



cariden

the economics of network control

RIPE51
Amsterdam

Estimating the Traffic Matrix in IP Networks

October 10th, 2005

Thomas Telkamp
telkamp@cariden.com

The Traffic Matrix

- The amount of data transmitted between every pair of network nodes
- For what purpose?
 - Analysis and Evaluation of other network states than the current:
 - Capacity Planning
 - network changes
 - Resilience Analysis
 - network under failure conditions
 - Optimization
 - OSPF/IS-IS metric optimization/TE
- Also interested in external traffic (Inter-AS)

What do we have?

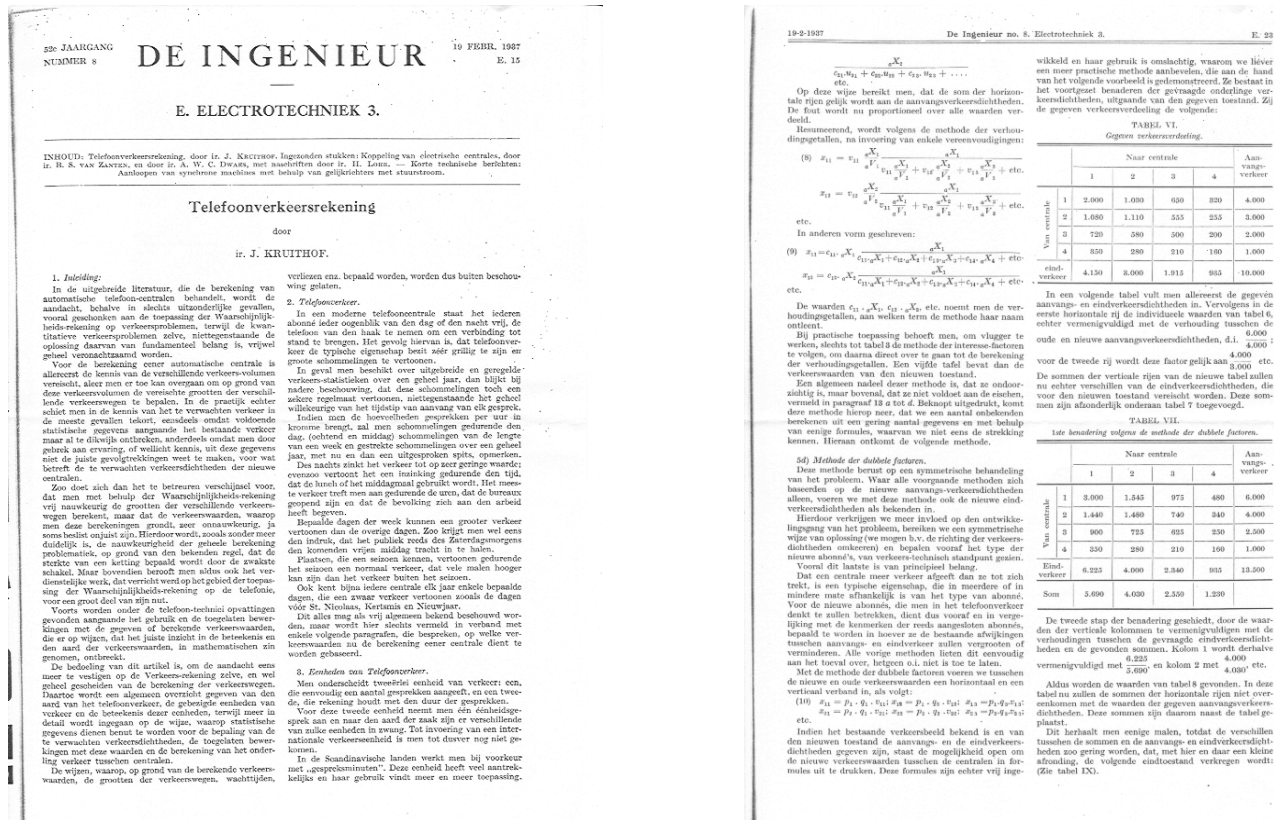
- NetFlow
 - v5 deployment is complex
 - newer versions (aggregation, v8) not complete
- DCU (Juniper)
 - only 16 classes
- BGP Policy Accounting
 - only 64 classes, BGP only
- MPLS
 - TE tunnels
 - requires full TE mesh
 - LDP counters
 - nice solution, only minor issues (see [4])

What do we want?

- Derive Traffic Matrix (TM) from easy to measure variables
 - No complex features to enable
- Link Utilization measurements
 - SNMP
 - easy to collect, e.g. MRTG
- Problem:
 - Estimate point-to-point demands from measured link loads*
- Network Tomography
 - Y. Vardi, 1996
 - Similar to: Seismology, MRI scan, etc.

Is this new?


- Not really...
- ir. J. Kruijthof: **Telefoonverkeersrekening, De Ingenieur, vol. 52, no. 8, feb. 1937 (!)**



Demand Estimation

- Underdetermined system:
 - N nodes in the network
 - $O(N)$ links utilizations (*known*)
 - $O(N^2)$ demands (*unknown*)
 - Must add additional assumptions (information)
- Many algorithms exist:
 - *Gravity model*
 - *Iterative Proportional Fitting (Kruithof's Projection)*
 - *Maximum Likelihood Estimation*
 - *Entropy maximization*
 - *Bayesian statistics (model prior knowledge)*
 - *Etc...!*
- Calculate the **most likely** Traffic Matrix

Example



y : link utilizations

A : routing matrix

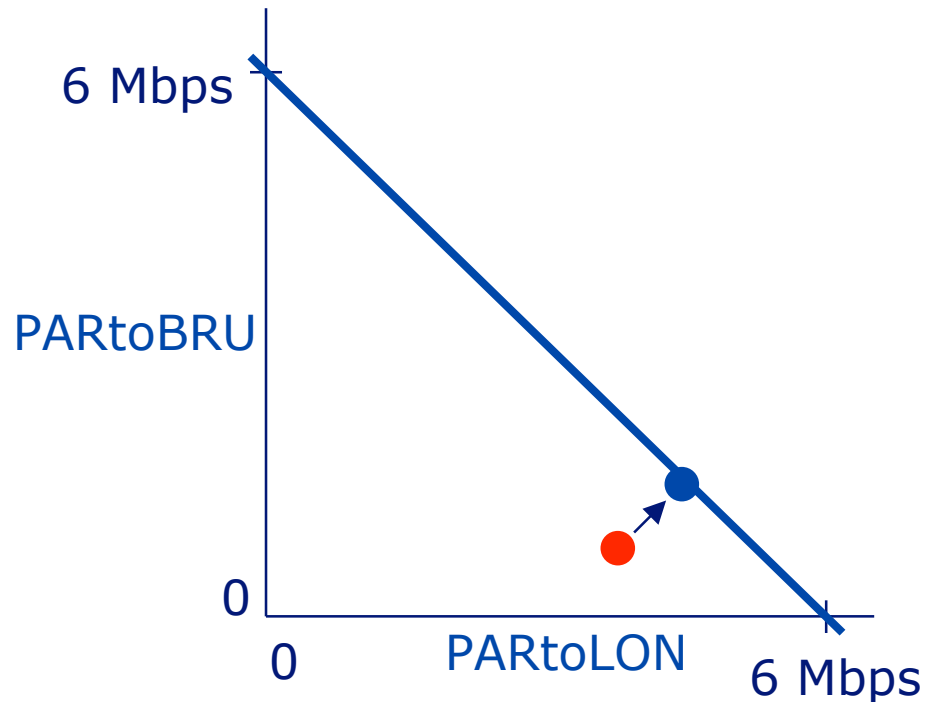
x : point-to-point demands

Solve: $y = Ax$

In this example: $6 = PARtoBRU + PARtoLON$

Example

$$\text{Solve: } \underline{y = Ax} \quad \rightarrow \quad \underline{6 = PARtoBRU + PARtoLON}$$



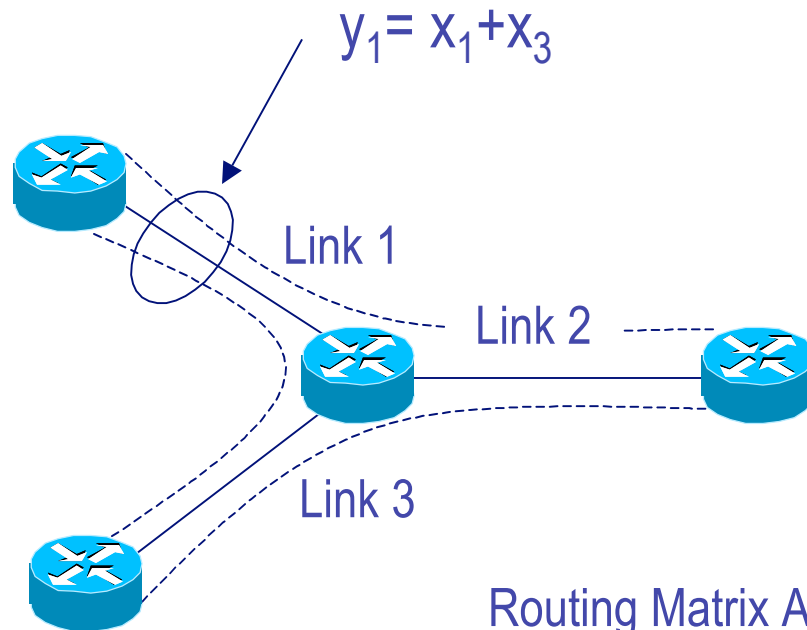
Additional information

E.g. Gravity Model (every source sends the same percentage as all other sources of its total traffic to a certain destination)

Example: Total traffic sourced at PAR is *50Mbps*. BRU sinks 2% of total traffic, LON sinks 8%:
PARtoBRU = 1 Mbps and
PARtoLON = 4 Mbps

Final Estimate: PARtoBRU = 1.5 Mbps and PARtoLON = 4.5 Mbps

General Formulation




- The total traffic on each link is the sum of all the source destination flows that route over that link
- Given Y and the routing matrix A solve for X

Routing Matrix A

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


Network Results: Estimated Demands

International Tier-1
IP Backbone



Estimated Link Utilizations!

International Tier-1
IP Backbone



Demand Estimation Results

- Individual demands
 - Inaccurate estimates...
- Estimated worst-case link utilizations
 - Accurate!
- Explanation
 - Multiple demands on the same path indistinguishable, but their sum is known
 - If these demands fail-over to the same alternative path, the resulting link utilizations will be correct


TM Estimation Deployment Case

- Large ISP network
 - about 80 Routers and 200 Circuits
 - 2550 TM entries
 - not all routers source/sink traffic (e.g. core)
- Known Traffic Matrix
 - Direct MPLS measurement
- Case-study will evaluate:
 - How does estimated TM compare to known TM?
 - How well does the estimated TM predict worst-case link utilizations?
 - How much can the estimated TM be improved by adding point-to-point measurements?
- TM estimation using Cariden MATE Software
 - Demand Deduction tool

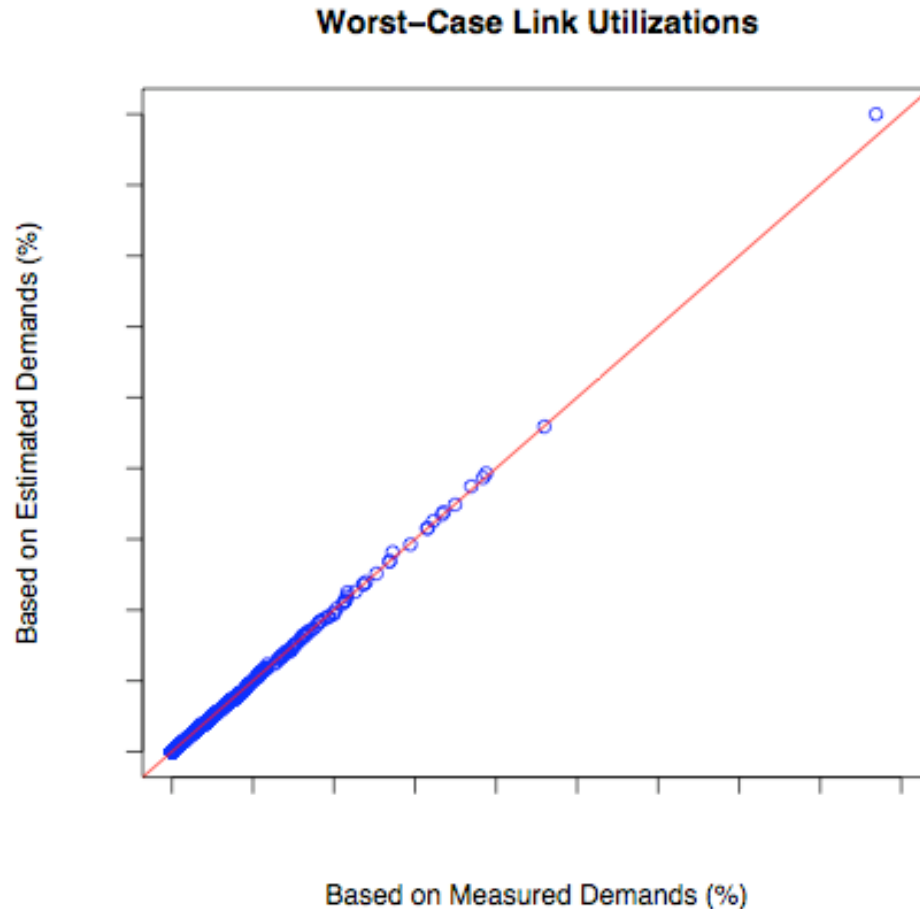
Procedure

- Start with current network and known TM
 - save as “PlanA” (with TM “Known”)
- IGP Simulation for non-failure
- Save Link Utilizations and Node In/Out traffic
- Estimate Traffic Matrix
 - New TM: “Estimated”
 - Save as “PlanB”
- Simulate IGP routing on both networks
 - single link failures
- Compare Results!

Estimated Demands




Worst-Case Link Utilization



- PlanA Traffic Matrix:
 - Measured
- PlanB Traffic Matrix:
 - Estimated
- IGP Simulation
 - Circuit failures
- Compare Worst-Case Link Utilizations (in %)

Add Measurements (1)

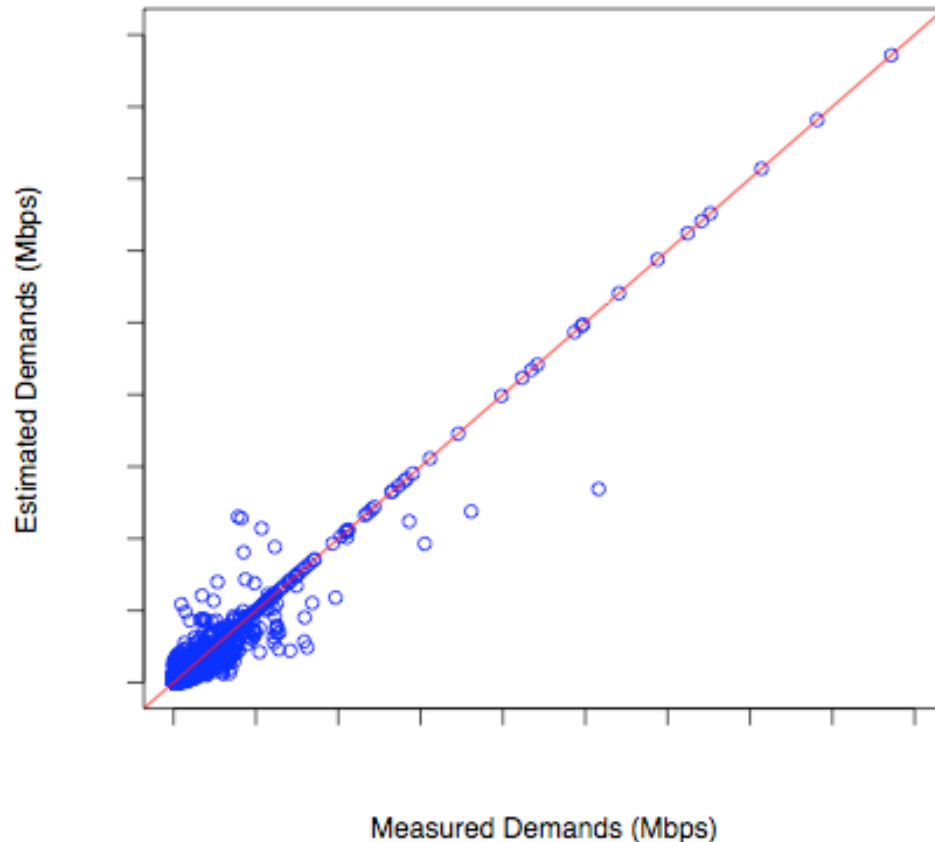
Demand Estimation: Measure top 100 Demands



- Select the top 100 demands from the estimated traffic matrix
- Setup Juniper DCU (or equivalent) to measure these demands
 - on 23 routers
 - 38% of total traffic
- Add measured data to the TM estimation process
- (FYI: 990 demands represent 90% of total traffic)

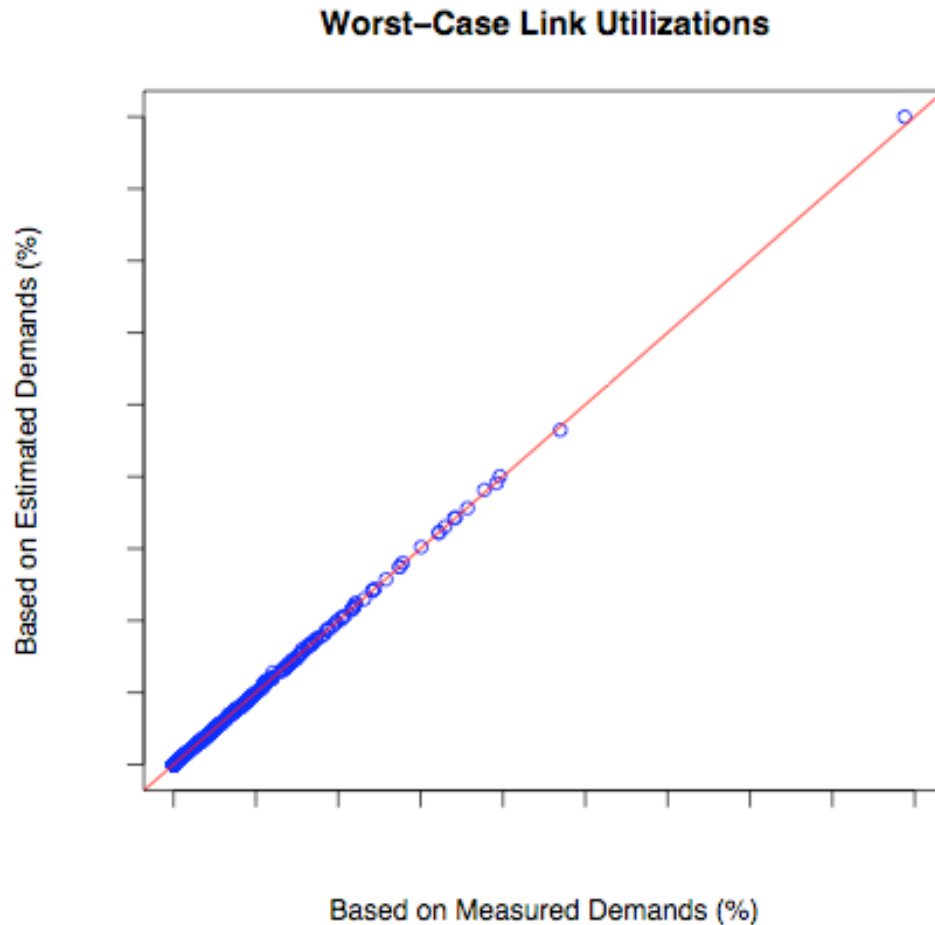
Add Measurements (2)

Demand Estimation: Measurements on 10 Routers



- Select the top 10 routers from the measured in/out traffic
- Select the top 16 demands on each of these routers (estimated)
- Setup Juniper DCU to measure these demands
 - 160 demands
 - 10 routers
 - 41% of total traffic
- Add measured data to the TM estimation process

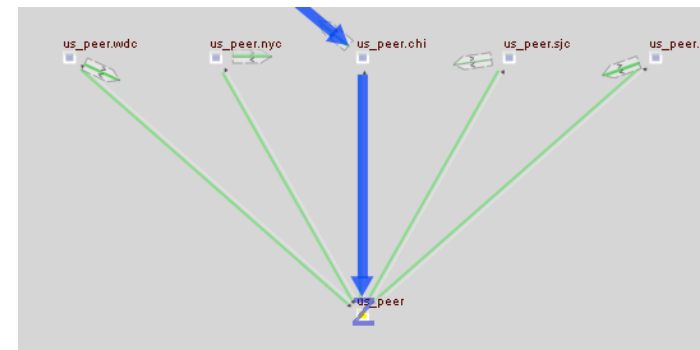
Worst-Case Link Utilization (2)



- Revisit the worst-case link utilizations, now with more accurate TM
- Similar results as before
- Hardly any room for improvement!


External/Inter-AS TM

- Traffic Matrix on a network will change due to core failures (closest-exit), or peering link failures
- Create *router-to-peer* TM
- Estimation procedure is similar
- Routing is different
 - policy restrictions:
e.g. no traffic from peer to peer
 - Virtual model of remote AS
- Estimation can make use of a known core TM



External Traffic Matrix

B. Claise, Cisco



From "Router to BGP AS", the router being the **AR** or **CR**
The external traffic matrix can influence the internal one

TM Estimation Summary

- Algorithms have been published
 - Implement yourself (e.g. IPF procedure)
 - Commercial tools are available
- Can be used in multiple scenarios:
 - Fully estimate Traffic Matrix
 - Combine with NetFlow/DCU/etc.
 - Measure large demands, estimate small ones
 - Estimate unknown demands in a network with partial MPLS mesh (LDP or RSVP)
 - Estimate Peering traffic when Core Traffic Matrix is known
- Also see AT&T work
 - E.g. Nanog29: *How to Compute Accurate Traffic Matrices for Your Network in Seconds* [2]

Questions?

Thomas Telkamp
telkamp@cariden.com

References

1. A. Gunnar, M. Johansson, and T. Telkamp, "Traffic Matrix Estimation on a Large IP Backbone - A Comparison on Real Data", *Internet Measurement Conference 2004*. Taormina, Italy, October 2004.
2. Yin Zhang, Matthew Roughan, Albert Greenberg, David Donoho, Nick Duffield, Carsten Lund, Quynh Nguyen, and David Donoho, "How to Compute Accurate Traffic Matrices for Your Network in Seconds", NANOG29, Chicago, October 2004.
3. AT&T Tomogravity page:
<http://public.research.att.com/viewProject.cfm?prjID=133/>
4. S. Schnitter, T-Systems; M. Horneffer, T-Com. "Traffic Matrices for MPLS Networks with LDP Traffic Statistics." Proc. Networks 2004, VDE-Verlag 2004.
5. Y. Vardi. "Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data." J.of the American Statistical Association, pages 365–377, 1996.