How to Compute Accurate Traffic Matrices for your Network in Seconds

Yin Zhang, Matthew Roughan, Carsten Lund, Nick Duffield, Albert Greenberg, Quynh Nguyen
– AT&T Labs-Research

David Donoho – Stanford

AT&T Shannon Lab

Problem

Have link traffic measurements (from SNMP)

Want to know demands from source to destination

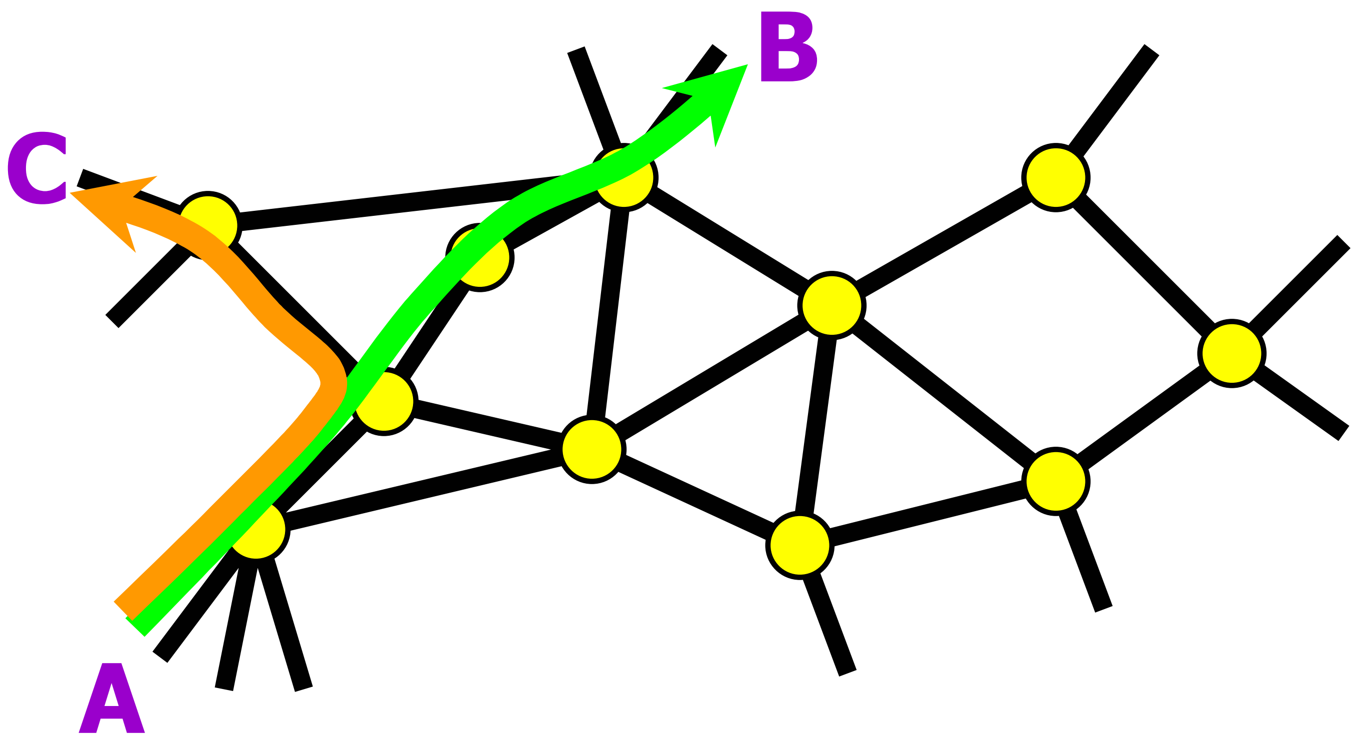

Example App: reliability analysis

Under a link failure, routes change
want to predict new link loads

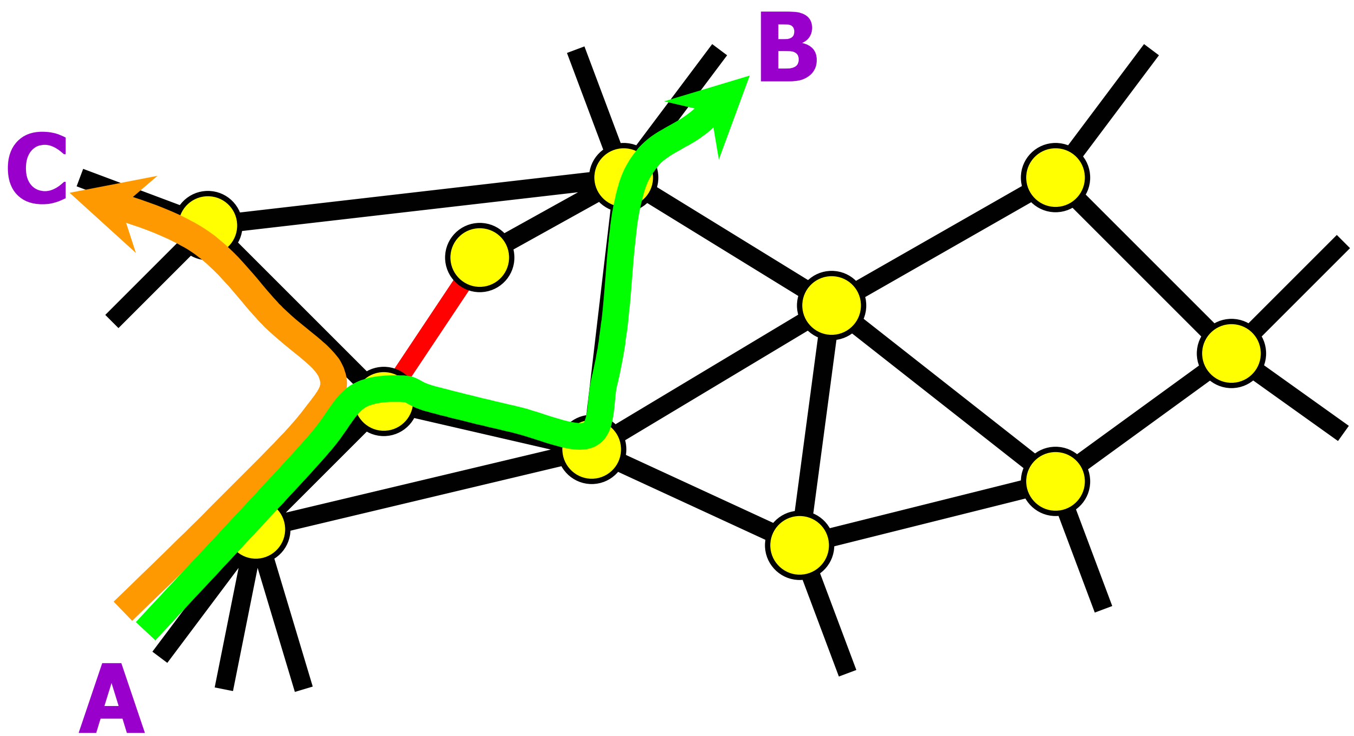

Network Engineering

- What you want to do
  - Reliability analysis
  - Traffic engineering
  - Capacity planning
- What do you need to know
  - ✓ Network and routing
  - ✓ Prediction and optimization techniques
  - ? Traffic matrix

Solution: Tomo-gravity

- Computes traffic matrices
  - input: SNMP, topology, routing policies
- Advantages
  - Today's data → no special instrumentation
  - Fast: a few seconds
  - Accurate: average 12% error
  - Scales: hundreds of nodes
  - Robust: copes easily with data glitches
  - Flexible: can incorporate more detailed data

Tomo-gravity

Tomography
- Astronomy
- Seismology
- Medical imaging
  - CAT in CATSCAN

+

Gravity modeling
- Econometrics
- Transportation
  - planes, trains, automobiles

Foundation: Information Theory

Tomo-gravity in a Nutshell

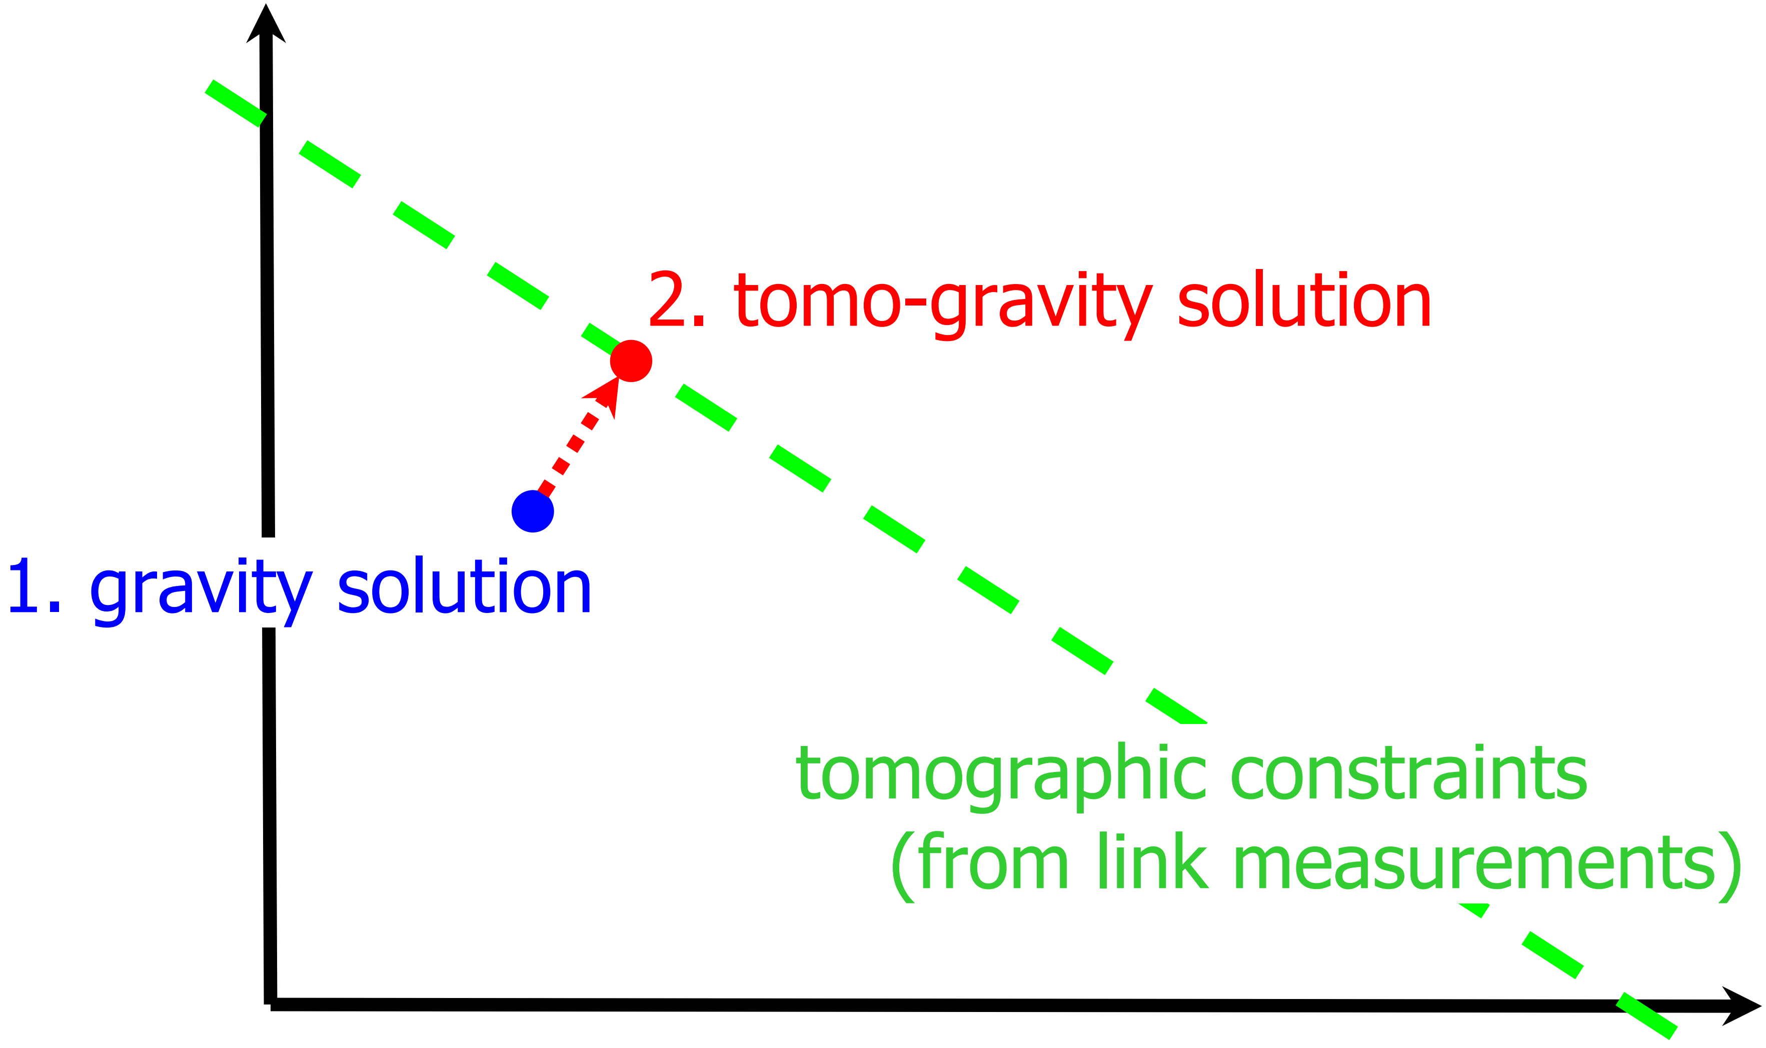

Tomo-gravity in practice

1. Get topology & routing
2. Measure SNMP link loads
3. Derive gravity solution
   - Uses edge loads
4. Compute tomo-gravity solution
   - Use internal link data
   - Matches observed link loads
   - Can incorporate more detailed measurements to boost accuracy

Real example

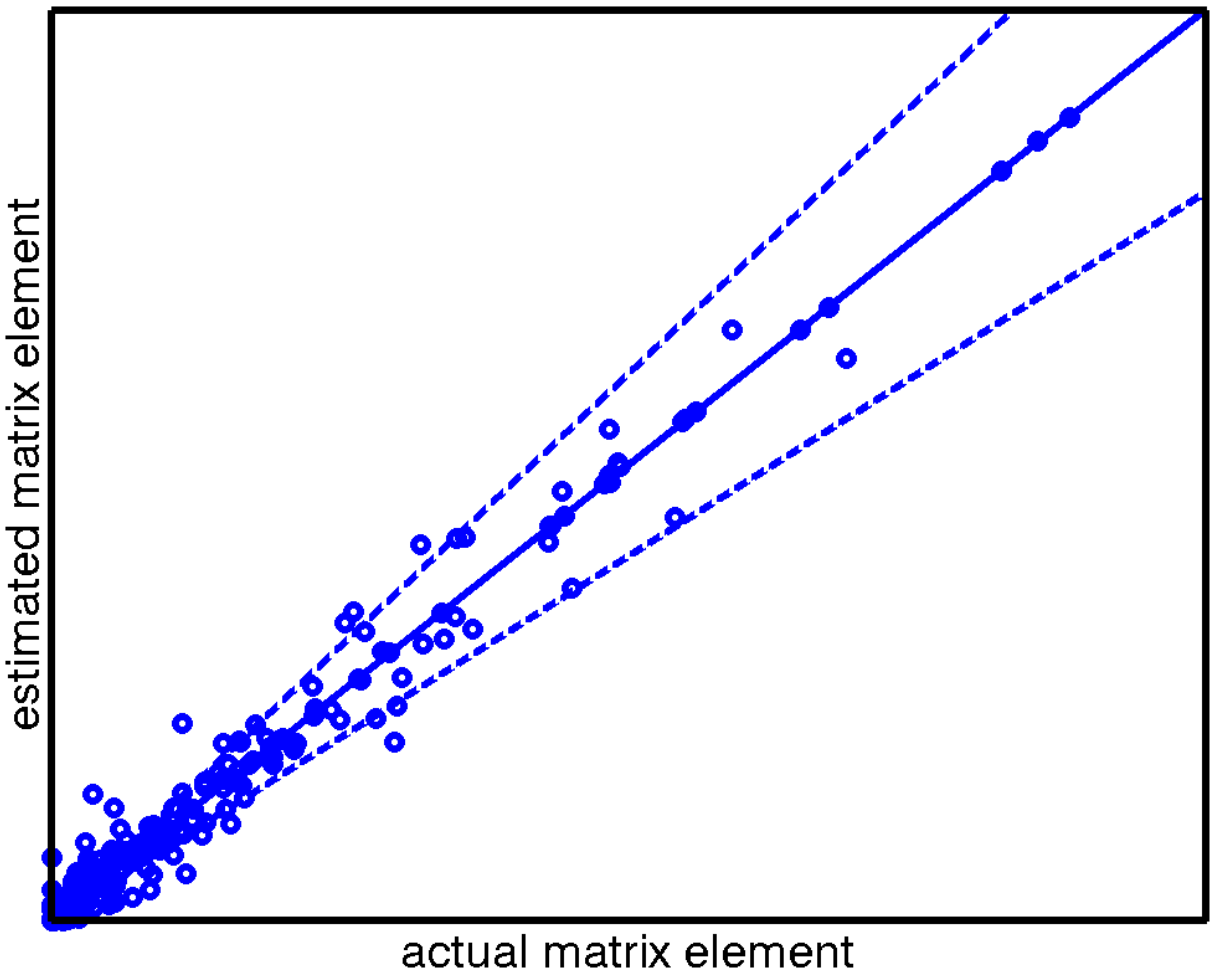

Tomo-gravity

Example use: reliability analysis

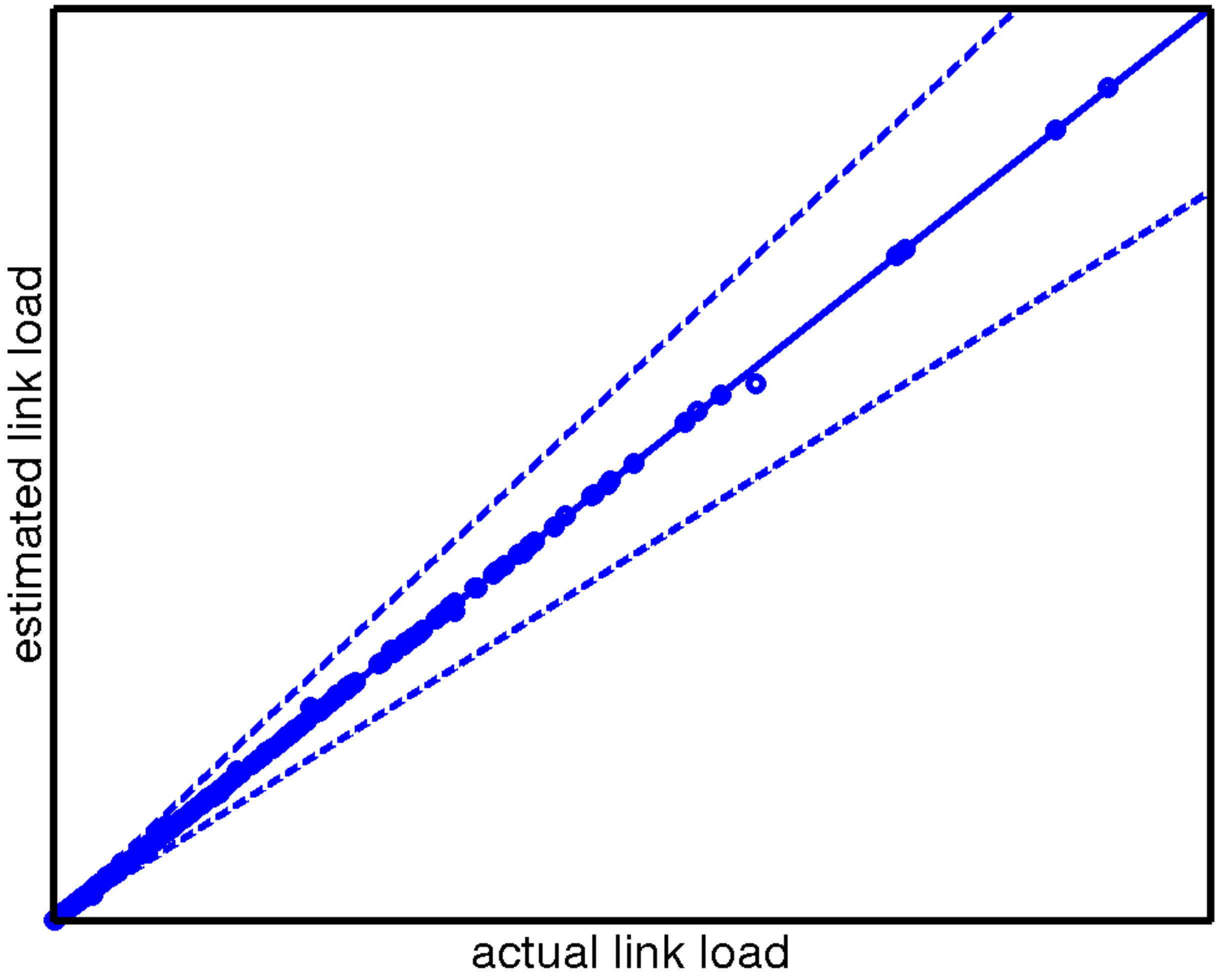

Tomo-gravity

Conclusion

- Tomo-gravity implemented
  - AT&T's IP backbone (AS 7018)
  - Hourly traffic matrices for > 1 year (in secs)
- For a number of applications
  - Reliability analysis (killer app...)
  - Traffic engineering
  - Capacity planning

http://www.research.att.com/~roughan/tomogravity.html

Key References

- "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.
- Both available at http://www.research.att.com/~roughan/papers.html

Additional Slides

Mathematical Formalism

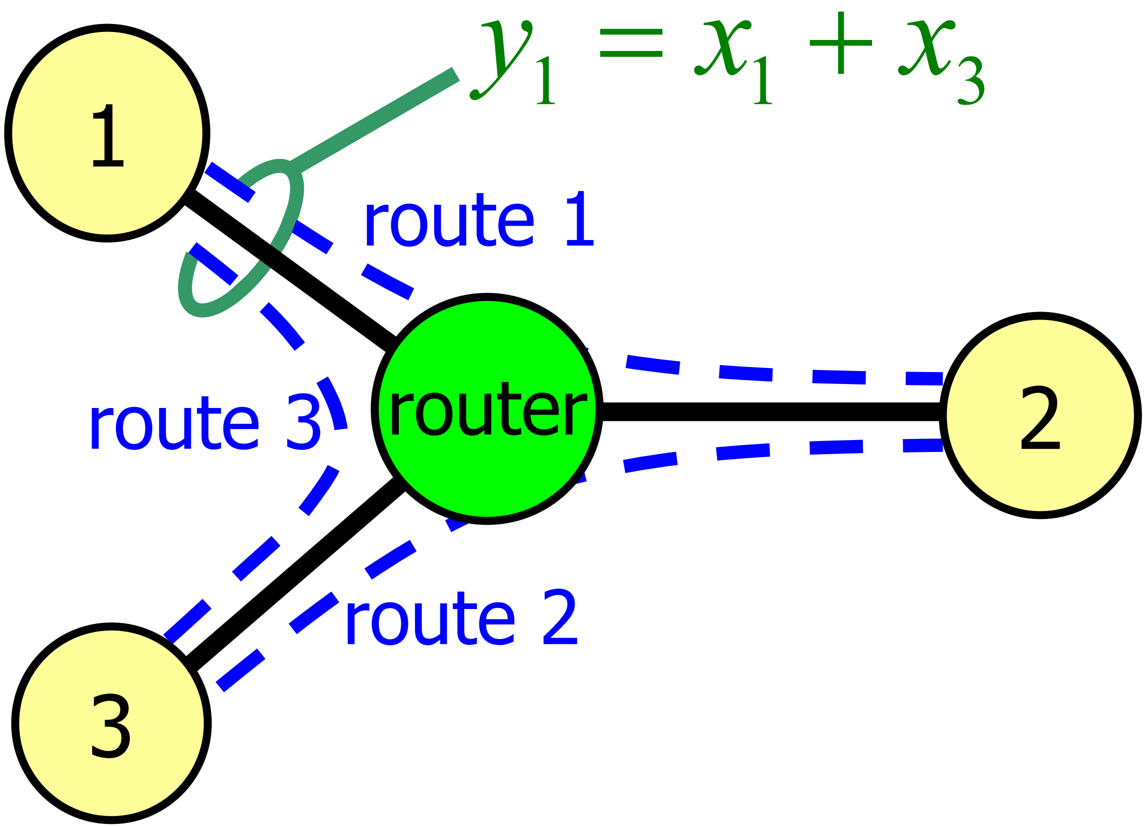

Equations

Traffic matrix

$$y = Ax$$

Link measurements

Routing matrix

Many more unknowns than measurements

Robustness (input errors)

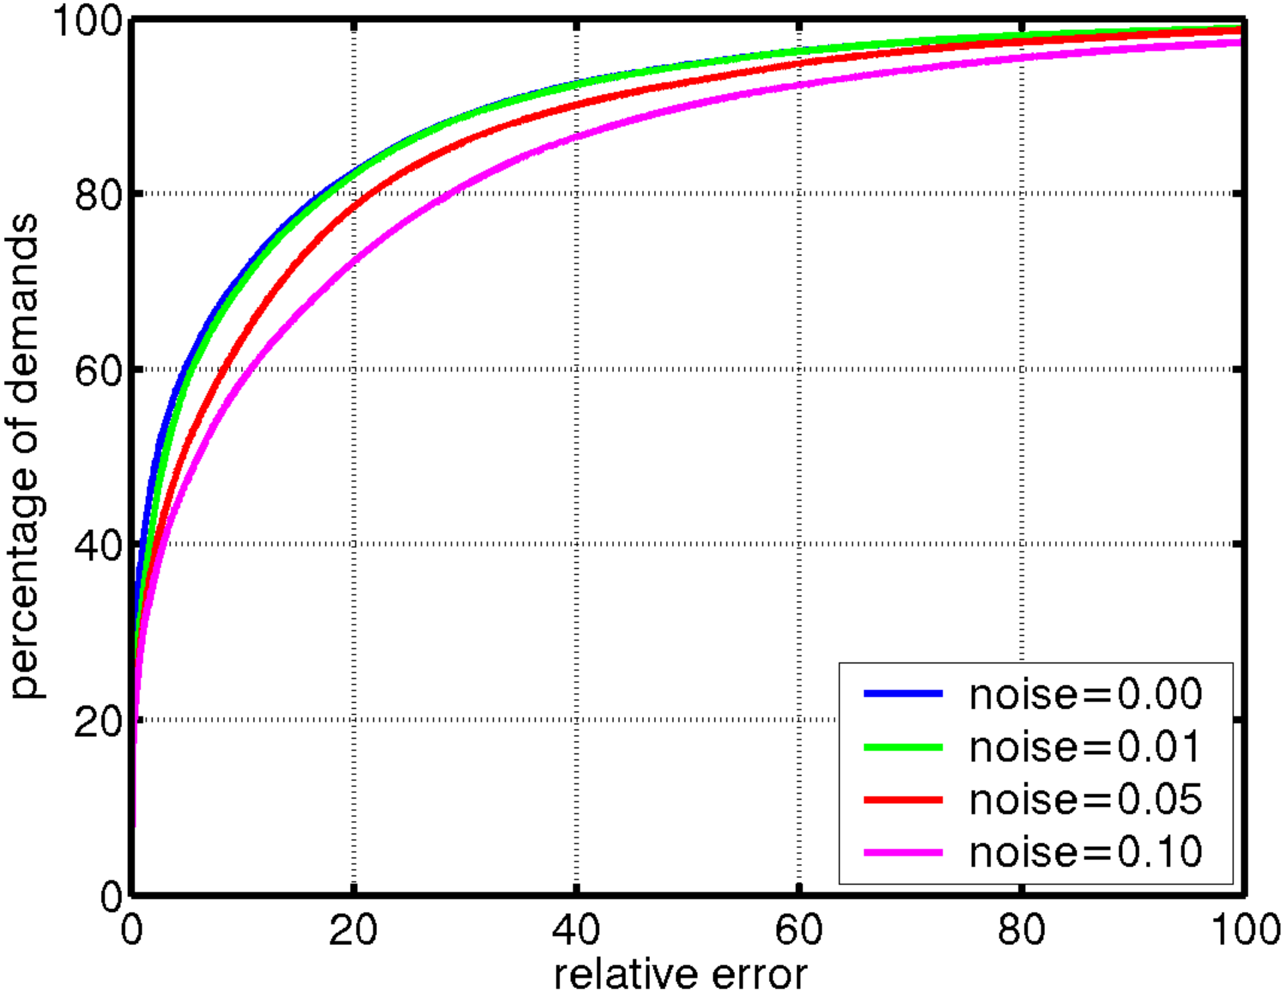

Tomo-gravity

Dependence on Topology

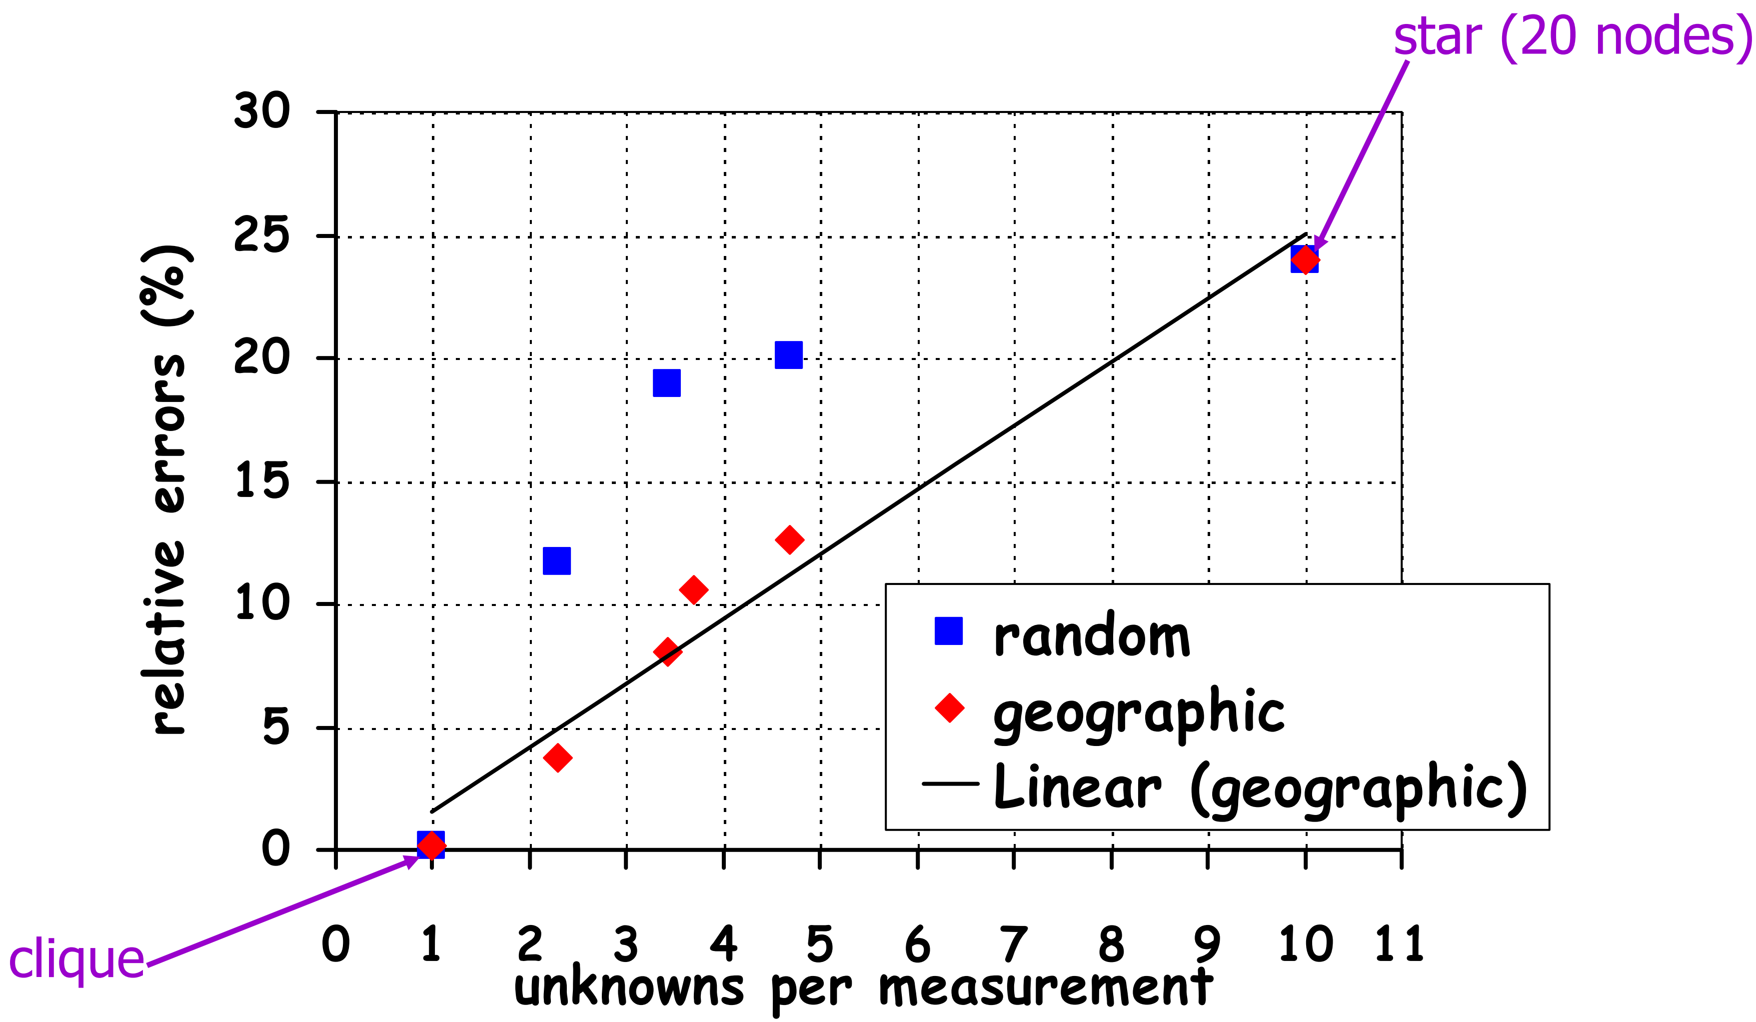

Additional information - Netflow

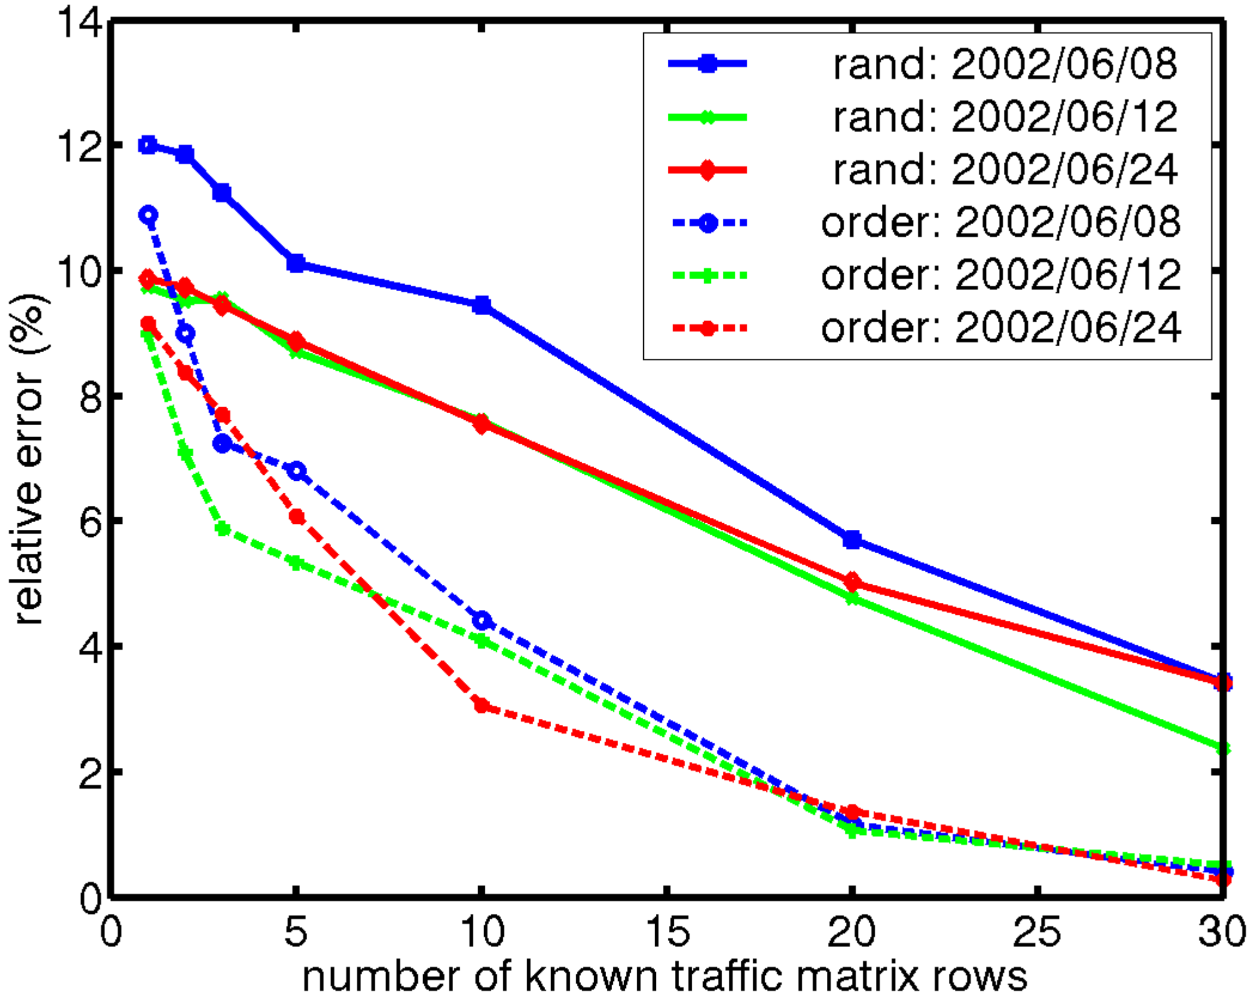

Tomo-gravity

Local traffic matrix (George Varghese)

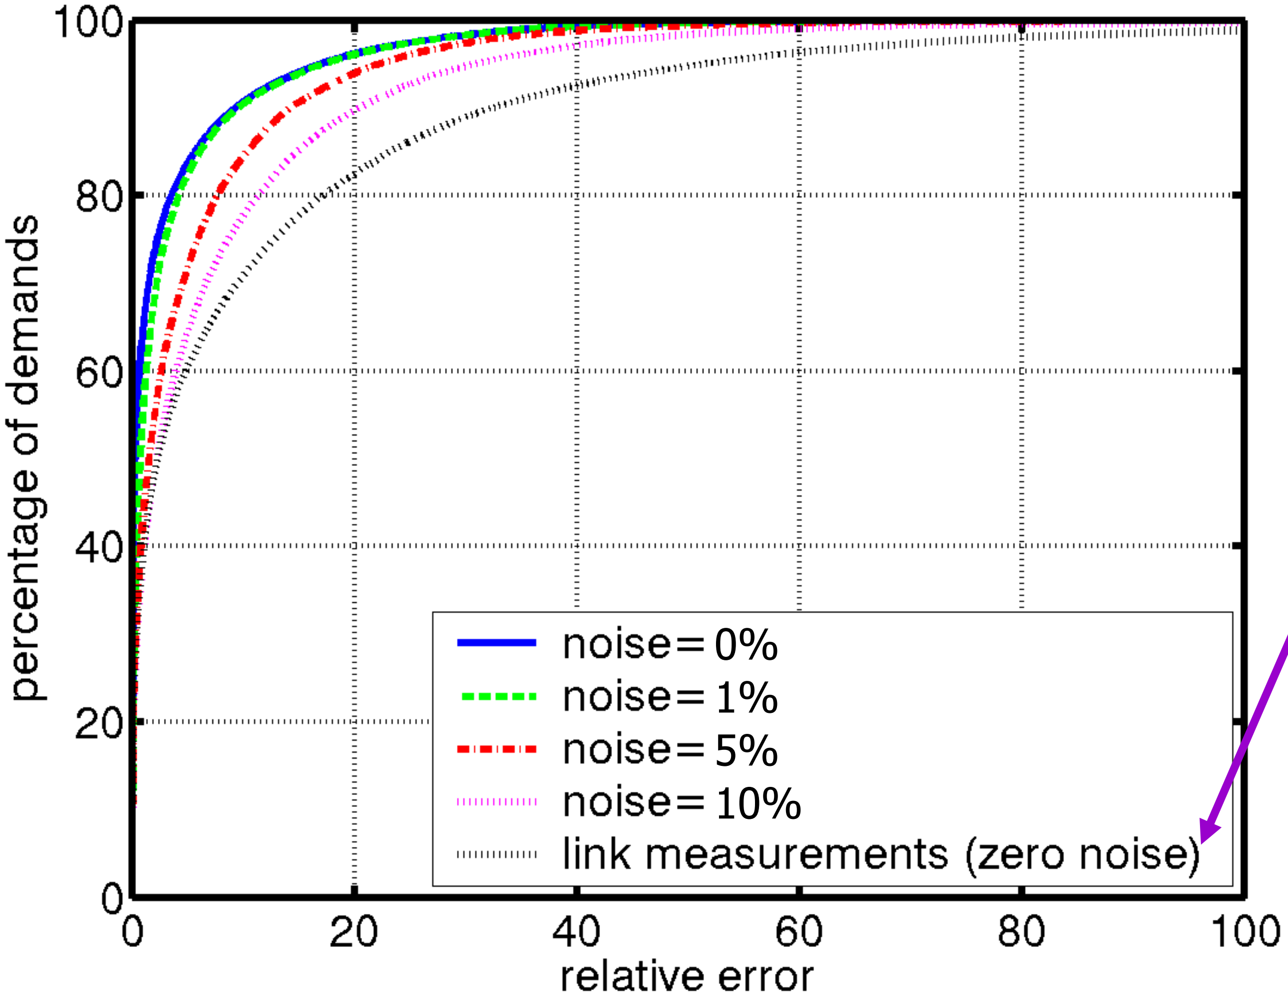

for reference
previous case

Tomo-gravity

Robustness (missing data)

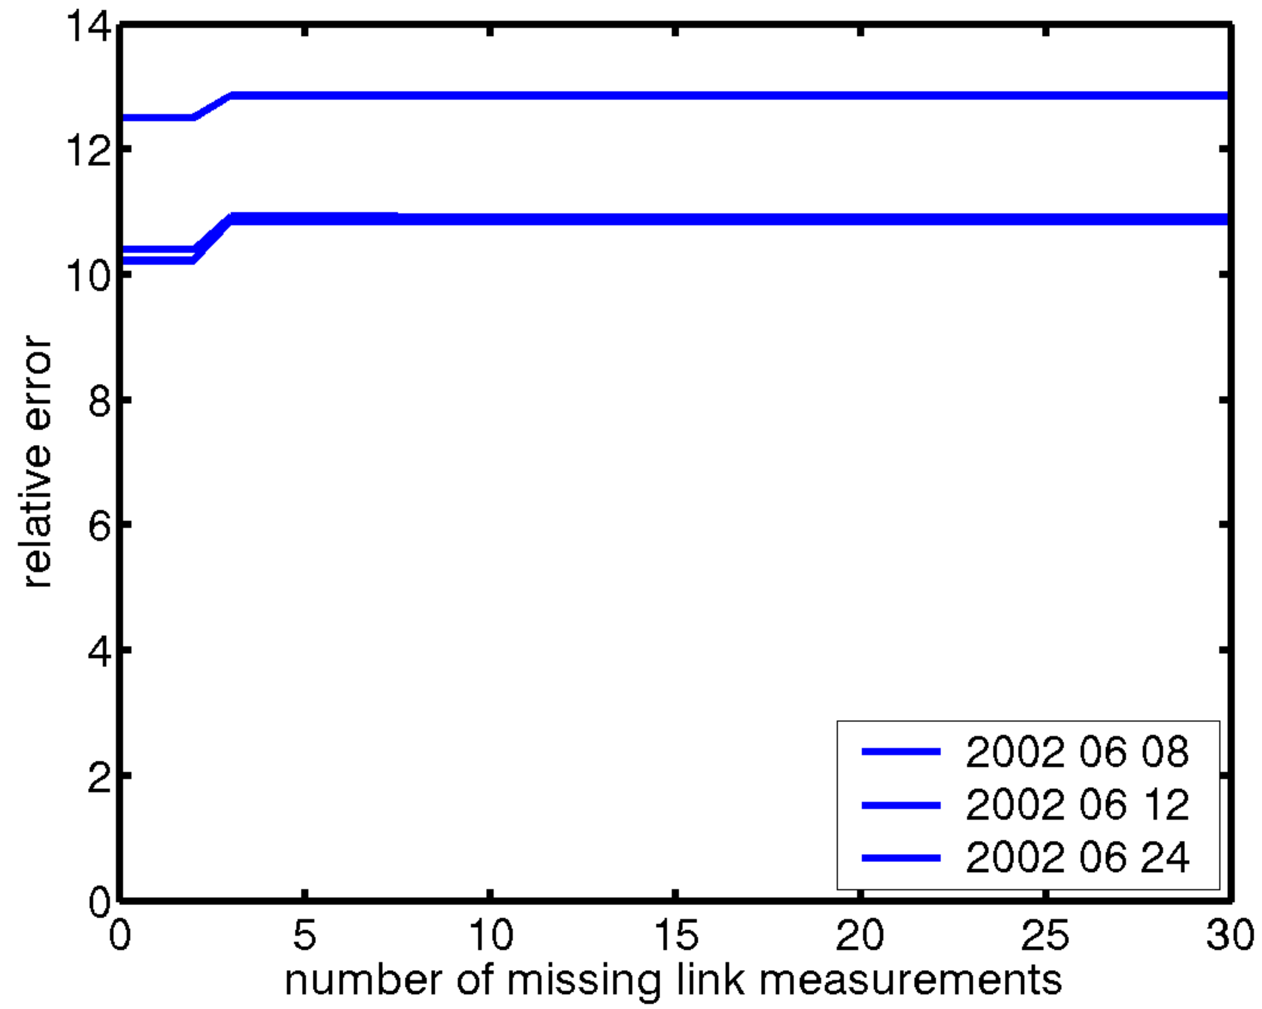

Tomo-gravity

Point-to-multipoint

We don't see whole Internet – What if an edge link fails?

Point-to-point traffic matrix isn't invariant

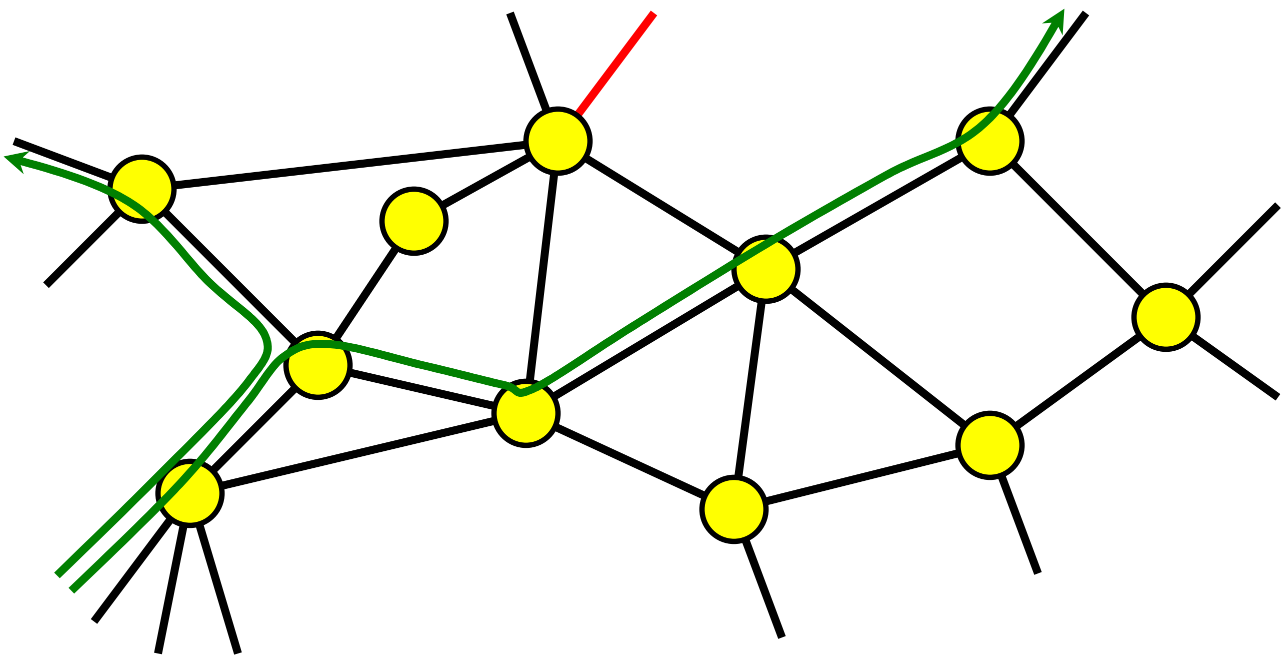

Point-to-multipoint

- Included in this approach
- Implicit in results above
- Explicit results worse
  - Ambiguity in demands in increased
  - More demands use exactly the same sets of routes
- use in applications is better

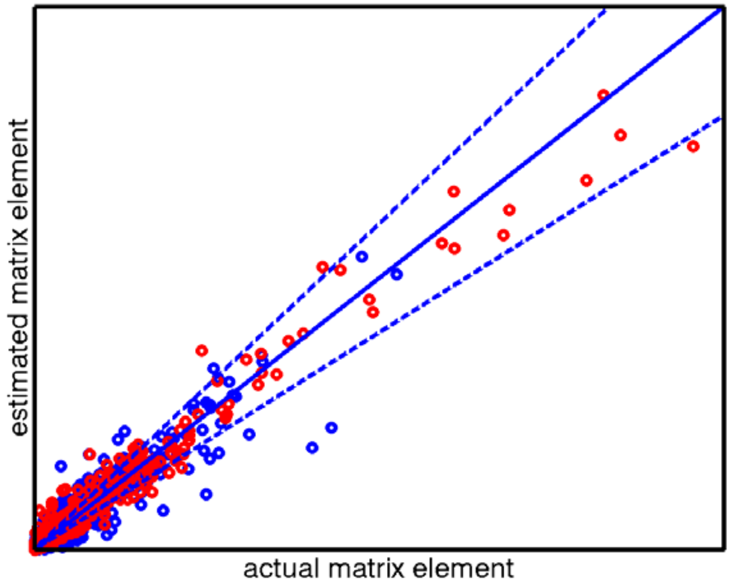

Link failure analysis