

On the Number of Distributed Measurement Points for Network Tomography

Joseph D. Horton U of New Brunswick

Alex López-Ortiz U of Waterloo

Motivation

- Tomography in real life:
 - NIMI
 - MINC
 - Matrix.net
 - SPAND
 - and many others...
- Why does it work?
- When does it work?

Previous Work

- Heuristics for location of the beacon set [Jamin et al. Infocomm 2000]
- Marginal utility of beacons [Barford, Bestavros, Byers, Crovella IMW 2001]
- Tomography in multicast networks of arbitrary topology [Bu, Duffield, Presti, Towsley Sigmetrics 2002]

Definitions and assumptions

- Strategically located computers called **beacon set**
- Ping/traceroute used to discover network topology
- BGP-like routing
- Goal: measure link state of every (advertised) link in the network

Theory

Difficult to find minimum set of beacons
in general

Theorem NP-hard to find minimum set of
beacons on arbitrary network topology
(reduction to set cover)

Theory

Even good approximations are hard

Theorem Hard to approximate optimum beacon set to factor of $\log n$ ($n = \#$ nodes).

Theory

Some networks require an impractically large number of beacons

Theorem There exist a network that requires $n/3$ beacons to recover the topology.


Practice

=> We could lose our ability to do tomography under **chaotic network growth**.

- However...
we know tomography works in the Internet today

Better model

- Focus only on nodes with alternative routing paths



- Forks in the road: Higher arity nodes

Theorem Higher arity nodes form a beacon set.

Higher Arity Nodes (HAN)

- Internet is *mostly* a tree => the number of HAN nodes is low
- In fact, it can be estimated from routing tables
- Somewhere around 10K
- Proposed cap at 32K [RFC draft, Savola 2003]
- Beacons set of about 10K within range of a commercial organization (e.g. Akamai, Google)

HANs on the RON

- Said beacon set covers every fork in the network, induces a RON
- => enables all paths routing, via forwarding
- ⇒ A commercial organization could provide all paths routing using this RON (Sockeye, Internap, RouteScience)

Conclusions

- Tomography does not work in general
- Works (mostly) in current internet, can be made exact using 10K nodes.
- It could stop working, if network growth were to be chaotic
- Higher arity nodes form a beacon set
- HANs form an all-paths routing RON via packet forwarding