Estimating Point-to-Multipoint Demand Matrices

Matthew Roughan, Yin Zhang

AT&T Shannon Lab

Problem: point-to-point TM estimation

Have link traffic measurements
Want to know demands from source to destination

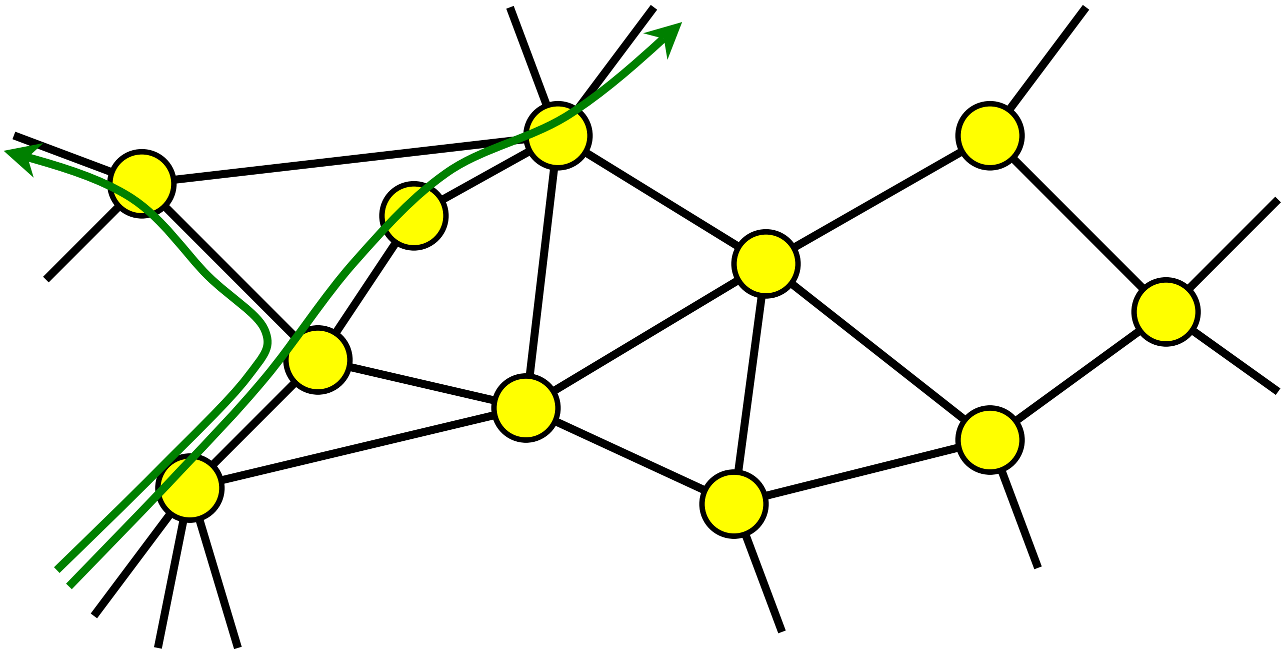

Example App: reliability analysis

Under a link failure, routes change

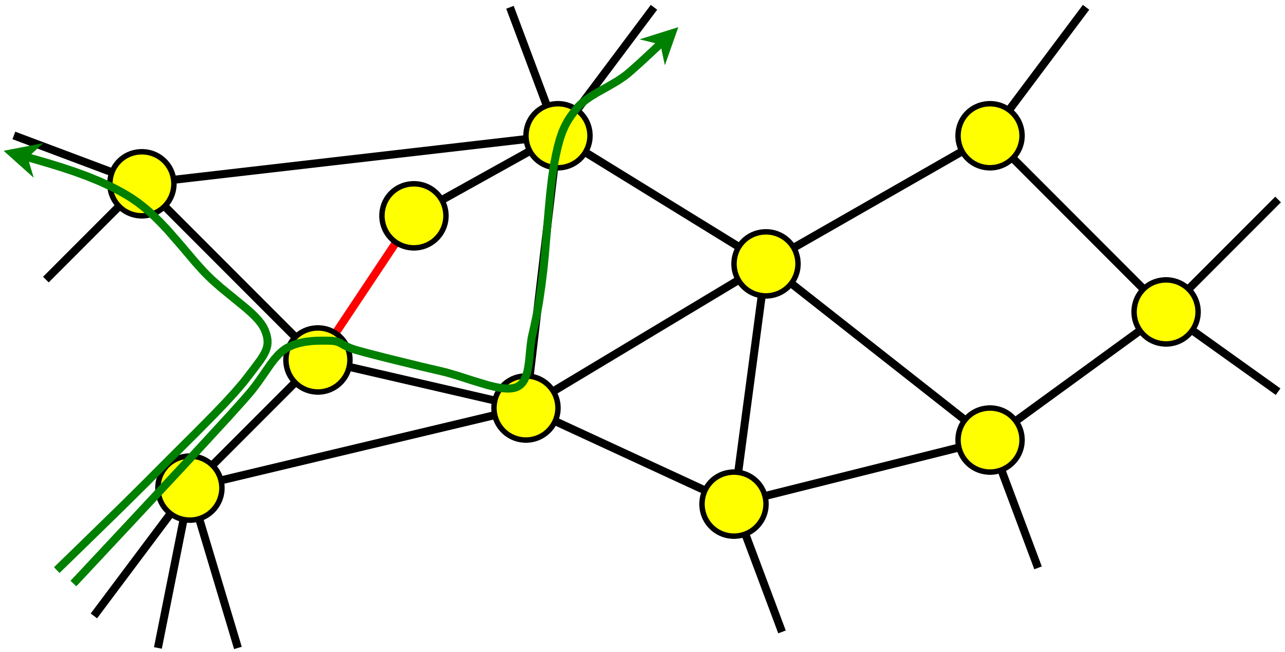

The problem

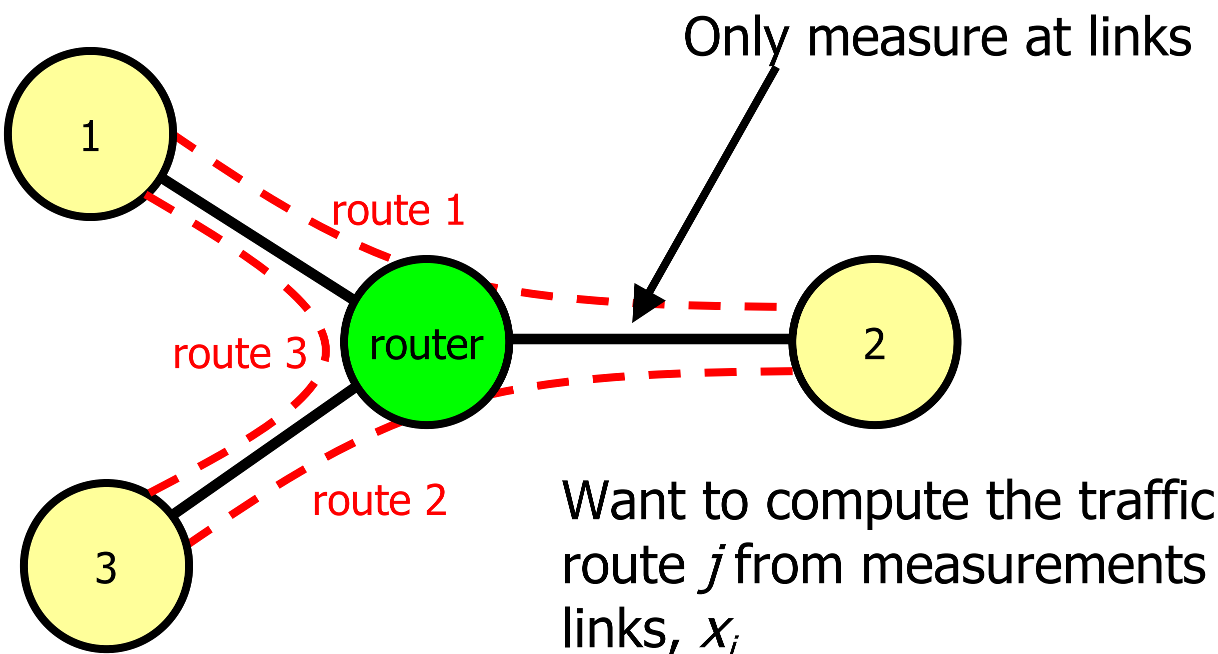

Want to compute the traffic y_j along route j from measurements on the links, x_i

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix}$$

The problem

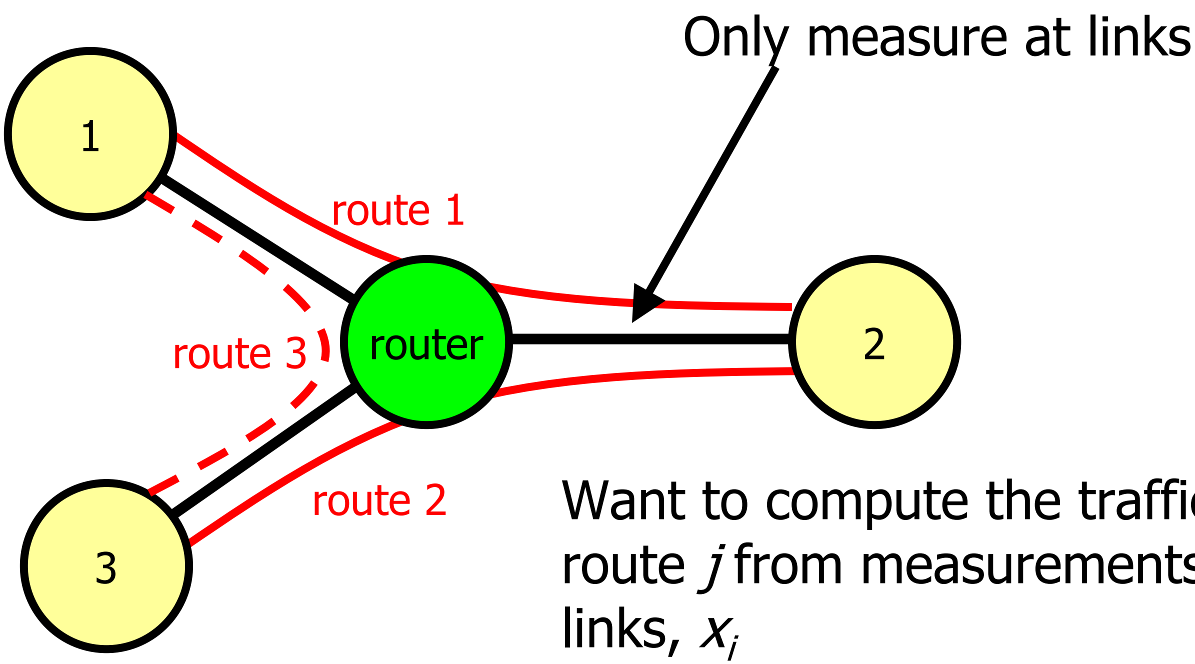

Want to compute the traffic t_j along route j from measurements on the links, x_i

$$\mathbf{y} = \mathbf{A}\ \mathbf{x}$$

Point-to-Multipoint

- We are trying to find an *invariant*
  - Something that doesn't change when the network changes
- But we only see one part of the network

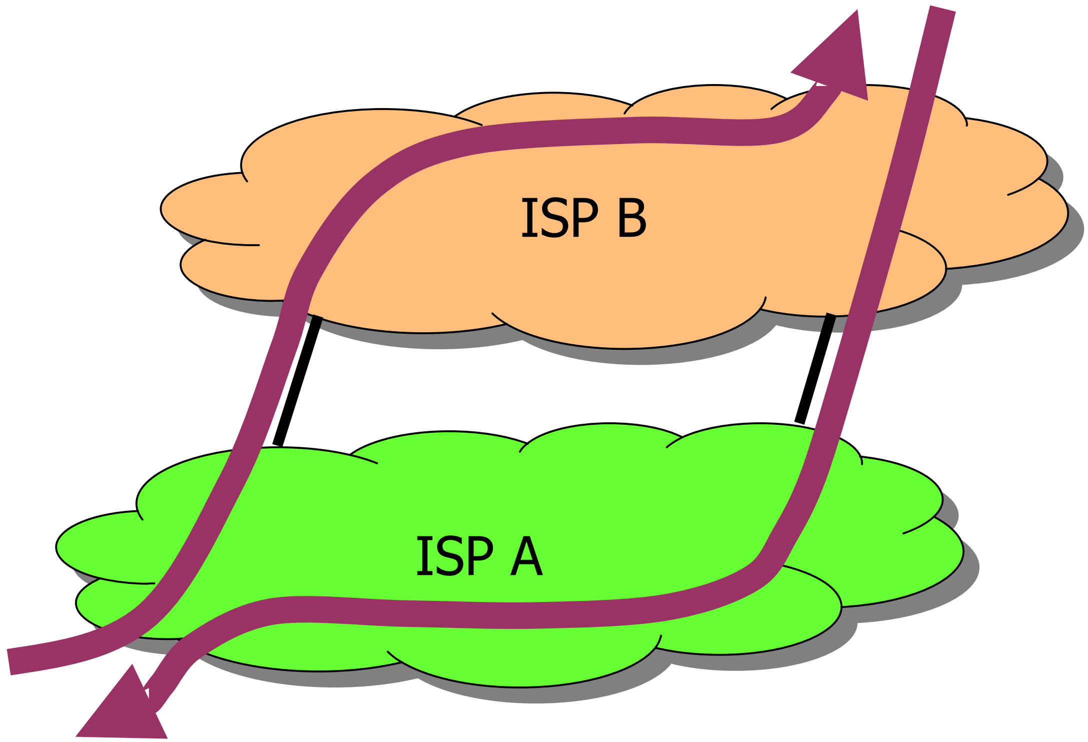

Peering link failure

- peering link failure so the traffic uses alternate
  - Traffic matrix changes

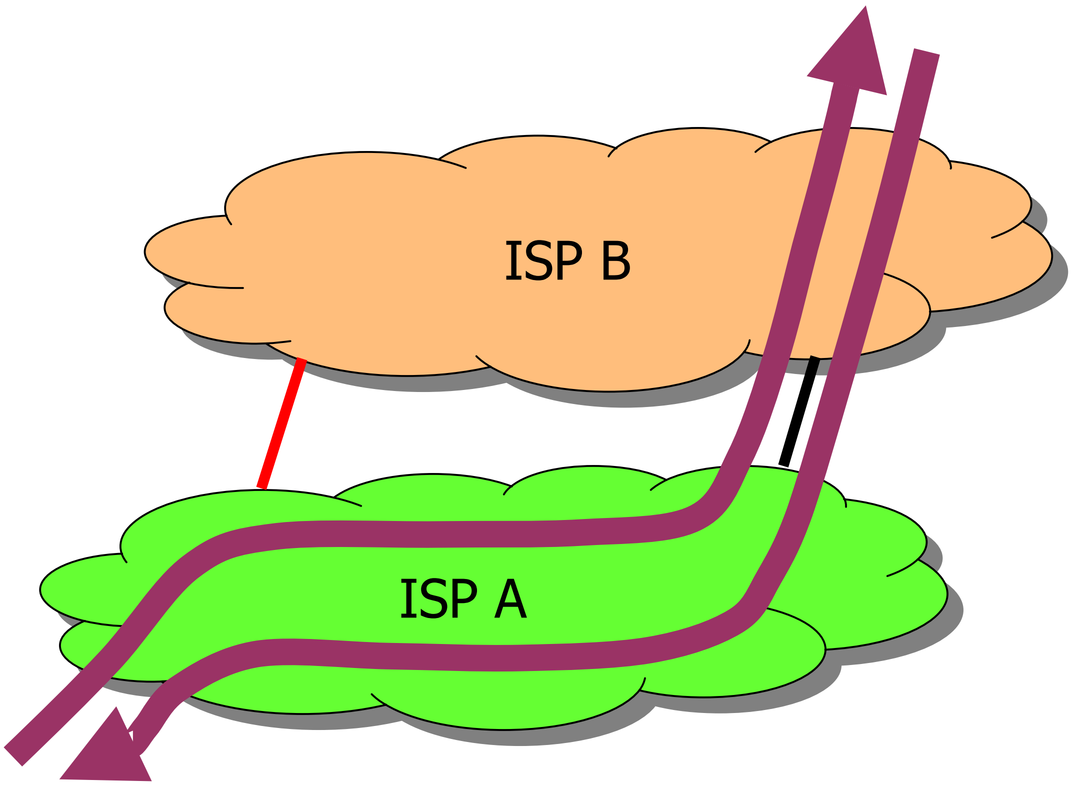

Invariant

- Traffic matrix changes
  - Not invariant to all network changes
- Point-to-Multipoint Demand Matrix
  - Remains unchanged
- Can be measured using netflow
- We can estimate it from link data
  - Much of multipoint traffic goes out a disjoint sets of exit points, grouped by peer
  - Larger peers all have private peering links
  - BGP policy between sets of peering links typically the same
- Basic trick
  - $y = A x$
  - But now x is the point to multipoint traffic matrix

Multipoint flows

- Inbound flows (peering -> access) still P-P
- Internal flows (access -> access) still P-P
- Outbound flows (access -> peering) become P-MP

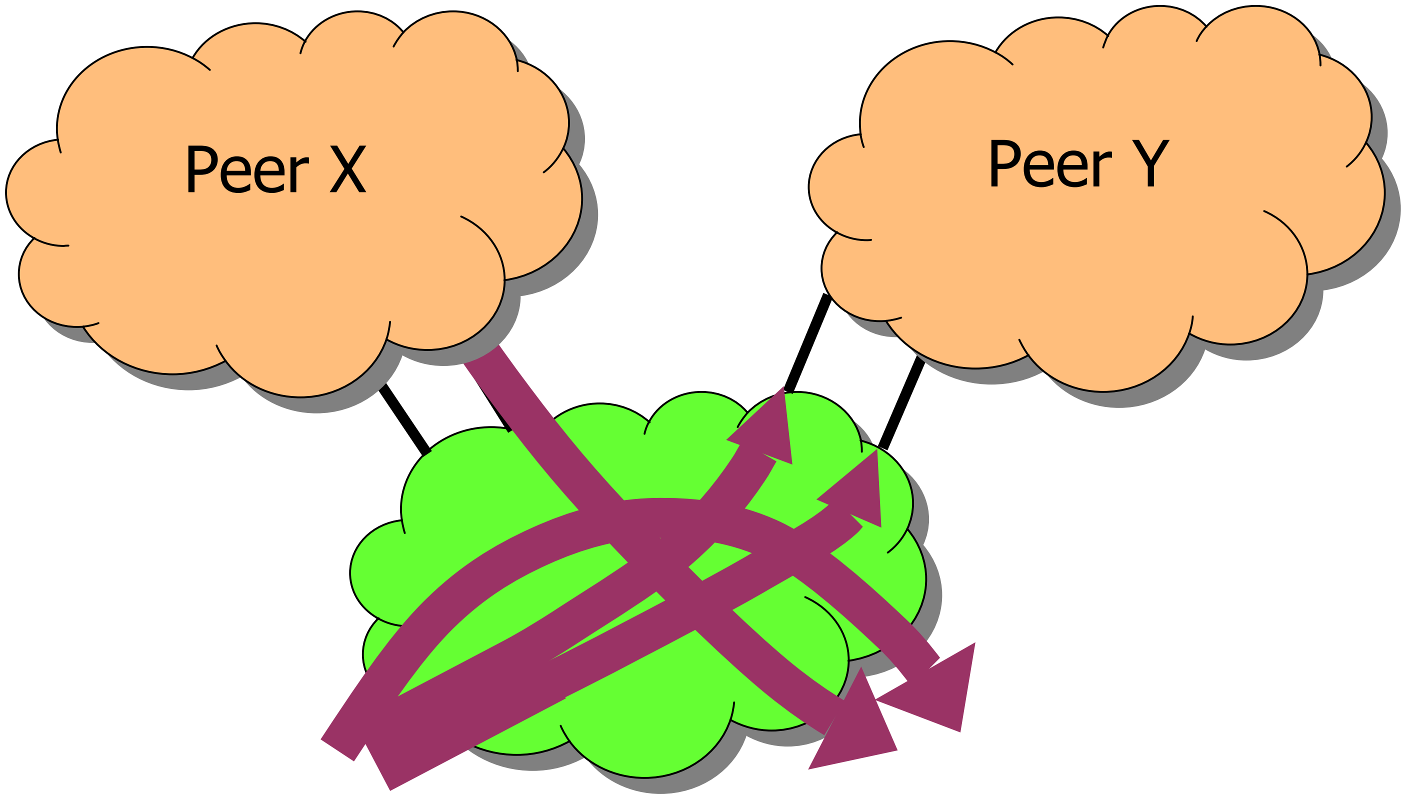

Gravity Model

- Assume traffic between sites is proportional to traffic at each site
  - Assumes there is no systematic difference between traffic in LA and NY
  - Only the total volume matters
  - Could include a distance term, but locality of information is not as important in the Internet as in other networks
- Equivalent to source/destination independence

$$\text{Prob}(S=s, D=d) = \text{prob}(S=s)\ \text{prob}(D=d)$$

$$\text{Prob}(D=d \mid S=s) = P(D=d)$$

Generalized gravity model

- Internet routing is asymmetric
  - Control over exit points
- Negligable transit traffic
  - No traffic from peer X to peer Y transits the backbone
- Leads to *conditionally independent* model
  - Independent conditional on the class of the ingress/egress points
  - Classes
    - Peering
    - Access

$$\text{Prob}(S=s, D=d \mid s \in C_s, d \in C_d) = \text{prob}(S=s \mid s \in C_s)\ \text{prob}(D=d \mid d \in C_d)$$

Combining gravity model and tomography

- In general the aren't enough constraints
- Constraints give a subspace of possible solutions

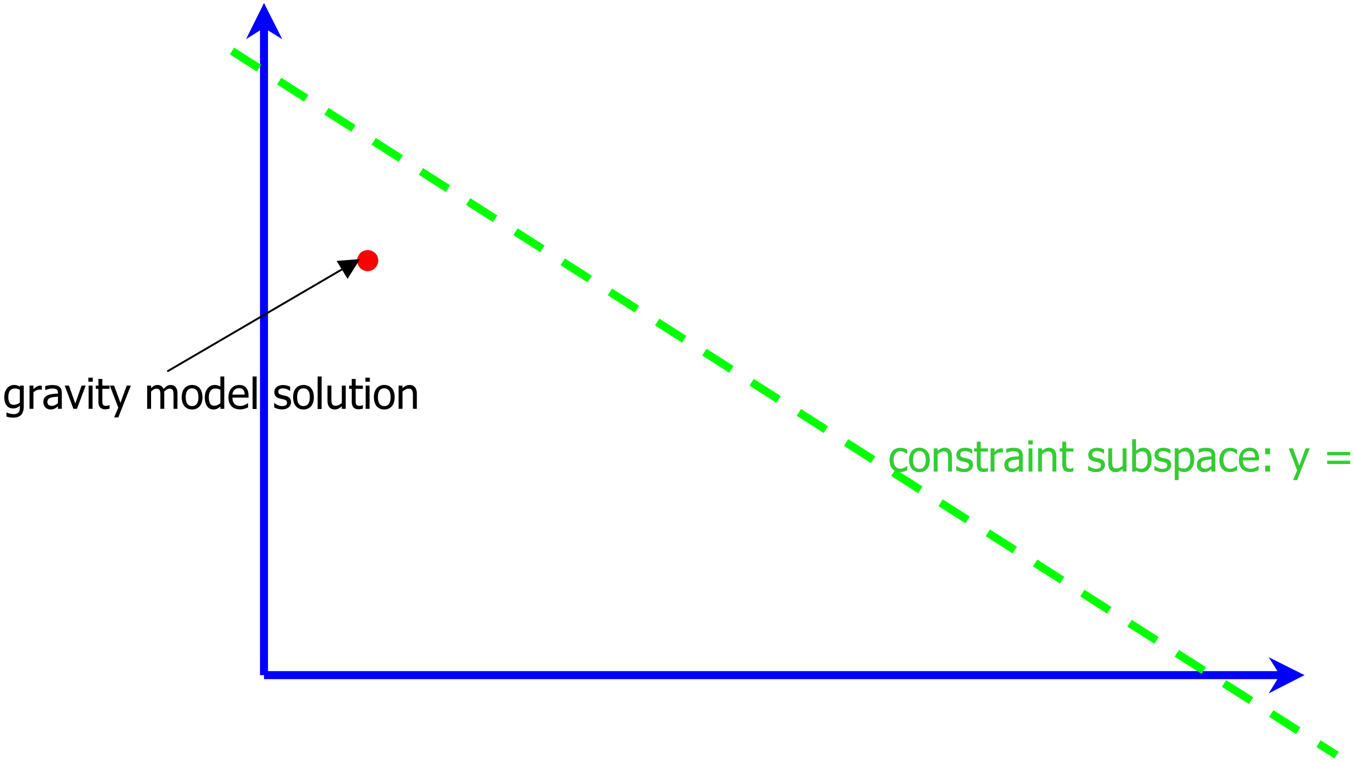

Solution

- Find a solution which
  - Satisfies the constraint
  - Is close to the gravity model *(Kullback-Liebler distance)*

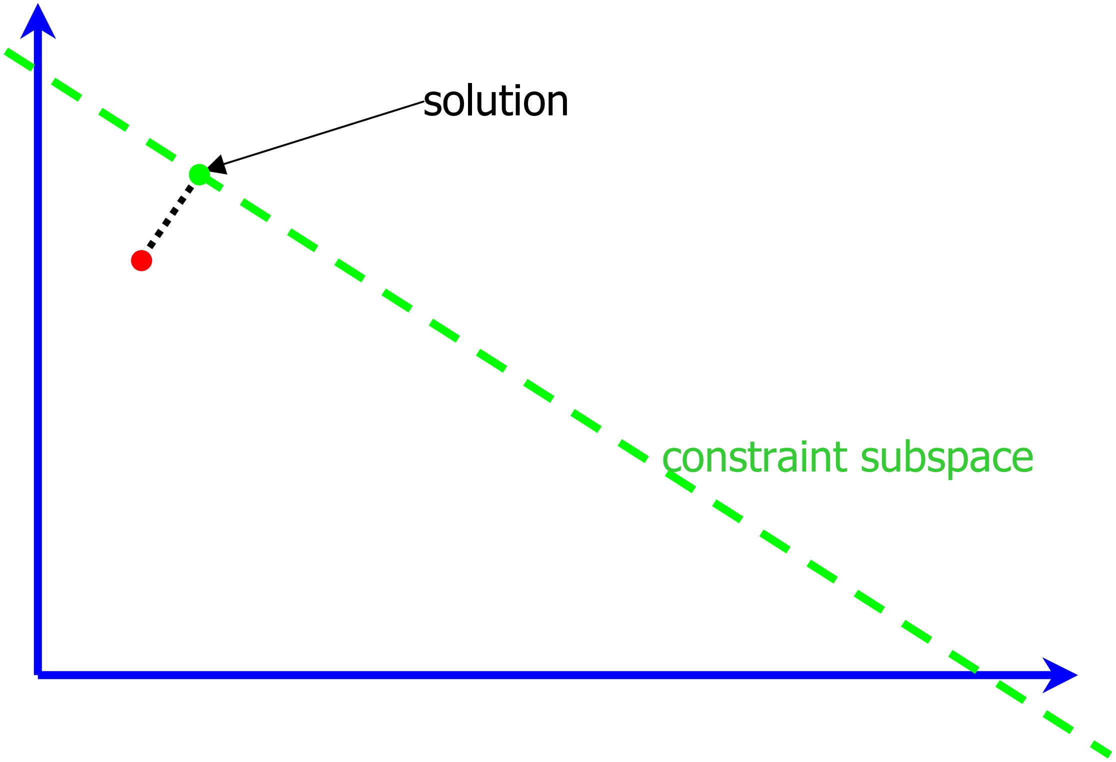

Validation for point-to-point

- Results good: ±20% bounds for larger flows
- Observables even better
- Robust
- Fast

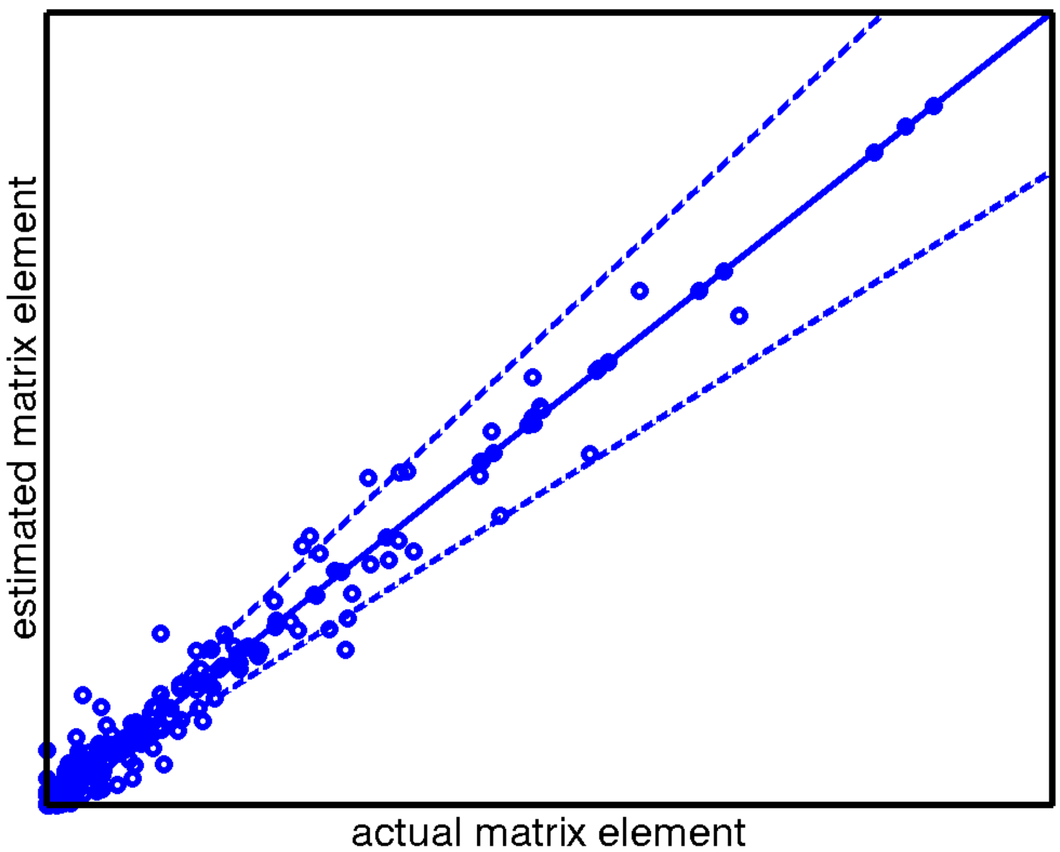

Validation for point-to-multipoint

- Not quite as good
- average error 20% compared to 12%

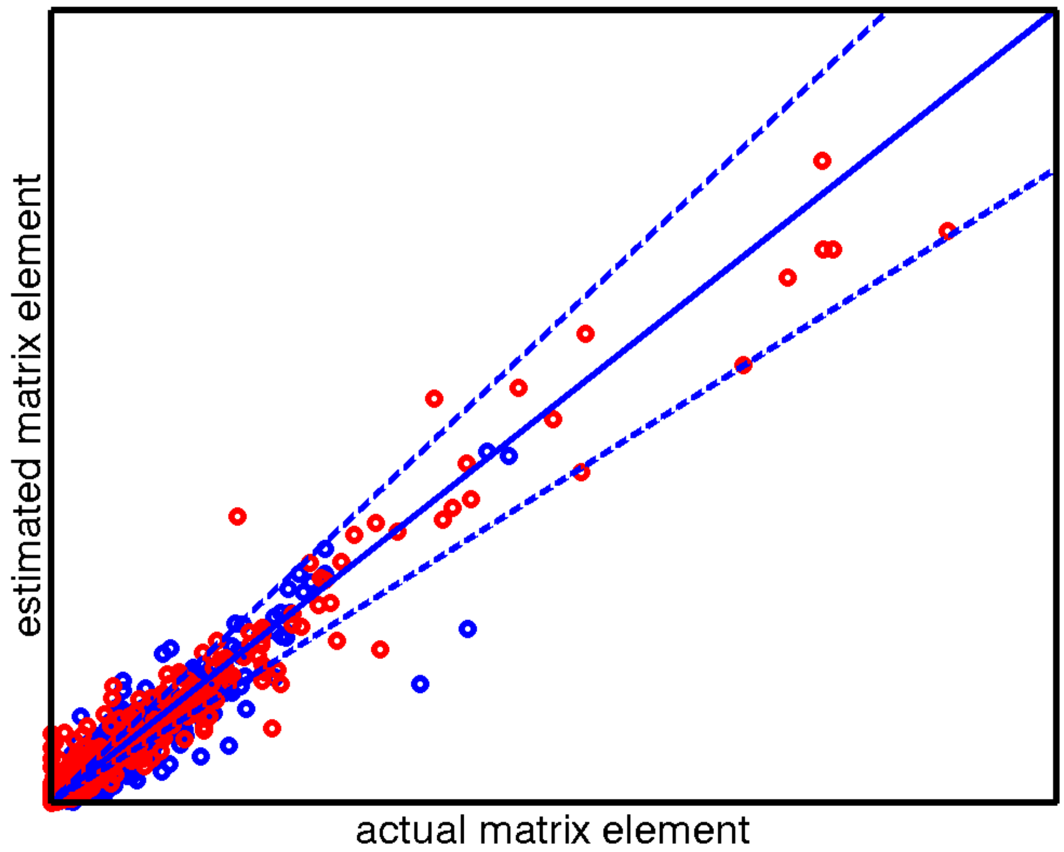

Conclusion

- Point-to-multipoint estimation is
  - Useful
    - Failure analysis of peering links
    - Failure analysis where IGP distance change closest exit point
  - Possible (from link stats)
- Results aren't yet as accurate as point-to-point
  - Need to run further experiments
  - Check parameters
  - Check we are comparing apples with apples