf{R}_o^T\mathbf{R}_o\right)^{-1}\mathbf{c}^{obs}$

(3) Recover the unknown path costs as

$$\mathbf{c}^{miss} = \mathbf{R}_m^T x_{\mathcal{R}(\mathbf{R}^T)} = \mathbf{R}_m^T\mathbf{R}_o\left(\mathbf{R}_o^T\mathbf{R}_o\right)^{-1}\mathbf{c}^{obs}$$

- For $N_p = N_v^2$, conjecture $\text{rank}(\mathbf{R}) = O(N_v \log N_v)$ [Chen et al '04]
  $\Rightarrow$ Almost order of magnitude savings in measurement overhead

Effective rank of $\mathbf{R}$

- Interpolation appealing if we can monitor $r = \text{rank}(\mathbf{R})$ paths
  - $\Rightarrow$ Cannot recover $\mathbf{c}^{miss}$ if a single measurement is missing
- Network kriging: recast problem as one of statistical prediction
  - $\Rightarrow$ Accurate even with $s \ll \text{rank}(\mathbf{R})$ measurements. How?
- Since $r = \text{rank}(\mathbf{R})$, can write the SVD of $\mathbf{R}^T$ as

$$\mathbf{R}^T = \sum_{k=1}^{r} \sigma_k \mathbf{u}_k \mathbf{v}_k^T \approx \sum_{k=1}^{s \ll r} \sigma_k \mathbf{u}_k \mathbf{v}_k^T$$

- Observation: often most of the smaller σ_k are close to zero
  - $\Rightarrow$ We say $\mathbf{R}$ is effectively of lower rank than r
  - $\Rightarrow$ Intuition: dependencies among links used by various paths

Example: Reduced dimensionality in Abilene

- Singular values of the Abilene routing matrix $\mathbf{R}$
  - $\Rightarrow$ $N_e = 30$ links and $N_p = 110$ paths. Plot shows $\text{rank}(\mathbf{R}) = 30$

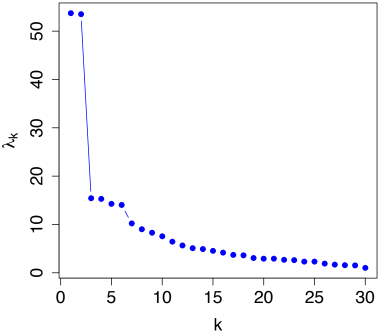

- Spectral gap apparent. Effective rank $s \in \{5, 10\}$, even $s = 2$?
  - $\Rightarrow$ Recover useful information about $\mathbf{c}$ from couple measurements

Routing matrix singular vectors

- Visualize top right singular vectors $\{\mathbf{v}_k\}_{k=1}^4$ of $\mathbf{R}^T$ (evecs. of $\mathbf{R}\mathbf{R}^T$)
  - $\Rightarrow$ Linearly independent "meta-paths" in "link space"
  - $\Rightarrow$ Intuition: shared patterns of links common to paths in $\mathbf{R}$

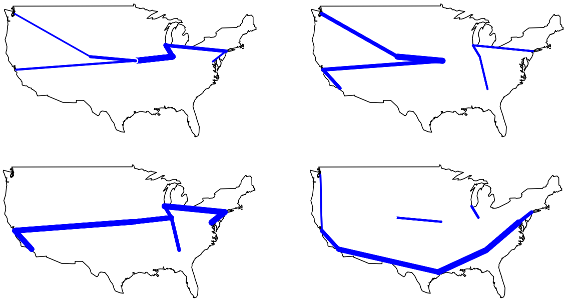

- Northern E-W meta-path $\{\mathbf{v}_k\}_{k=1}^3$, and southern E-W meta-path $\mathbf{v}_4$

Network kriging

- Consider predicting an arbitrary linear summary $\mathbf{a}^T\mathbf{c}$ of $\mathbf{c}$
- Ex: network-wide average path cost $\mathbf{a} = \mathbf{1}/N_p$, or c_{ij} where $\mathbf{a} = \mathbf{e}_{ij}$

- Let $\mathbf{x}$ be a realization of $\mathbf{X}$, with mean $\boldsymbol{\mu}$ and $\text{var}\left[\mathbf{X}\right] = \boldsymbol{\Sigma}$

  $\Rightarrow$ Because $\mathbf{C} = \mathbf{R}^T\mathbf{X}$, then $\mathbb{E}\left[\mathbf{C}\right] = \mathbf{R}^T\boldsymbol{\mu}$ and $\text{var}\left[\mathbf{X}\right] = \mathbf{R}^T\boldsymbol{\Sigma}\mathbf{R}$

- Given $s \leq \text{rank}(\mathbf{R})$ measured path costs $\mathbf{c}^{obs}$, find

$$\hat{p}(\mathbf{c}^{obs}) = \arg\min_{p} \mathbb{E}\left[(\mathbf{a}^T\mathbf{C} - p(\mathbf{C}^{obs}))^2\right]$$

  $\Rightarrow$ Minimum mean-squared error (MMSE) predictor, given by

$$\hat{p}(\mathbf{c}^{obs}) = \mathbb{E}\left[\mathbf{a}^T\mathbf{C} \,\middle|\, \mathbf{C}^{obs} = \mathbf{c}^{obs}\right] = \mathbf{a}_o^T\mathbf{c}^{obs} + \mathbb{E}\left[\mathbf{a}_m^T\mathbf{C}^{miss} \,\middle|\, \mathbf{C}^{obs} = \mathbf{c}^{obs}\right]$$

LMMSE predictor

- Restrict attention to linear (L)MMSE predictors $\hat{p}(\mathbf{c}^{obs}) = \hat{\mathbf{a}}^T\mathbf{c}^{obs}$

$$\hat{\mathbf{a}}^T\mathbf{c}^{obs} = \mathbf{a}_o^T\mathbf{c}^{obs} + \mathbf{a}_m^T\boldsymbol{\mu} + \mathbf{a}_m^T\mathbf{V}_{mo}\mathbf{V}_o^{-1}\left(\mathbf{c}^{obs} - \mathbf{R}_o^T\boldsymbol{\mu}\right)$$

$\Rightarrow$ Used (cross-)covariances $\mathbf{V}_o = \mathbf{R}_o^T\boldsymbol{\Sigma}\mathbf{R}_o$ and $\mathbf{V}_{mo} = \mathbf{R}_m^T\boldsymbol{\Sigma}\mathbf{R}_o$

- Estimate $\boldsymbol{\mu}$ from the data via generalized LS, i.e.,

$$\hat{\boldsymbol{\mu}} = \left(\mathbf{R}_o\mathbf{V}_o^{-1}\mathbf{R}_o^T\right)^{\dagger}\mathbf{R}_o\mathbf{V}_o^{-1}\mathbf{c}^{obs}$$

- Substitution of $\hat{\boldsymbol{\mu}}$ yields the network kriging predictor [Chua et al '06]

$$\hat{\mathbf{a}}^T\mathbf{c}^{obs} = \mathbf{a}_o^T\mathbf{c}^{obs} + \mathbf{a}_m^T\mathbf{V}_{mo}\mathbf{V}_o^{-1}\mathbf{c}^{obs}$$

- SVD-based path selection to minimize $\mathbb{E}\left[(\mathbf{a}^T\mathbf{C} - \hat{\mathbf{a}}^T\mathbf{C}^{obs})^2\right]$

$\Rightarrow$ Like the QR decomposition with pivoting in [Chen et al '04]

Example: Abilene path delays

- Abilene backbone: $N_v = 11$ PoPs, $N_e = 30$ links, $N_p = 110$ paths
- Measure link delays $\mathbf{x}_{1:\tau}$ for $\tau = 6 \times 24 \times 3 = 432$ time slots

  $\Rightarrow$ Router sampling every 10 mins., three days in 2003
- Abilene routing matrix $\mathbf{R} \in \{0, 1\}^{30\times 110}$ given, time invariant

  $\Rightarrow$ Pseudo-measurements: path costs $\mathbf{c}(t) = \mathbf{R}^T\mathbf{x}(t)$, $t = 1, \ldots, \tau$
- Applied the network kriging predictor to a subset $\mathbf{c}^{obs}(t)$

$$\hat{\mathbf{a}}^T\mathbf{c}^{obs}(t) = \mathbf{a}_o^T\mathbf{c}^{obs}(t) + \mathbf{a}_m^T\mathbf{V}_{mo}\mathbf{V}_o^{-1}\mathbf{c}^{obs}(t),\ t = 1, \ldots, \tau$$

  $\Rightarrow$ Various choices of $s \leq \text{rank}(\mathbf{R})$, SVD-based path selection

  $\Rightarrow$ Covariance $\boldsymbol{\Sigma}$ assumed diagonal, estimated from data

Path delay predictons

- Average path delay in Abilene predicted with $s = 3, 5, 7$, or 9 paths
  - $\Rightarrow$ Actual delay via interpolation of $s = 30 = \text{rank}(\mathbf{R})$ paths

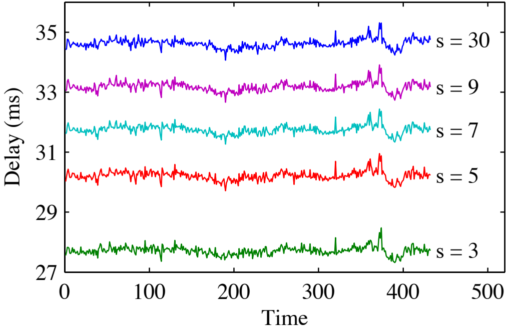

- Biased predictions, missing link information in approximated $\mathbf{R}$
  - $\Rightarrow$ Can be compensated if allowed to measure 30 paths once
- Predictions capture well the delay dynamics, for all s

Case study

Network flows, measurements and statistical analysis

Gravity models

Traffic matrix estimation

Case study: Internet traffic matrix estimation

Estimation of network flow costs

Case study: Dynamic delay cartography

Delay monitoring

- Motivating reasons
  - Assess network health
  - Fault diagnosis
  - Network planning
- Application domains
  - Old 8-second rule for WWW
  - Content-delivery networks
  - Peer-to-peer networks
  - Multiuser games
  - Dynamic server selection

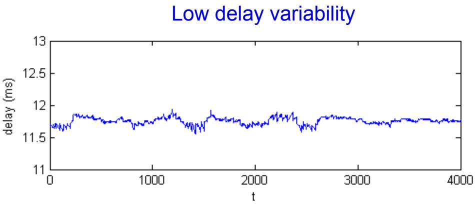

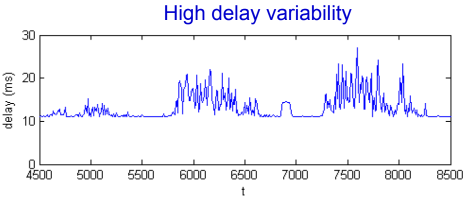

- Goal: infer path delays from limited end-to-end measurements

Predicting path delays

- Consider a network graph $G(V, E)$. Let P be the set of paths in G
- Several challenges in measuring all end-to-end path delays

  $\Rightarrow$ Overhead: number of paths $N_p = O(N_v^2)$

  $\Rightarrow$ Congested routers may drop packets

- Q: Can fewer measurements suffice?
- A: Yes! Most paths share multiple links $\Rightarrow$ Correlations [Chua'06]

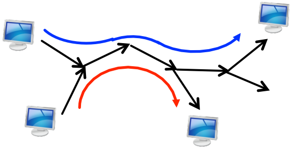

- End-to-end delay prediction problem: Given delay measurements $\mathbf{c}^{obs}$ in paths $P^{obs} \subset P$, predict $\mathbf{c}^{miss}$ in $P^{miss} = P \setminus P^{obs}$

Network kriging prediction

- Given (cross-)covariances $\mathbf{V}_o = \text{cov}[\mathbf{c}^{obs}]$ and $\mathbf{V}_{mo} = \text{cov}[\mathbf{c}^{miss}, \mathbf{c}^{obs}]$
- The universal kriging predictor is

$$\hat{\mathbf{c}}^{miss} = \mathbf{V}_{mo}\mathbf{V}_o^{-1}\mathbf{c}^{obs}$$

$\Rightarrow$ To obtain $\mathbf{V}_o$ and $\mathbf{V}_{mo}$, adopt a linear model for the path delays

$$\mathbf{c} = \mathbf{G}\mathbf{x} = \mathbf{R}^T\mathbf{x}, \qquad [\mathbf{G}]_{pl} = \begin{cases} 1, & \text{link } l \in \text{path } p \\ 0, & \text{otherwise} \end{cases}$$

- Link delays $\mathbf{x} \in \mathbb{R}^{N_e}$ and $\boldsymbol{\Sigma} = \text{cov}[\mathbf{x}] \Rightarrow$ From model $\text{cov}[\mathbf{c}]$ is

$$\begin{bmatrix} \mathbf{c}^{obs} \\ \mathbf{c}^{miss} \end{bmatrix} = \begin{bmatrix} \mathbf{S}_o \\ \mathbf{S}_m \end{bmatrix} \mathbf{G}\mathbf{x} \Rightarrow \begin{bmatrix} \mathbf{V}_o & \mathbf{V}_{om} \\ \mathbf{V}_{mo} & \mathbf{V}_m \end{bmatrix} = \begin{bmatrix} \mathbf{S}_o \\ \mathbf{S}_m \end{bmatrix} \mathbf{G}\boldsymbol{\Sigma}\mathbf{G}^\top \begin{bmatrix} \mathbf{S}_o \\ \mathbf{S}_m \end{bmatrix}^\top$$

$\Rightarrow$ Sampling matrix $\mathbf{S} = [\mathbf{S}_o^\top, \mathbf{S}_m^\top]^\top$ known, selected heuristically

Spatio-temporal prediction

- Network kriging prediction for a single temporal snapshot of delays
- D. Chua et al, "Network kriging," *IEEE J. Sel. Areas Communications*, vol. 24, pp. 2263-2272, 2006
- Wavelet-based approach for spatio-temporal delay prediction
  - Diffusion wavelet matrix constructed from the topology of G
  - Can capture temporal correlations, up to τ time slots
  - High complexity $O(\tau^3 P^3) \Rightarrow$ Challenging for $\tau > 10$
- M. Coates et al, "Compressed network monitoring for IP and all-optical networks," *Proc. ACM Internet Measurement Conference*, 2007
- Q: Should the same set of paths be measured every time slot?

  $\Rightarrow$ Low balancing? Effectiveness of random path selection?
- Low-complexity spatio-temporal inference with online path selection

Simple delay model

- Model delay $c_p(t)$ measured on path $p \in P$ at time t as

$$c_p(t) = \chi_p(t) + \nu_p(t) + \epsilon_p(t)$$

- Component $\chi_p(t)$ captures queuing delays, traffic dependent
  - Nonstationary: Random walk with driving noise covariance $\mathbf{C}_\eta$

$$\boldsymbol{\chi}(t) = \boldsymbol{\chi}(t-1) + \boldsymbol{\eta}(t)$$

- Component $\nu_p(t)$ lumps propagation, transmission, processing delays
  - Traffic independent, temporally white with covariance $\mathbf{C}_\nu = \alpha \mathbf{G}\mathbf{G}^\top$
- Measurement noise $\epsilon_p(t)$ i.i.d. over paths and time, $\text{var}\left[\epsilon_p(t)\right] = \sigma^2$

Kriged Kalman filter formulation

- Paths measured on subset $P^{obs} \subset P$, use sampling matrix $\mathbf{S}_o(t)$

$$\mathbf{c}^{obs}(t) = \mathbf{S}_o(t)\boldsymbol{\chi}(t) + \boldsymbol{\nu}^{obs}(t) + \boldsymbol{\epsilon}(t), \qquad \boldsymbol{\nu}^{obs}(t) := \mathbf{S}_o(t)\boldsymbol{\nu}(t)$$

- Kriged Kalman filter (KKF) state and measurement equations

$$\begin{aligned}\boldsymbol{\chi}(t) &= \boldsymbol{\chi}(t) + \boldsymbol{\eta}(t)\\ \mathbf{c}^{obs}(t) &= \mathbf{S}_o(t)\boldsymbol{\chi}(t) + \boldsymbol{\nu}^{obs}(t) + \boldsymbol{\epsilon}(t)\end{aligned}$$

- Goal: given historical data $\mathcal{H}(t) = \{\mathbf{c}^{obs}(\tau)\}_{\tau=1}^{t}$, predict $\mathbf{c}^{miss}(t)$
- K. Rajawat et al, "Dynamic network delay cartography," *IEEE Trans. Info. Theory,* vol. 60, pp. 2910-2920, 2014

Kriged Kalman filter updates

- State and covariance update recursions

$$\hat{\boldsymbol{\chi}}(t) := \mathbb{E}\left[\boldsymbol{\chi}(t) \,\middle|\, \mathcal{H}(t)\right] = \hat{\boldsymbol{\chi}}(t-1) + \mathbf{K}(t)[\mathbf{c}^{obs}(t) - \mathbf{S}_o(t)\hat{\boldsymbol{\chi}}(t-1)]$$

$$\mathbf{M}(t) := \mathbb{E}\left[(\hat{\boldsymbol{\chi}}(t) - \boldsymbol{\chi}(t))(\hat{\boldsymbol{\chi}}(t) - \boldsymbol{\chi}(t))^\top\right] = [\mathbf{I} - \mathbf{K}(t)\mathbf{S}_o(t)][\mathbf{M}(t-1) + \mathbf{C}_\eta]$$

- KKF gain

$$\mathbf{K}(t) = [\mathbf{M}(t-1)+\mathbf{C}_\eta]\mathbf{S}_o^\top(t)[\mathbf{S}_o(t)(\mathbf{M}(t-1)+\mathbf{C}_\eta+\mathbf{C}_\nu)\mathbf{S}_o^\top(t)+\sigma^2\mathbf{I}]^{-1}$$

- Kriging predictor $\hat{\mathbf{c}}^{miss}(t) = \mathbf{S}_m(t)\hat{\boldsymbol{\chi}}(t) + \hat{\boldsymbol{\nu}}^{miss}(t)$, where

$$\hat{\boldsymbol{\nu}}^{miss}(t) := \mathbf{S}_m(t)\mathbf{C}_\nu\mathbf{S}_o^\top(t)[\mathbf{S}_o(t)\mathbf{C}_\nu\mathbf{S}_o^\top(t)+\sigma^2\mathbf{I}]^{-1}(\mathbf{c}^{obs}(t)-\mathbf{S}_o(t)\hat{\boldsymbol{\chi}}(t))$$

Kriging covariance models

- Q: How do we find the spatial covariance $\mathbf{C}_\nu$?
- Idea: paths sharing multiple links should be highly correlated

  $\Rightarrow$ Linear model: $\mathbf{C}_\nu = \alpha \mathbf{G}\mathbf{G}^\top$

  $\Rightarrow$ Graph Laplacian model: $\mathbf{C}_\nu = \mathbf{L}^\dagger$

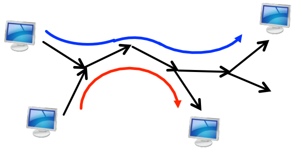

- Similar principles used to define graph kernels
- Can also handle route changes, especially incremental changes

Selection of measured paths

- KKF can model and track network wide delays given sample paths
- Q: Practical sampling of paths? Optimal measurements? Criterion?
- Error covariance matrix (define $\mathbf{\Phi}(t) = [\mathbf{M}(t-1) + \mathbf{C}_\nu + \mathbf{C}_\eta]/\sigma^2$)

$$\mathbf{M}^{miss}(t) = \mathbb{E}\left[(\mathbf{c}^{miss}(t) - \hat{\mathbf{c}}^{miss}(t))(\mathbf{c}^{miss}(t) - \hat{\mathbf{c}}^{miss}(t))^\top\right]$$
$$= \sigma^2\mathbf{I} + \sigma^2\mathbf{S}_m(t)\left[\mathbf{\Phi}^{-1}(t) + \mathbf{S}_o^\top(t)\mathbf{S}_o(t)\right]^{-1}\mathbf{S}_m^\top(t)$$

- Optimal experimental design

$$\hat{P}^{obs}(t) := \arg\min_{P^{obs}\subset P} \log\det(\mathbf{M}^{miss}(t)), \quad \text{s. to } |P^{obs}| = N_p^{obs}$$

- Criterion: D-optimal design, i.e., entropy of a Gaussian RV

  $\Rightarrow$ Cost depends on P^{obs} via sampling matrix $\mathbf{S}_o(t)$ in $\mathbf{M}^{miss}(t)$

Greedy algorithm

- Simple greedy algorithm to select observed paths P^{obs}
- Repeat $|P^{obs}|$ times: $P^{obs} \leftarrow P^{obs} \cup \arg\max_{p \notin P^{obs}} \delta_{P^{obs}}(p)$, where

$$\delta_{\emptyset}(p) = -\log\left(1 + [\mathbf{M}(t-1) + \mathbf{C}_\eta + \mathbf{C}_\nu]_{p,p}\right)$$
$$\delta_{P^{obs}}(p) = -\log\left(1 + \left[((\mathbf{M}(t-1) + \mathbf{C}_\eta + \mathbf{C}_\nu)^{-1} + \mathbf{S}^\top \mathbf{S})^{-1}\right]_{p,p}\right)$$

$\Rightarrow$ Submodular, monotonic $\rightarrow$ Greedy solution $(1 - e^{-1})$ optimal

- Increments $\delta_{P^{obs}}(p)$ efficiently evaluated in $O(|P||P^{obs}|^3)$

$\Rightarrow$ Operational complexity can be reduced further [Krause'11]

- Can be modified to handle cases when
  (i) Few nodes measure delays on all paths. Which nodes to choose?
  (ii) All nodes measure delay on only one path. Which paths to chsose?

Empirical validation: Internet2

- Internet2 backbone: 72 paths, lightly loaded network

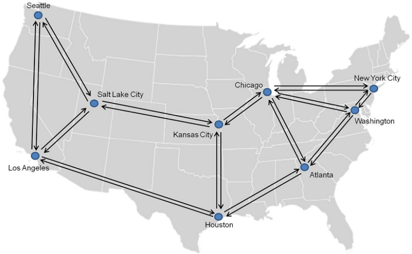

- One-way delay measurements collected using `OWAMP`

  $\Rightarrow$ Every minute for 3 days in July 2011 $\sim$ 4500 samples
- Training phase employed to estimate $\mathbf{C}_\eta$, α [Myers'76]
  - Modified estimators to handle measurements on subsets of paths
  - First 1000 samples on 50 random paths used for training

Network delay cartography: Internet2

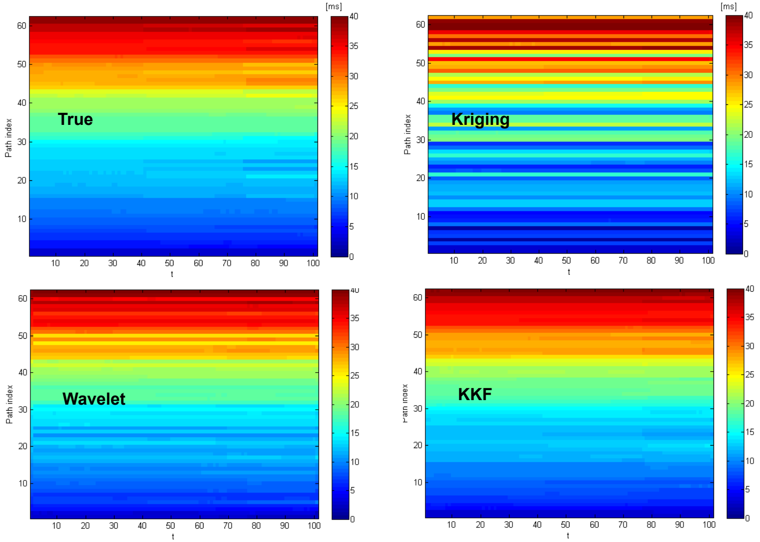

Prediction error: Internet2

- Normalized mean-square prediction error as figure of merit

$$\text{NMSPE} = \frac{1}{T|P^{miss}|}\sum_{t=1}^{T}\left\|\hat{\mathbf{c}}^{miss}(t) - \mathbf{c}^{miss}(t)\right\|^2$$

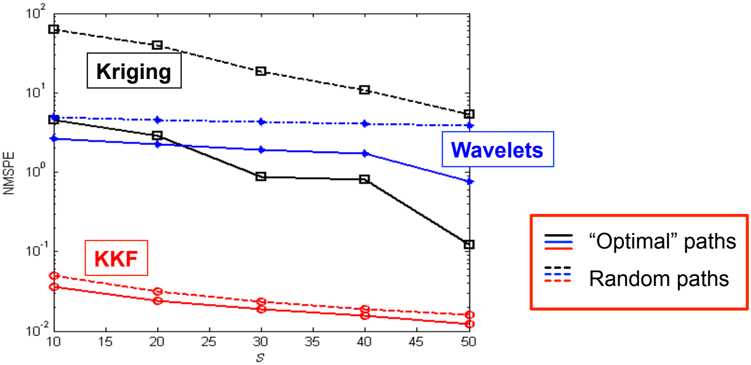

Empirical validation: NZ-AMP

- NZ-AMP delay dataset: 186 paths, heavily loaded network

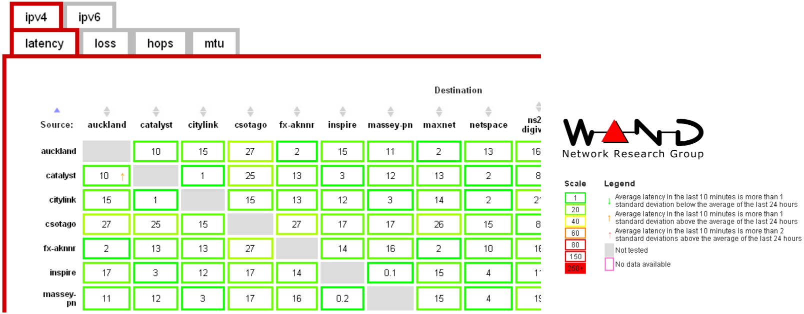

Source:	auckland	catalyst	citylink	csotago	fx-aknnr	inspire	massey-pn	maxnet	netspace	ns2 digiv
auckland		10	15	27	2	15	11	2	13	16
catalyst	10 ↑		1	25	13	3	12	13	2	8
citylink	15	1		15	13	12	3	14	2	2
csotago	27	25	15		27	17	17	26	15	8
fx-aknnr	2	13	13	27		14	16	2	10	1
inspire	17	3	12	17	14		0.1	15	4	1
massey-pn	11	12	3	17	16	0.2		15	4	1

- Round-trip-times measured using `ICMP`, paths via `scamper`
  - $\Rightarrow$ Every 10 minutes in August 2011 $\sim$ 4500 samples

Prediction error: NZ-AMP

Random path selection

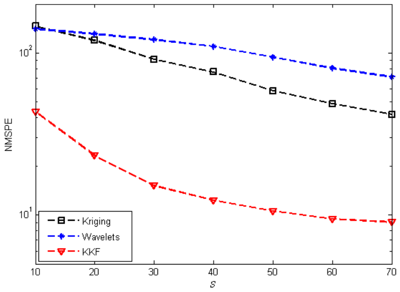

"Optimal" path selection

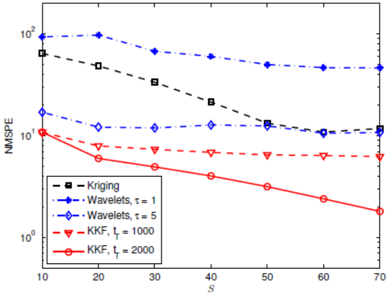

- NMSPE order of magnitude larger than for the Internet2 data

  $\Rightarrow$ Attributed to the markedly higher delay variability here

Delay scatter plots: NZ-AMP

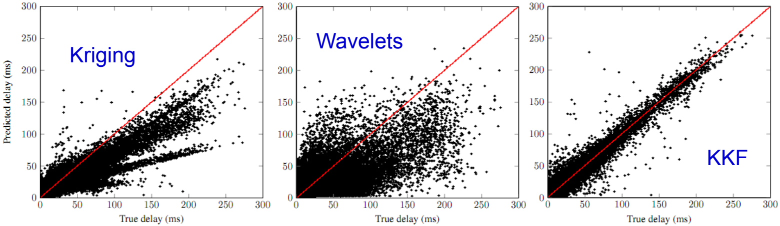

- Prediction of path delays. Plot $\hat{c}_{ij}^{miss}$ vs c_{ij}^{miss}
  - $\Rightarrow$ Fairly linear trend for KKF, variability $\nearrow$ for short delays
  - $\Rightarrow$ Network kriging and diffusion wavelets biased down

Glossary

- Network traffic flows
- Routing matrix
- Traffic matrix
- Link counts
- Network flow costs
- Network monitoring
- Gravity model
- Generalized linear model
- Traffic matrix estimation
- Network tomography
- Poisson traffic models
- Entropy minimization
- Tomogravity
- Kalman filter
- End-to-end measurements
- Active network tomography
- Network kriging
- Path-cost interpolation
- Identifiability
- Effective rank
- (L)MMSE predictor
- Path selection
- Diffusion wavelets
- Kriged Kalman filter
- Optimal experimental design
- Submodular function