Traffic Matrix Estimation

Matthew Roughan

http://www.research.att.com/~roughan
<roughan@research.att.com>

Problem

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

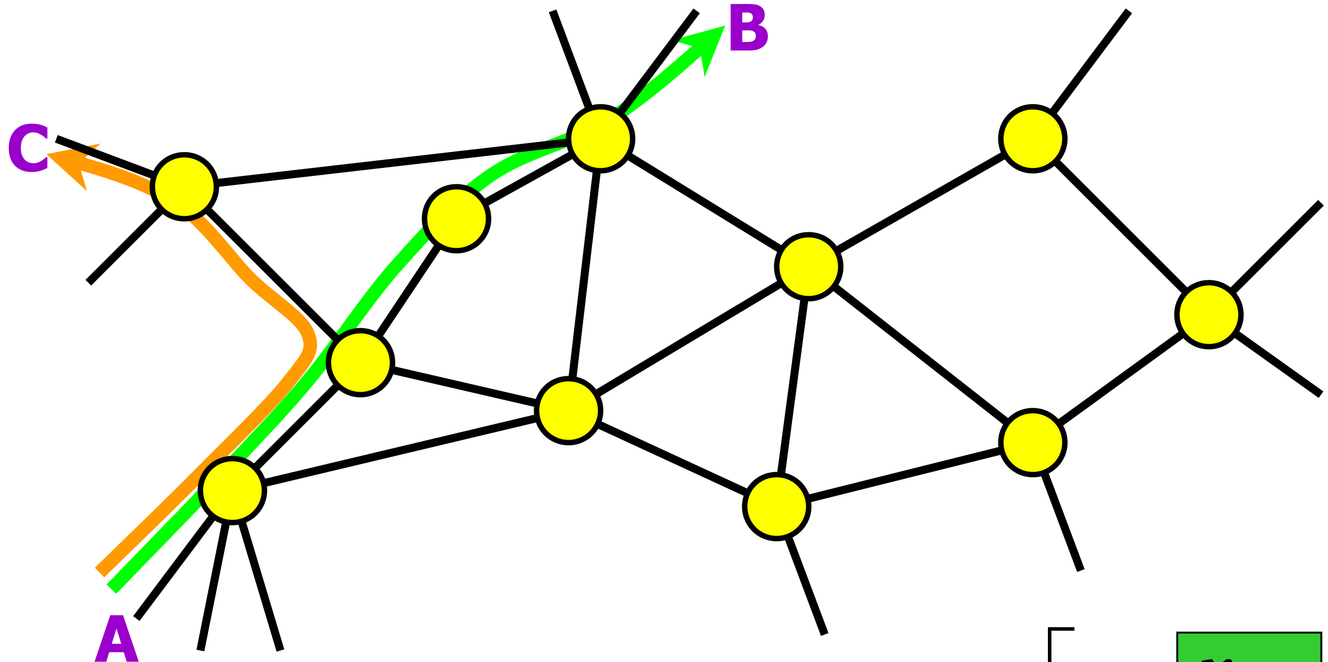

$$TM = \begin{bmatrix} - & x_{A,B} & x_{A,C} & \cdots \\ \cdot & \cdot & \cdot & \cdots \\ \cdot & \cdot & \cdot & \cdots \\ \cdot & \cdot & \cdot & \cdots \end{bmatrix}$$

Goals

- Need a traffic matrix
  - Capacity planning
  - Traffic engineering (choosing OSPF weights)
  - Reliability analysis
  - Detecting anomalies
  - Understanding traffic over the whole network
    - To run realistic simulations
- Don't have direct data
  - Netflow can provide direct estimates
  - Not currently available over whole edge of network
  - SNMP data is available over almost all network
- Want to use SNMP measurements to get a TM
  - Maybe we can also use Netflow where available?

Traffic Data

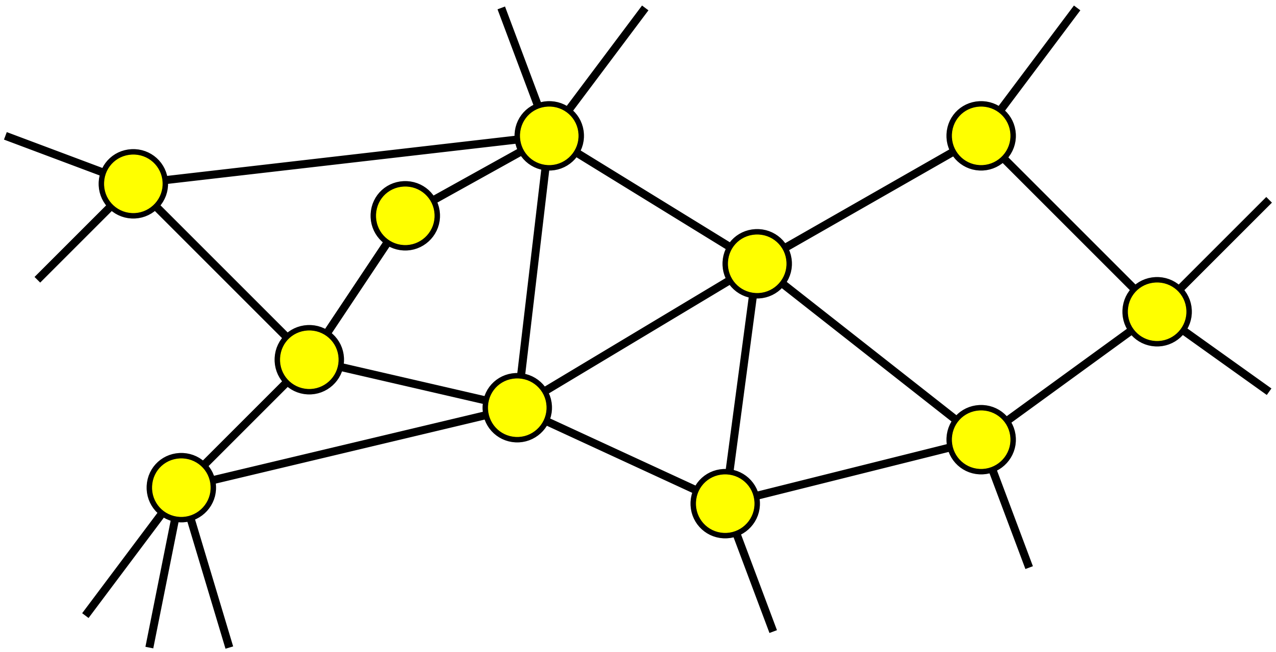

Data Availability - packet traces

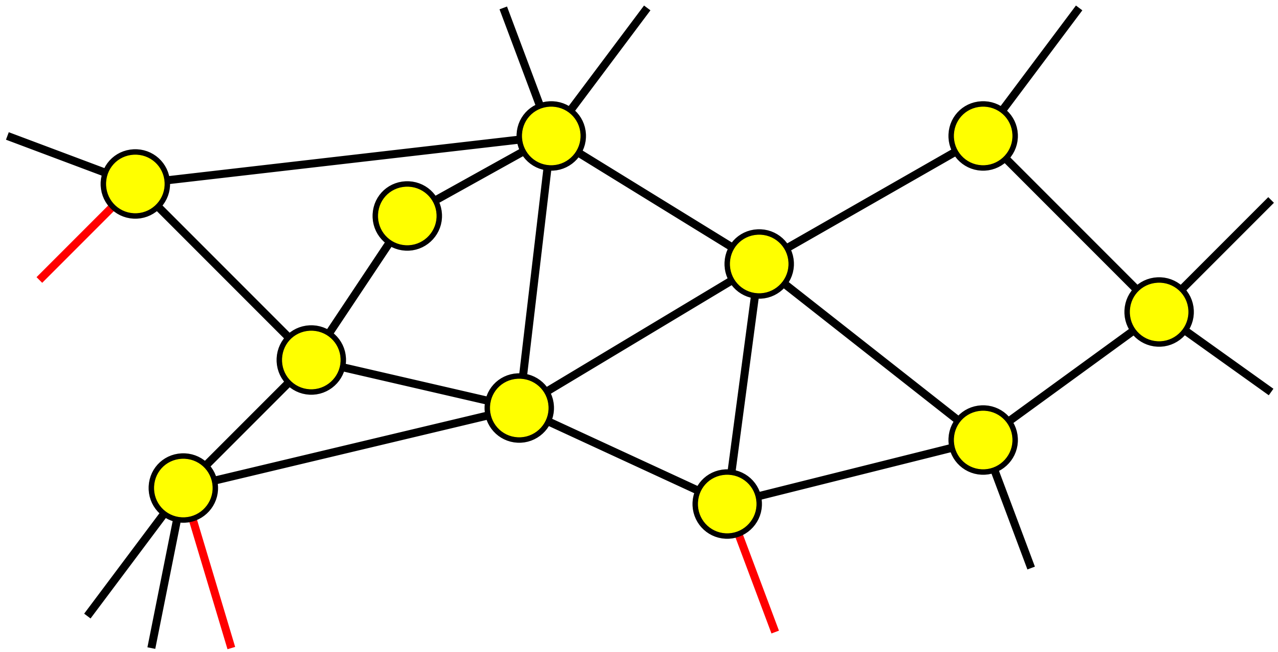

Packet traces limited availability – like a high zoom snap shot
• special equipment needed (O&M expensive even if box is cheap)
• lower speed interfaces (only recently OC48 available, no OC192)
• huge amount of data generated

Data Availability – flow level data

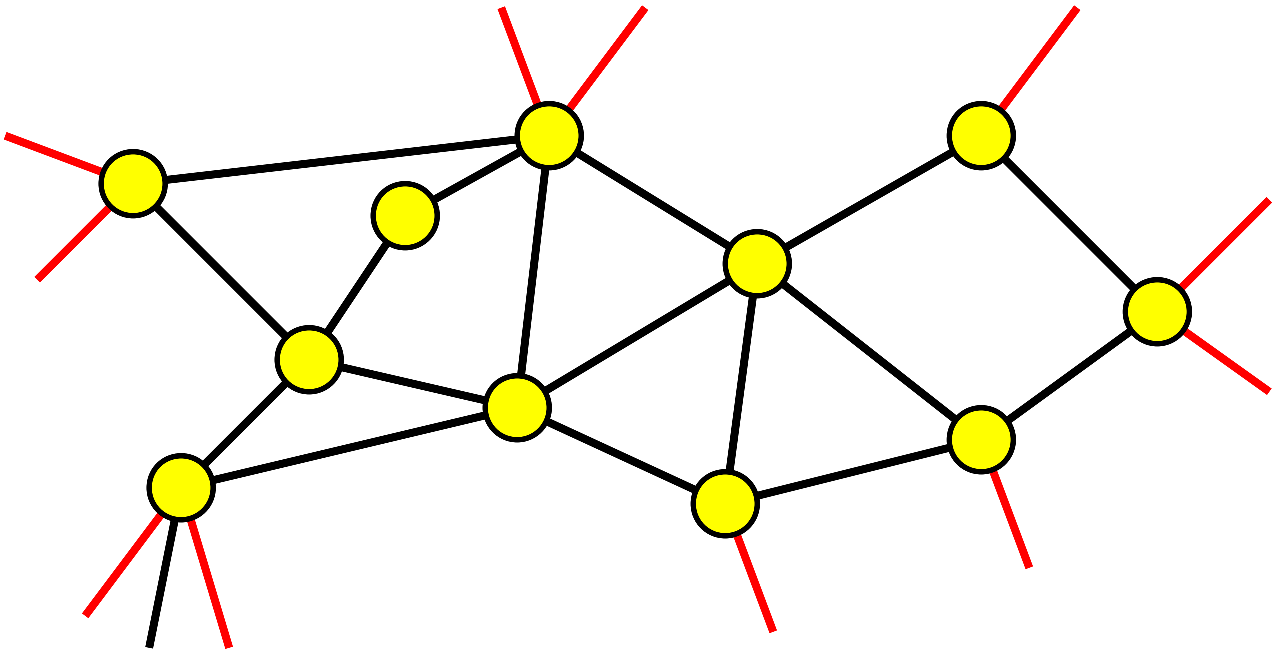

Flow level data not available everywhere – like a home movie of the network

- historically poor vendor support (from some vendors)
- large volume of data (1:100 compared to traffic)
- feature interaction/performance impact

Netflow Measurements

- Detailed IP flow measurements
  - Flow defined by
    - Source, Destination IP,
    - Source, Destination Port,
    - Protocol,
    - Time
  - Statistics about flows
    - Bytes, Packets, Start time, End time, etc.
  - Enough information to get traffic matrix
- Semi-standard router feature
  - Cisco, Juniper, etc.
  - not always well supported
  - potential performance impact on router
- Huge amount of data (500GB/day)

Data Availability - SNMP

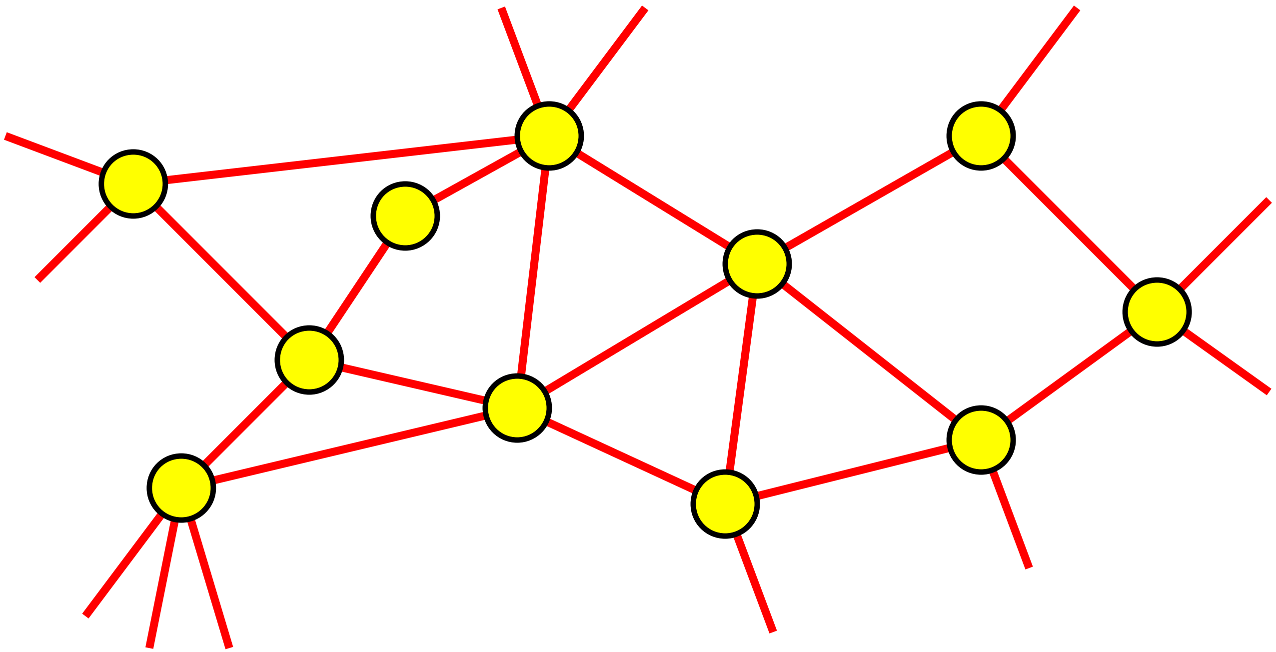

SNMP traffic data – like a time lapse panorama

- MIB II (including IfInOctets/IfOutOctets) is available almost everywhere
- manageable volume of data
- no significant impact on router performance

SNMP

- Pro
  - Comparatively simple, little infrastructure needed
  - Relatively low volume, low overhead
  - It is used already (lots of historical data)
- Con
  - Data quality – an issue with any data source
    - Ambiguous
    - Missing data
    - Irregular sampling
  - Octets counters only tell you link utilizations
    - Hard to get a traffic matrix
    - Can't tell what type of traffic
    - Can't easily detect DoS, or other unusual events
  - Coarse time scale – 5 min is typical

Topology and configuration

- Topology
  - Based on downloaded router configurations, every 24 hours
    - Links/interfaces
    - Location (to and from)
    - Function (peering, customer, backbone, …)
    - BGP configurations
- Routing
  - Forwarding tables (FIB)
  - OSPF weight and BGP table dumps
  - OSPF or BGP route monitors
- Routing simulations
  - Simulate IGP and BGP to get routing matrices
- Gives the Routing Matrix A

Example App: reliability analysis

Under a link failure, routes change
want to find an traffic invariant

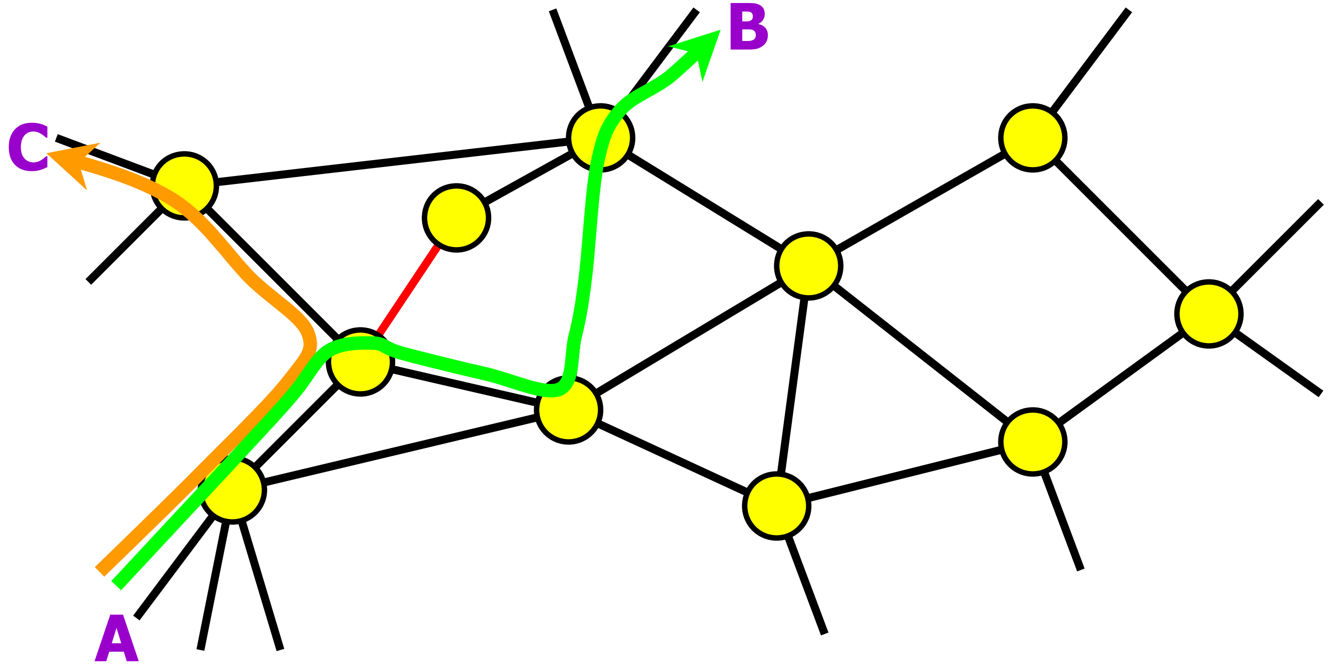

Example App: traffic engineering

- Route Optimization
  - Choosing route parameters that use the network most efficiently
  - Measure efficiency by maximum utilization
- Methods
  - Shortest path IGP weight optimization
    - OSPF/IS-IS
    - Choose 'weights'
  - Multi-commodity flow optimization
    - Implementation using MPLS
    - Explicit route for each origin/destination pair

Mathematical Formalism

Only measure traffic at links

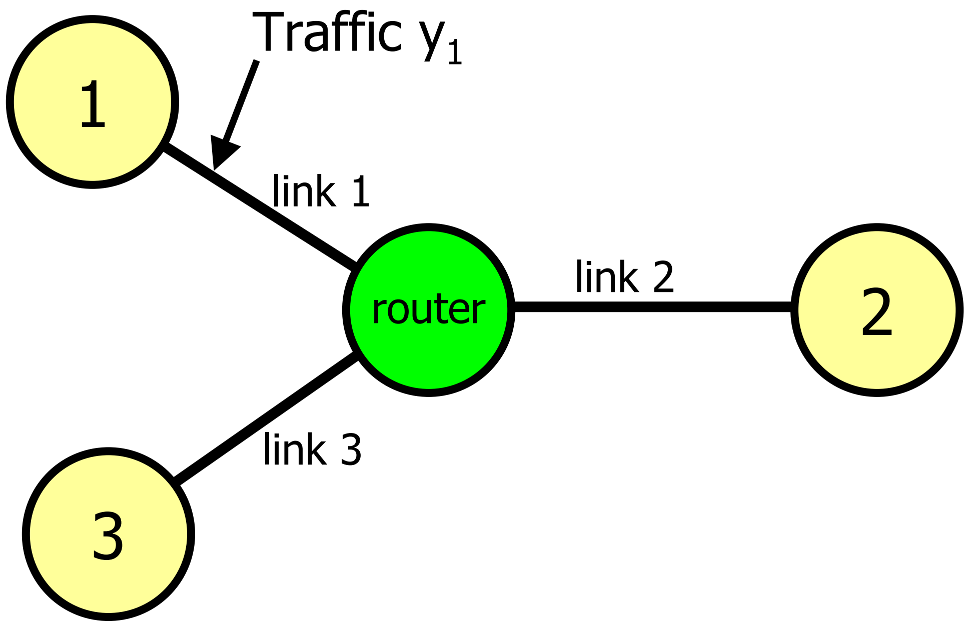

Mathematical Formalism

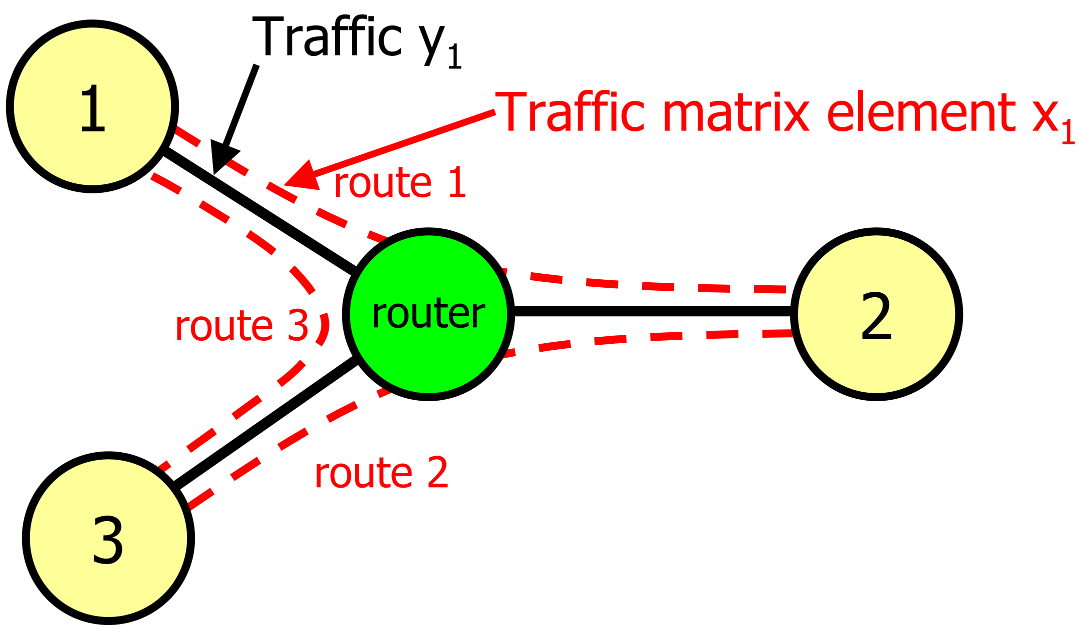

Problem: Estimate traffic matrix (x's) from the link measurements (y's)

Mathematical Formalism

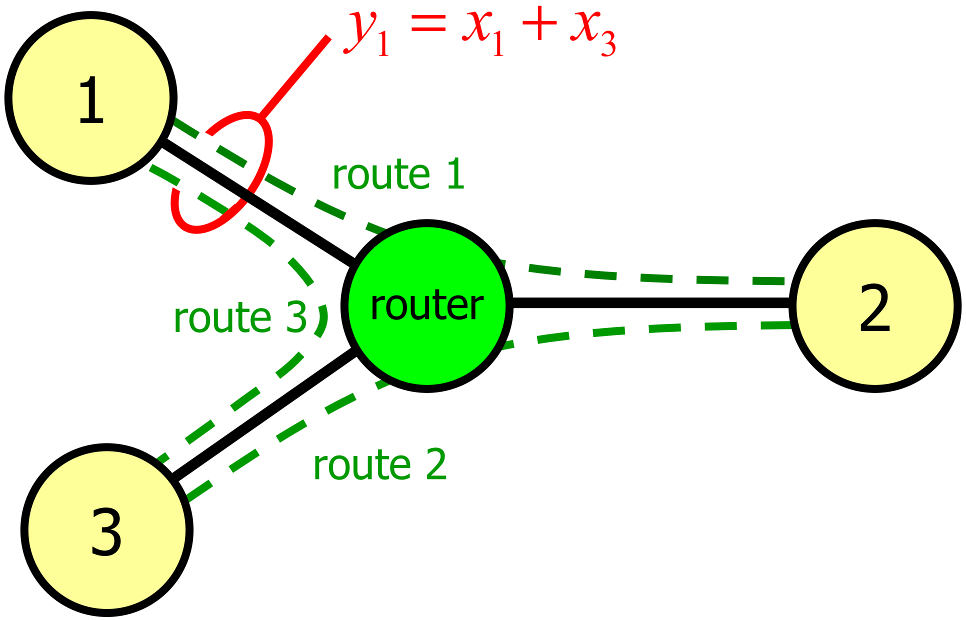

Problem: Estimate traffic matrix (x's) from the link measurements (y's)

Mathematical Formalism

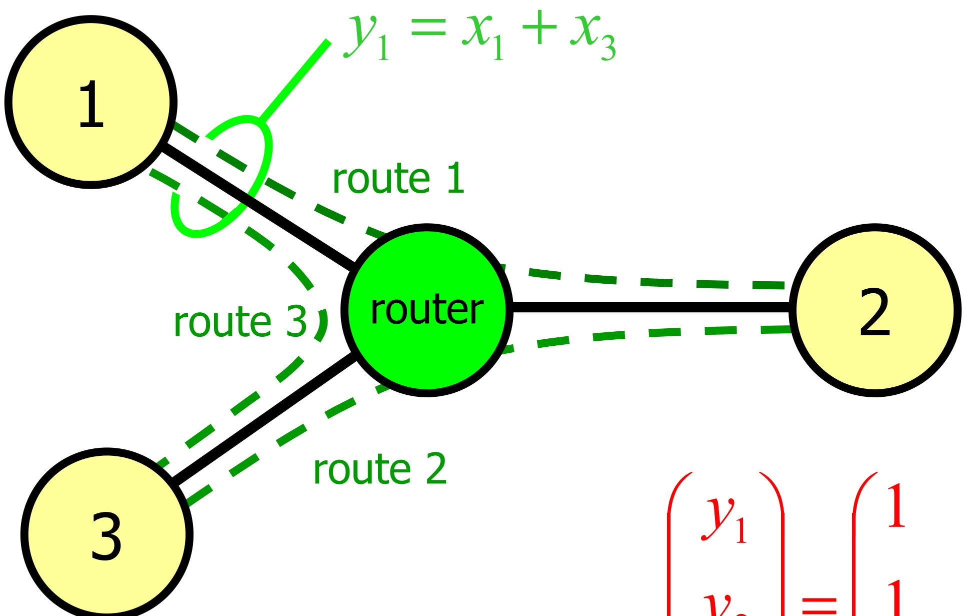

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$$

Problem: Estimate traffic matrix (x's) from the link measurements (y's)

Mathematical Formalism

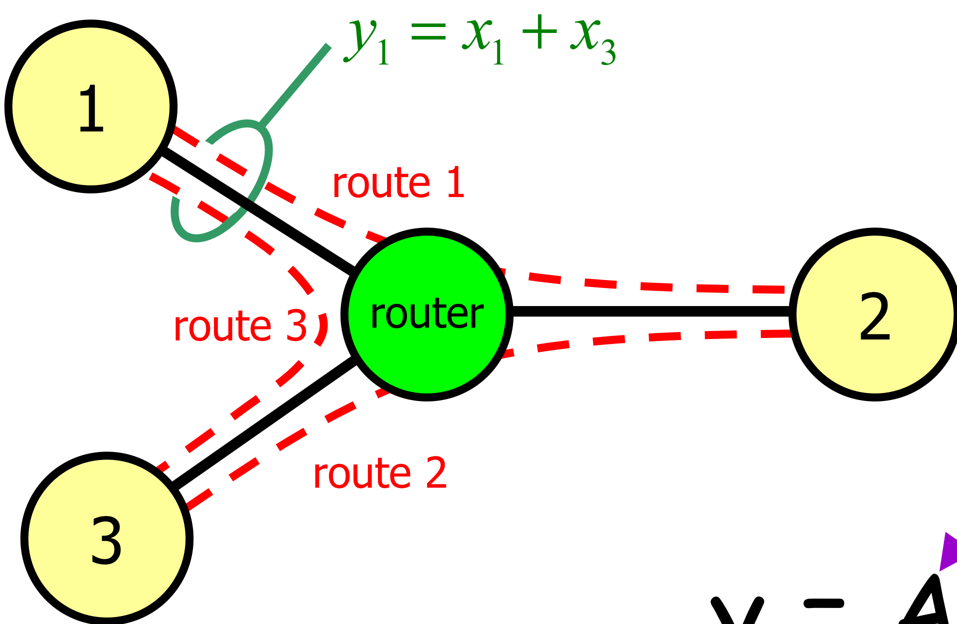

Routing matrix

$$y = A^T x$$

For non-trivial network
UNDERCONSTRAINED

Approaches to TM estimation

- Measurement
  - "Deriving traffic demands for operational IP networks: methodology and experience", A.Feldman, A.Greenberg, C.Lund, N.Reingold, J.Rexford, F.True, ACM SIGCOMM 2000.
- MLE/EM
  - "Network tomography: estimating source-destination traffic intensities from link data", Y.Vardi, J.Am.Statist.Assoc., 91, pp. 265—377, 1996.
  - "Time-varying network tomography: router link data", J.Cao, D.Davis, S.V.Wiel and B.Yu, J.Am.Statist.Assoc., 95, pp. 1063—1075, 2000.
- Bayesian
  - "Bayesian inference on network traffic using link count data", C.Tebaldi, and M.West, J.Am.Statist.Assoc., 93, pp. 557—576.
  - "Iterative Bayesian Estimation of the Origin Destination Traffic Matrix", Sandrine Vaton, INTiMaTE, Paris, France, 2003.
  - "Network Tomography: an iterative Bayesian analysis", Proceedings of the International Teletraffic Congress (ITC-18) 2003.
- Choice models/gravity
  - "Traffic matrix estimation: existing techniques and new directions", A.Medina, N.Taft, Ksalmatian, S.Bhattacharyya, and C.Diot, ACM SIGCOMM, 2002.
  - "Experience measuring backbone traffic variability: models, metrics, measurements and meaning", M.Roughan, A.Greenberg, C.Kalmanek, M.Rumsewicz, J.Yates and Y.Zhang, abstract in ACM SIGCOMM Internet Measurement Workshop, 2002.
- Minimum Mututal Information (MMI)
  - "Fast, accurate computation of large-scale IP traffic matrices from link measurements", Y.Zhang, M.Roughan, N.Duffield and A.Greenberg, ACM SIGMETRICS 2003.
  - "An information theoretic approach to traffic matrix estimation", Y.Zhang, M.Roughan, C.Lund and D.Donoho, ACM SIGCOMM 2003.

Maximum likelihood estimation

- Assume a particular model for the traffic
  - Vardi => Poisson
  - Cao et al => Gaussian
- From the model, infer relationship between Mean and variance:
  - Poisson: mean = variance
  - Gaussian: mean $\propto$ variance^c, $c = 1$, or 2
- Use the relationship to derive extra equations
  - Problem is no longer underconstrained
  - May actually be over-constrained
- Trick is then efficient estimation
  - EM algorithm
  - Iterative Proportional Fitting

Bayesian

- Start with a prior model
  - E.G. Poisson model
- Standard Bayesian inference
  - MCMC, Gibb's sampling
- More recent work (Vaton and Gravey)
  - Uses more sophisticated prior models
    - Multi-level model (Markov modulated Poisson process)

Gravity/choice model

$$T_{SD} = \frac{R_S \ A_D}{f_{SD}}$$

$$R_S \quad = \quad \text{repulsion}$$
$$A_D \quad = \quad \text{attraction}$$
$$f_{SD} \quad = \quad \text{friction}$$

Simple Gravity

$$T_{SD} = \frac{T_S^{in} \; T_D^{out}}{T}$$

T_S^{in} = traffic into the network at S

T_D^{out} = traffic out of the network at D

T = total traffic

Simple gravity continued

Equivalent

- Simple gravity
- Independent S and D
  - $P(S,D) = p(S)\,p(D)$
  - $P(D|S) = P(D)$
- Mutual information between S and D is zero
  - $I(S,D) = 0$

Simple gravity is not great

- Not terrible either (very simple)
- Only uses edge data
- Can be improved using conditional independence
  - Model hot-potato routing

Minimum Mutual Information (MMI)

- Mutual Information I(S,D)=0
  - Information gained about S from D

$$I(S,D)=\sum_{s,d} p(S=s,D=d)\log\frac{p(S=s,D=d)}{p(S=s)\,p(D=d)}$$

  - I(S,D) = relative entropy with respect to independence
  - Can also be given by Kullback-Leibler information divergence
- Why this model
  - In the absence of information, let's assume no information
    - Example of maximum entropy principle
    - When you have data, minimize subject to constraints
  - Minimal assumption about the traffic
  - Large aggregates tend to behave like overall network

MMI in practice

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

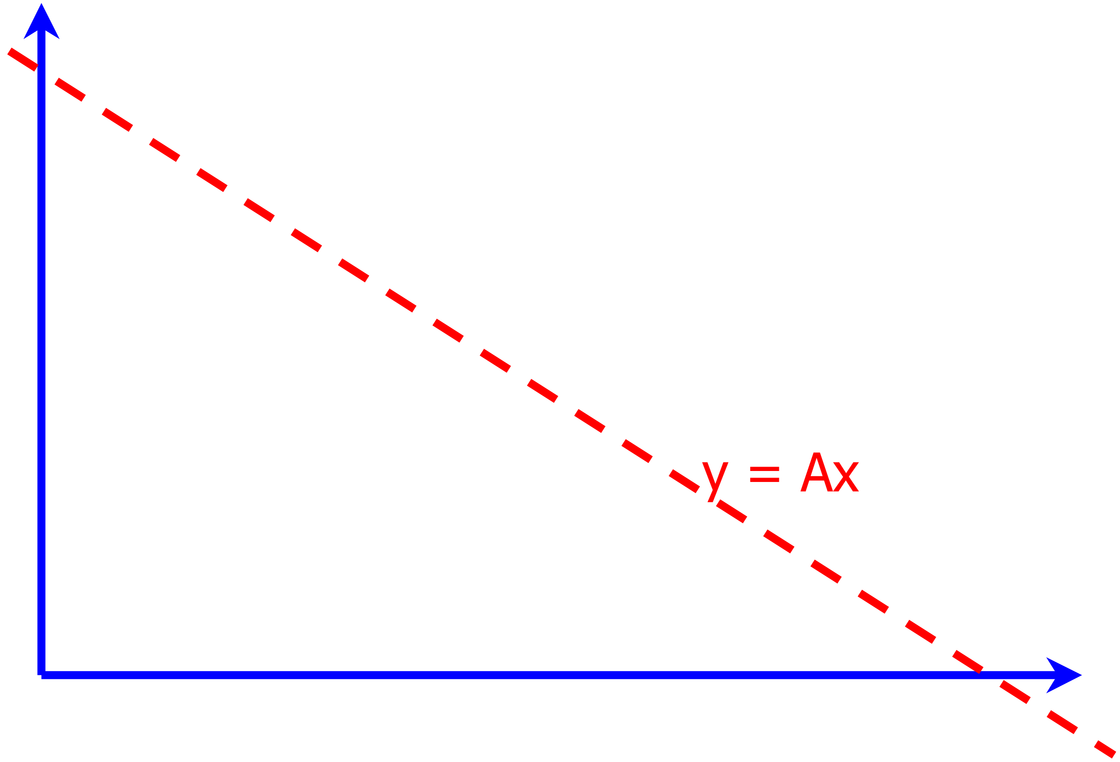

MMI in practice

- Independence gives us a starting point

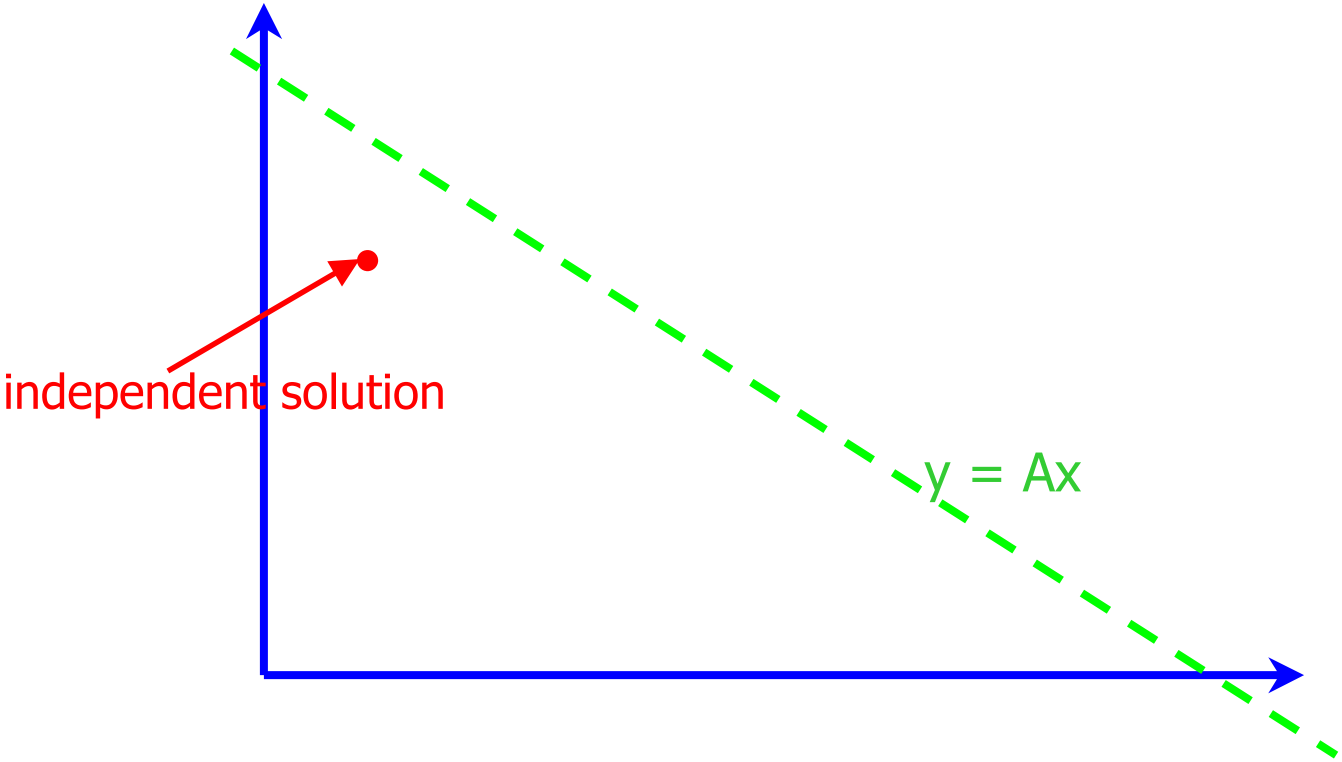

MMI in practice

- Find a solution which
  - Satisfies the constraint
  - Is *closest* to the independent solution

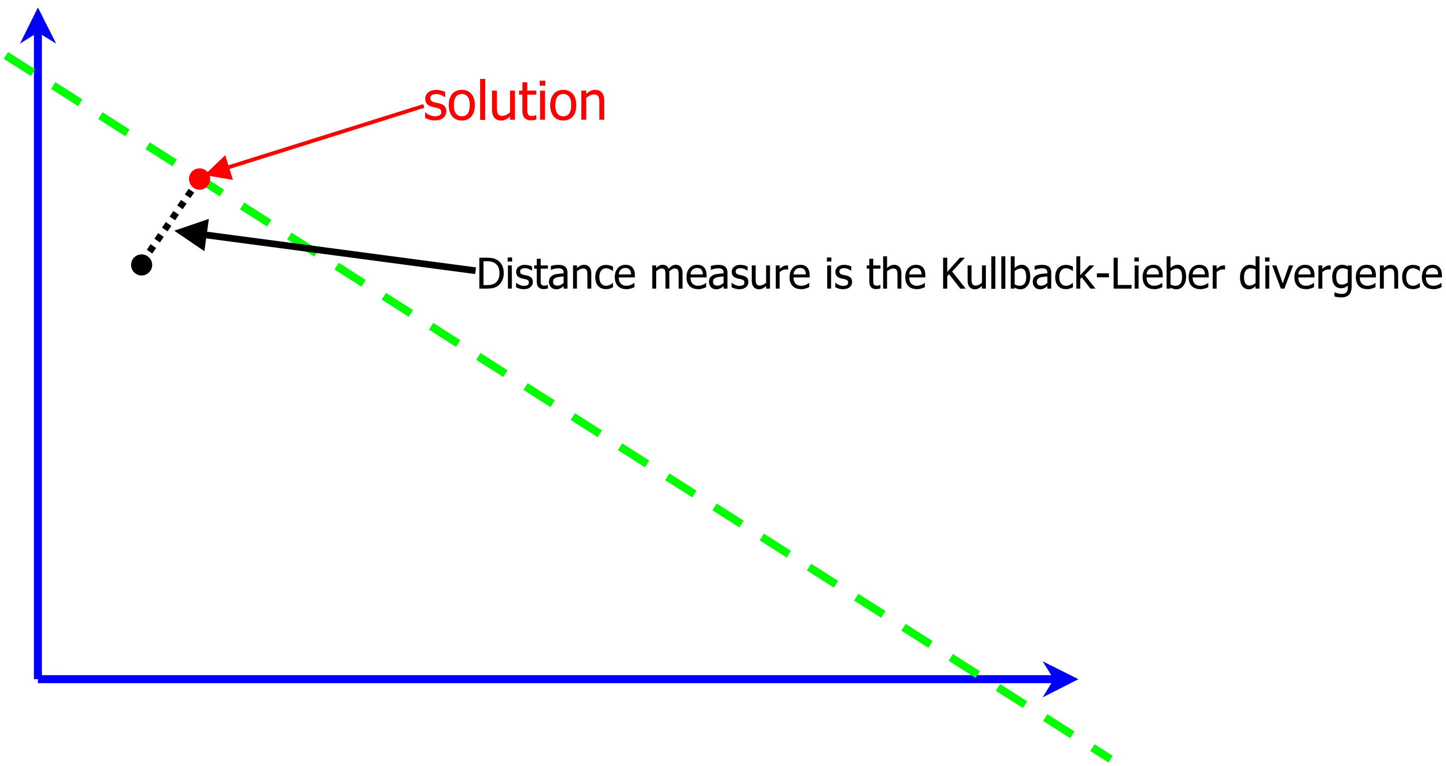

Complications

- Level of aggregation
  - Prefix to prefix O(100k)
  - Ingress-link to egress-link O(10k)
  - Ingress-router to egress-router O(1k)
  - Backbone-router to backbone-router O(100)
  - PoP to PoP O(10)
- Hot-potato routing
- Point-to-multipoint

Hot potato routing

- Internet routing is asymmetric
- A provider can control exit points for traffic going to peer networks

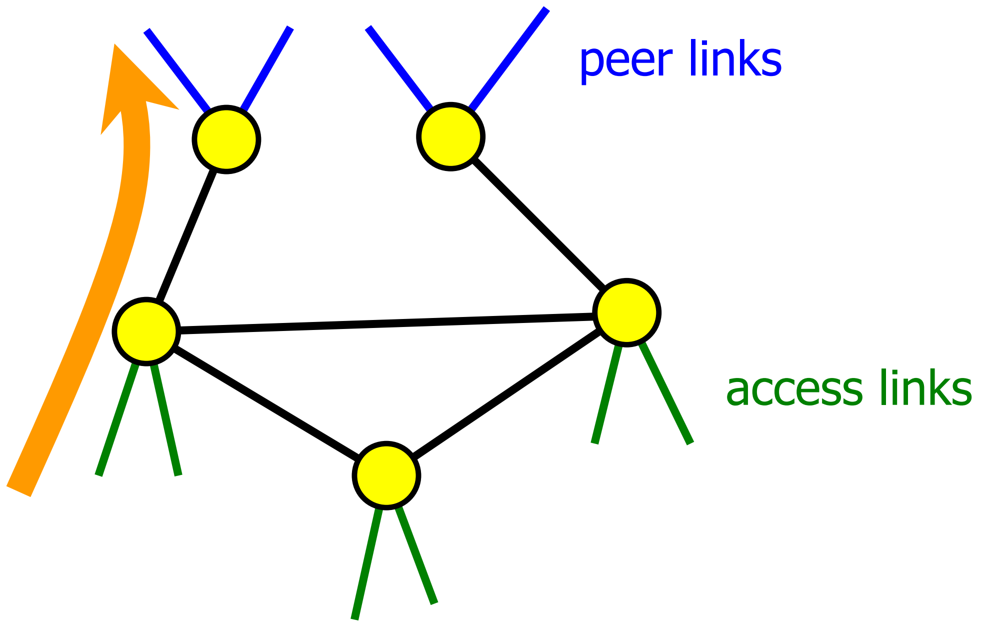

Hot potato routing

- Internet routing is asymmetric
- A provider can control exit points for traffic going to peer networks
- Have much less control of where traffic enters

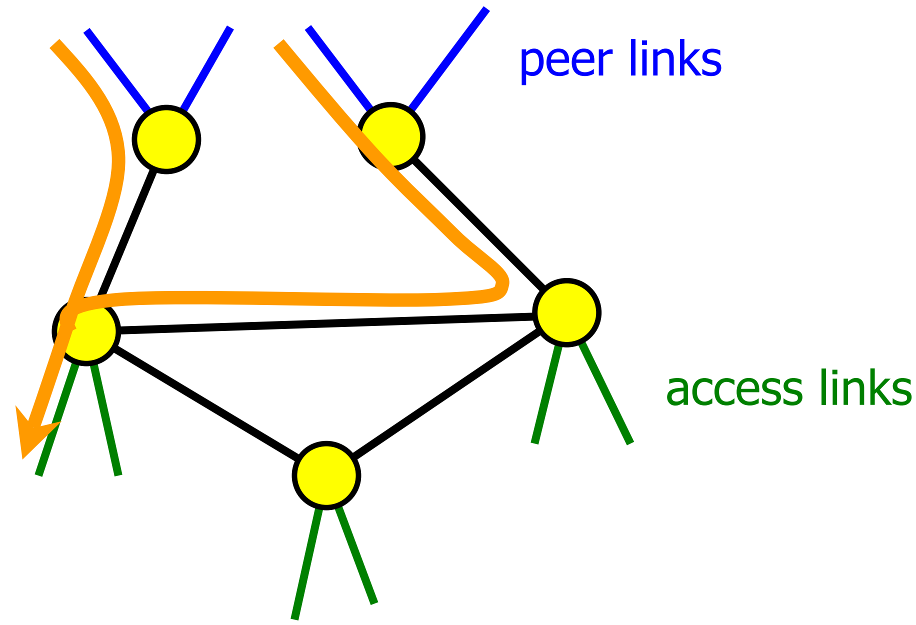

Point-to-Multipoint

- 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

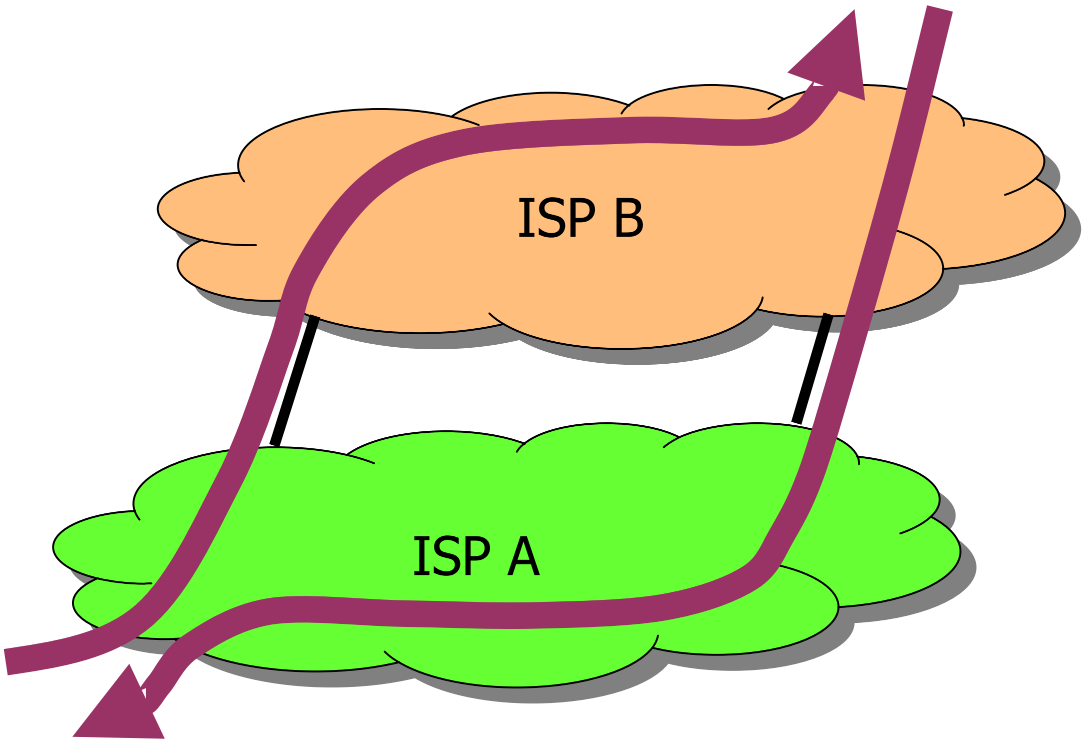

Peering link failure

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

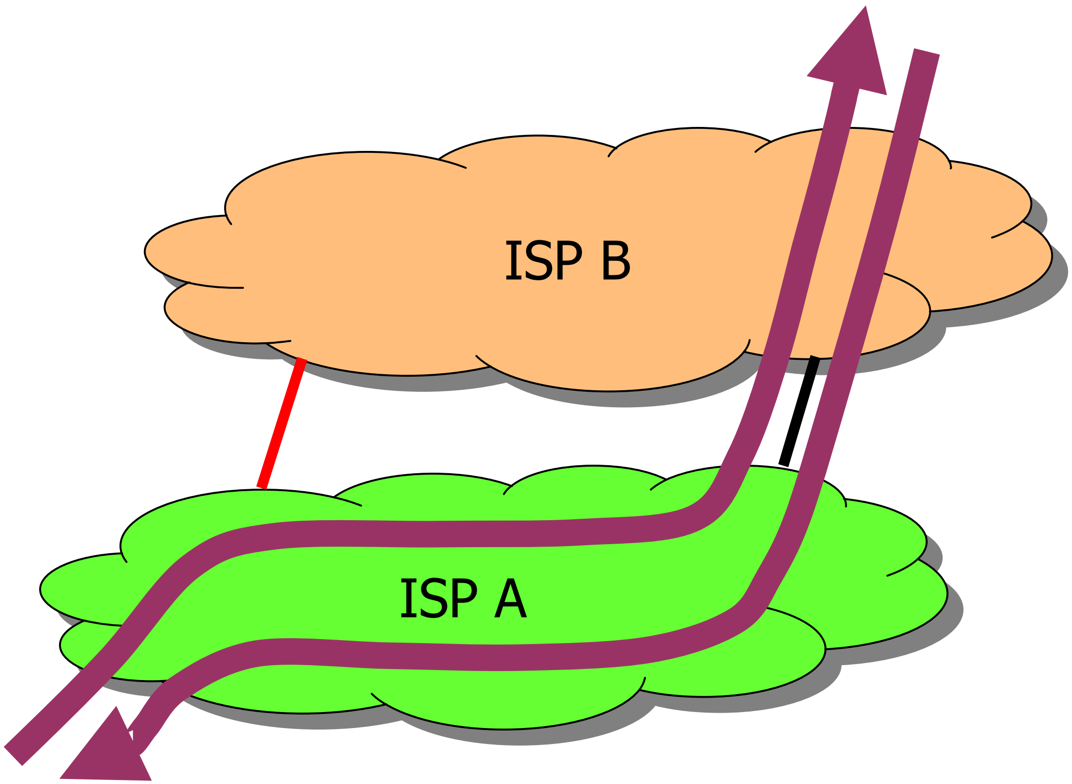

Point-to-multipoint

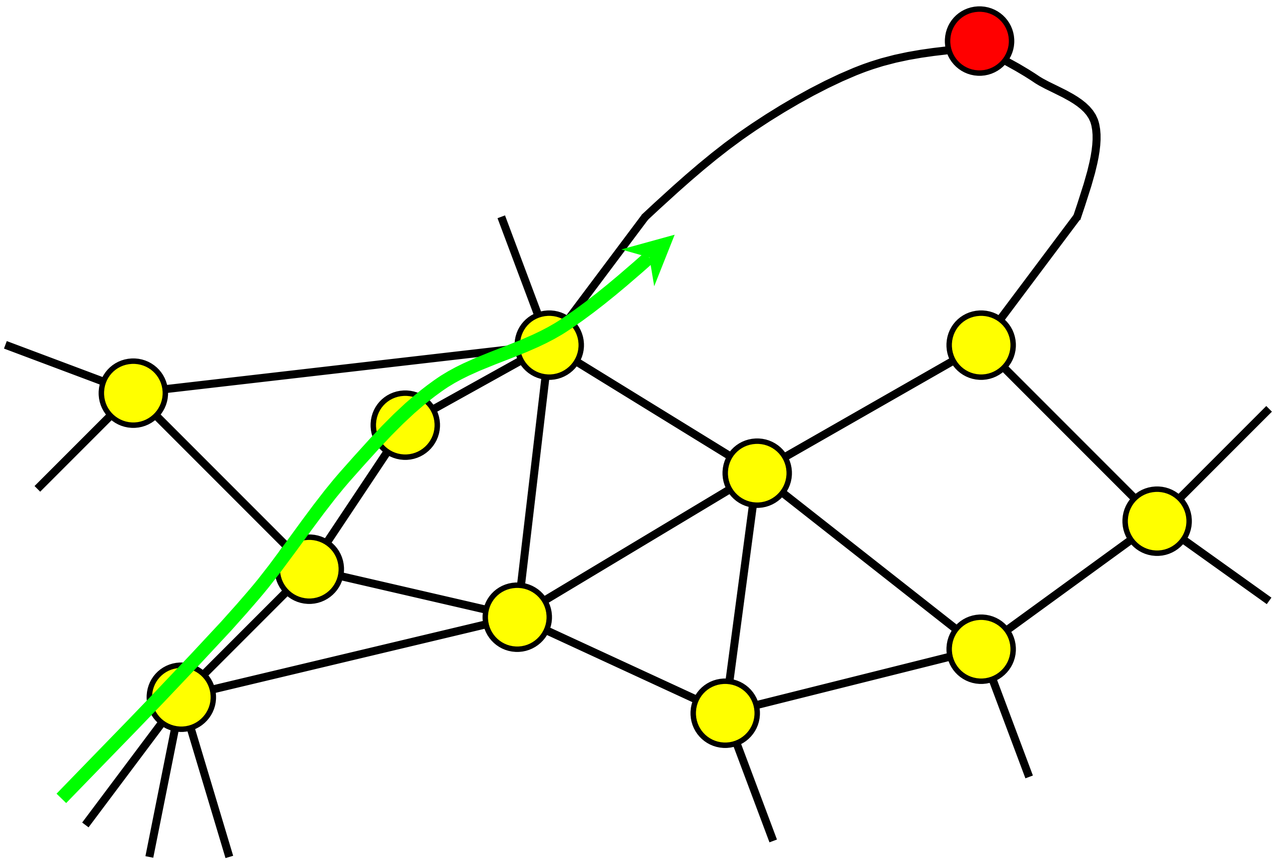

Point-to-multipoint

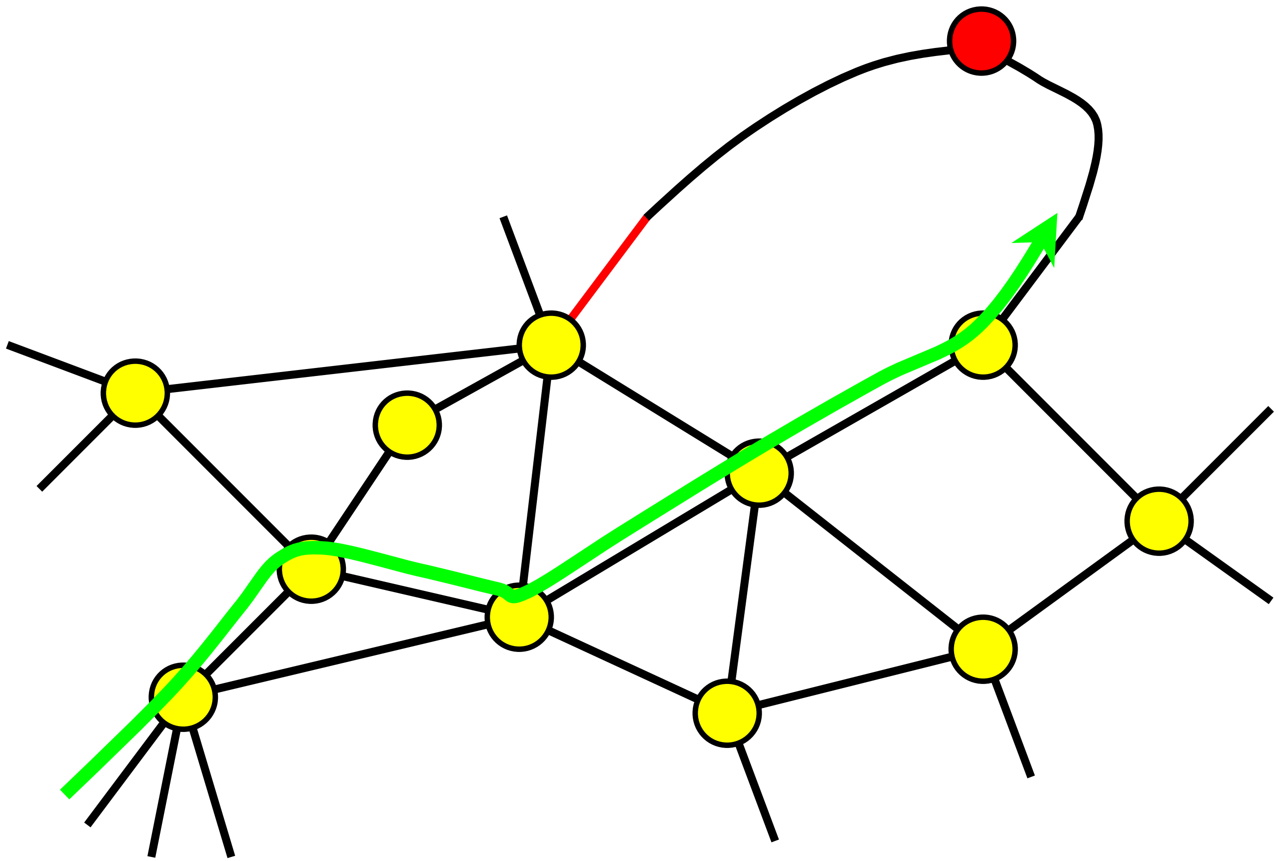

Conclusion

- Problem:
  estimate end-to-end demands from link measurements

- Several methods available
- There has been limited cross-comparison
  - Lack of common
    - Data sources
    - Implementations
    - Data formats

- Abilene data (soon)