Geographic Properties of Internet Routing

Lakshminarayanan Subramanian
University of California, Berkeley

Venkat N.Padmanabhan
Microsoft Research, Redmond

Randy H.Katz
University of California, Berkeley

Properties we study

- **Circuitousness of Geographic Paths**
  - **Can help to flag anomalous paths?**
- **Multi-ISP routing properties**
  - **Hot-Potato/ Cold-Potato Routing?**
  - **Can we find sub-optimal peering among ISPs?**
- **Geographic Fault Tolerance**
  - **Can ISPs take geographic node failures?**

Outline

- **What is a Geographic Path?**
- **GeoTrack and its properties**
- **Experimental Methodology**
- **Circuitousness of Routes**
- **Routing across multiple ISPs**
- **Geographic Fault Tolerance**
- **Conclusions**

What is a Geographic Path?

- Consider the *traceroute* path to an end-host. The *geographic path* is the ordered list of unique geographic locations of the routers in the path.
- Geographic path does not cover the entire $L2$ path.
- Example:

  BERK–BERK-7507.POS.calren2.net Berkeley+CA
  acr1-serial2-3-0-0.SanFranciscosfd.cw.net San+Francisco+CA
  acr1-loopback.Seattlesel.cw.net Seattle+WA
  208.172.83.118 unknown
  sl-bb10-sea-7-0.sprintlink.net Seattle+WA
  sl-bb21-pen-5-2.sprintlink.net Pennsauken+NJ
  at-bb4-nyc-4-0-0.appliedtheory.net New+York+NY
  at-bb1-nyc-6-0-0-OC3.appliedtheory.net New+York+NY

- Geographic Path: Berkeley, San Francisco, Seattle, Pennsauken, New York

GeoTrack

- Recognize router's geographic location using *codes* that may be embedded in its DNS name.
- Example:
  - Code *pen* in *sl-bb21-pen-5-2.sprintlink.net* refers to Pennsauken+NJ.
- Coverage of the Tool
  - US + 26 countries in Europe.
  - More than 2000 location codes for US and Europe.
- Recognizability of routers
  - Recognize 70% of *.net* router labels(7842 out of 11296).
  - Among 13 major ISPs, recognizability is 87%.

Recognizability of Geographic Paths

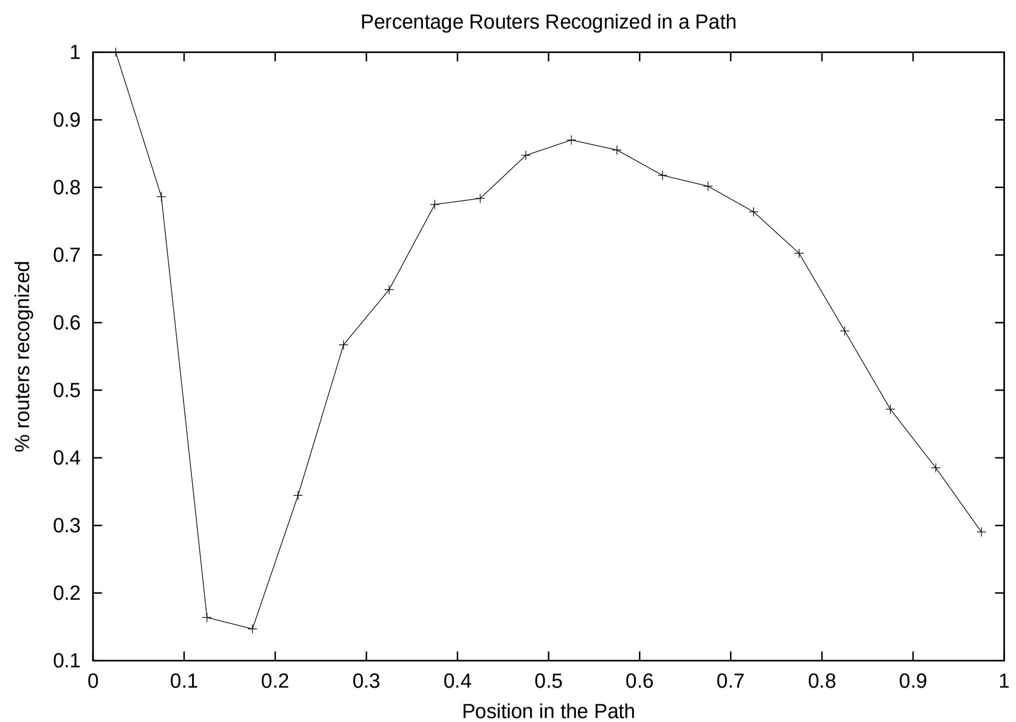

- Can recognize most of the routers when in transit between major ISPs. This forms a large portion of the Geographic Path.

Experimental Methodology

- Measurement Test bed
  - 20 geographically dispersed probe points(17 in US, 3 in Europe)
- Variety of Destination End-hosts
  - UnivHosts: 265 universities in the US.
  - LibWeb: 1205 public libraries in the US.
  - EuroWeb: 1092 web servers in Europe.
  - TVHosts: 3100 clients of an online TV program guide.
- Approximately 85000 traceroutes in total.
- Vern Paxson's traceroute dataset in 1995.
- ISP topologies from CAIDA's Mapnet tool.

Metrics

- Circuitousness of Paths
  - Compute *linearized distance* of a geographic path as the sum of hop-hop geographic distances in the path.
  - *Distance Ratio* = linearized distance /end2end geographic distance.
- Multi-ISP routing
  - Compute *fraction of path in ISP* as the fraction of the linearized distance of the path within the ISP.
- Geographic fault tolerance
  - Number of components in the graph.

Circuitousness: Effect of Source Location

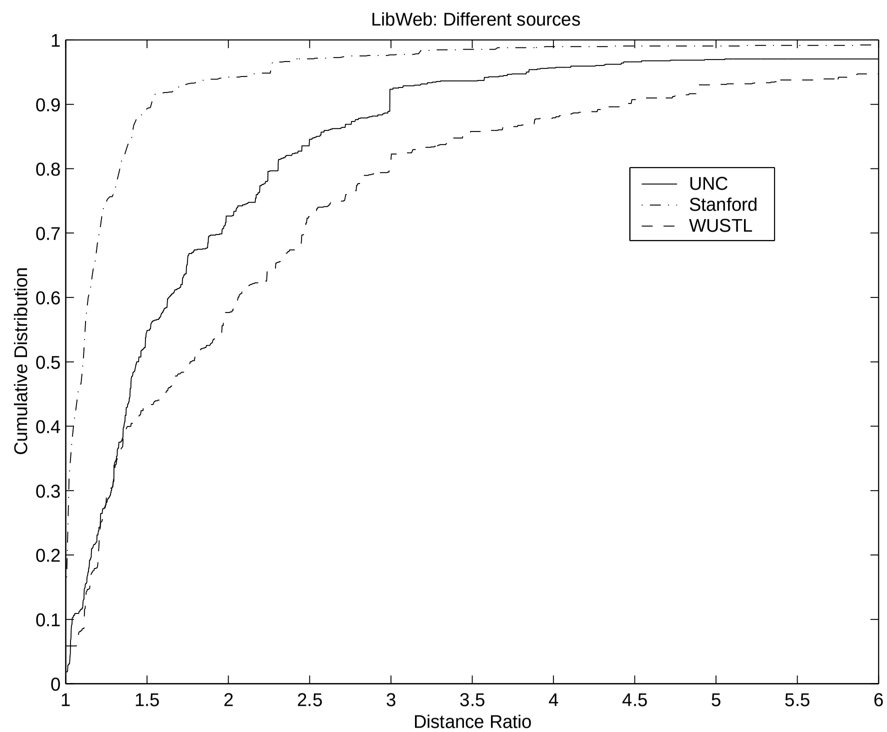

- **If source is close to *well connected* geographic location, distance ratio is expected to be smaller.**

Routing in US vs Routing in Europe

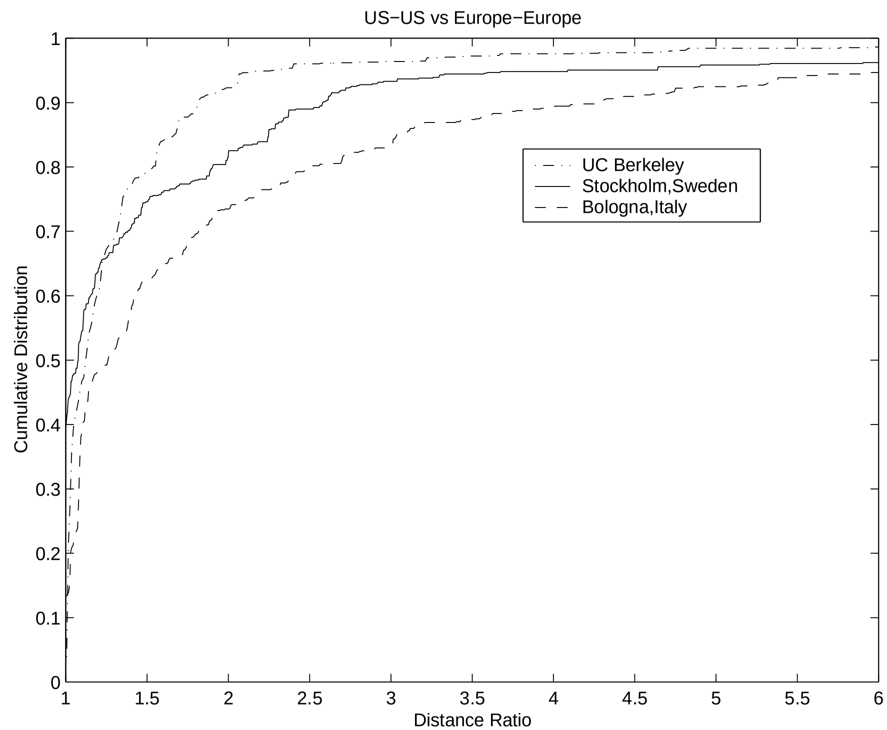

- There are many small (national) networks in Europe which makes routing circuitous. So many small ISPs are traversed in a path.

Distance ratio vs Geographic Distance

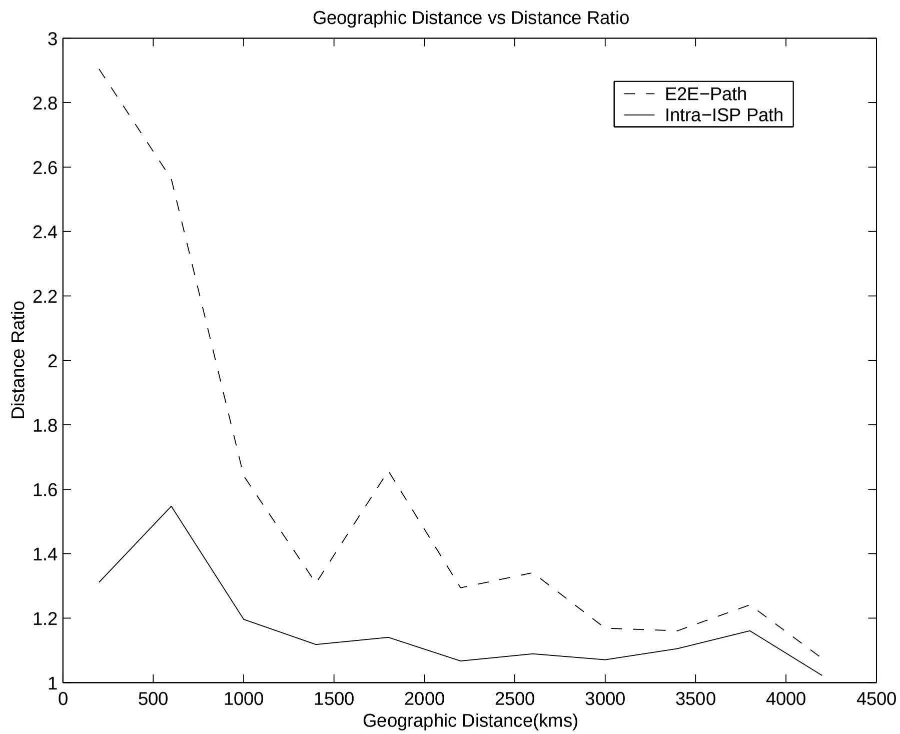

- **End2End paths are much more circuitous than intra-ISP paths.**

Multi-ISP routing

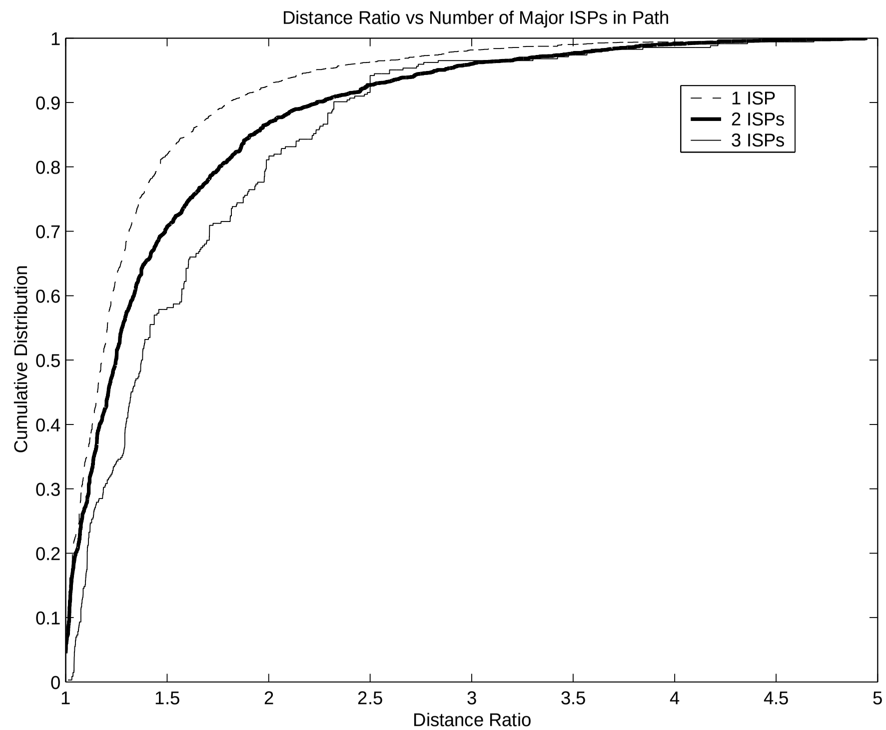

- **The distance ratio seems to increase as the number of *major* ISPs in a path increase. Does this suggest sub-optimal peering?**

Hot-Potato routing

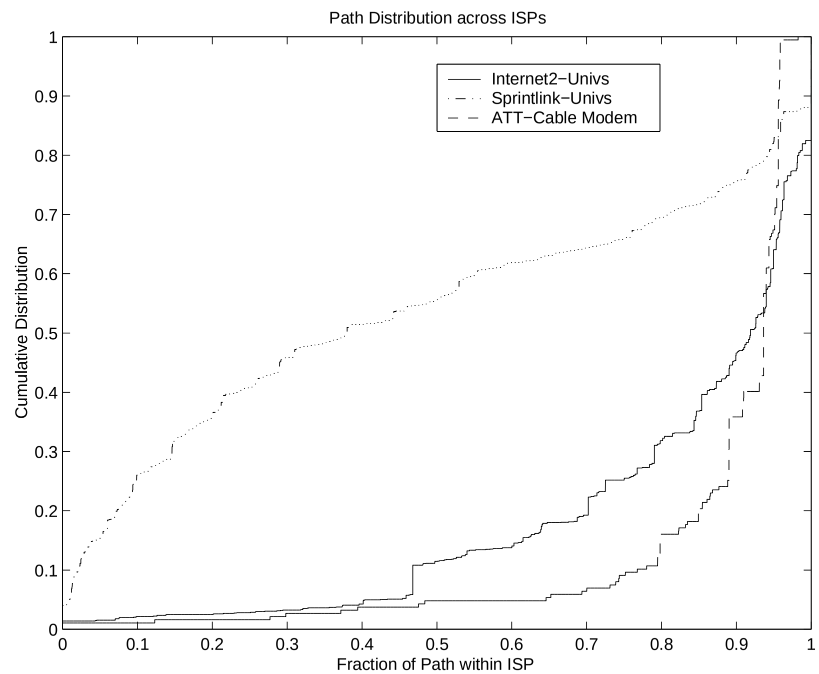

- Some major ISPs(Sprintlink) seem to perform hot-potato routing while some others(Internet2) performs cold-potato routing.
- Peering relationships may have a higher influence on routing than hot-potato/cold-potato policies.

Summary: Circuitousness, Multi-ISP routing

- Circuitousness depends on:
  - Source and Destination location (proximity to networking hubs).
  - Network structure(small networks vs Large ISPs).
- Intra-ISP paths are less circuitousness Inter-ISP paths.
- Circuitous paths normally traverse multiple major ISPs.
- Geography helps in determining:
  - Presence of Hot-potato/Cold-Potato routing in ISPs
  - Can potentially detect sub-optimal peering?

Geographic Fault Tolerance

- Consider an ISP's topology with N geographic nodes. i.e. collapse all routers in one geographic location into one node.
- A connected graph has $N \times (N+1)/2$ communicating pairs of geographic nodes (including intra-node communication).
- When node failure(s) occurs:
  - Graph may be sub-divided into components.
  - A component with k nodes has $k \times (k+1)/2$ communicating pairs.
  - Determine number of communicating pairs in the new topology.
- Worst-case scenario: Nodes with high degree in the graph fail.

ISP: Degree Distributions

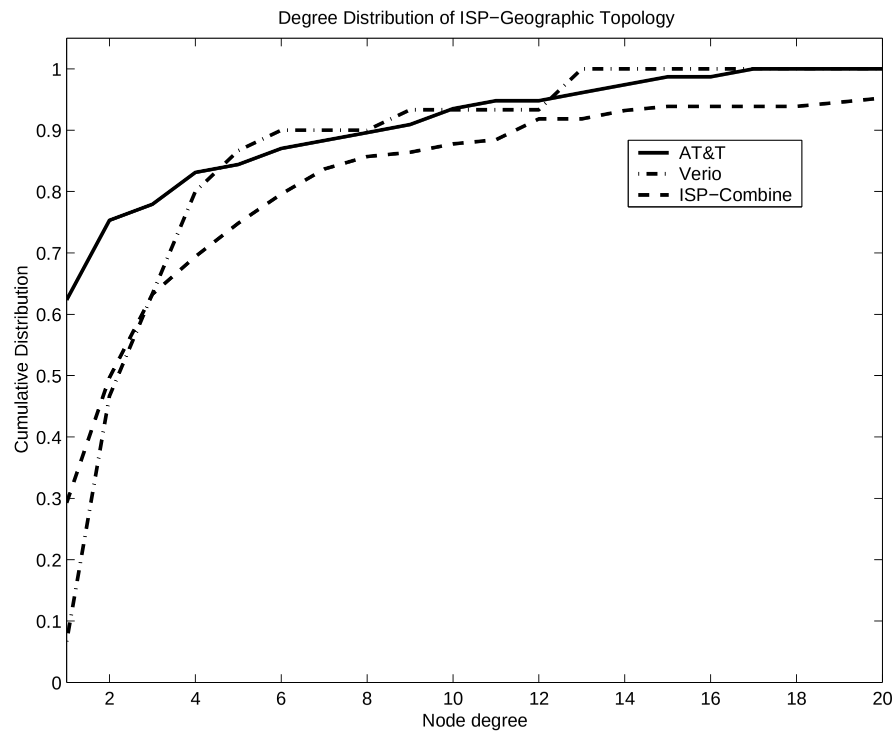

- Many major ISPs have some nodes of very high degree making them vulnerable to failures. (New York, Chicago)
- Even the combined topology of 9 ISPs has a skewed degree distribution.

Worst case Failure Scenario

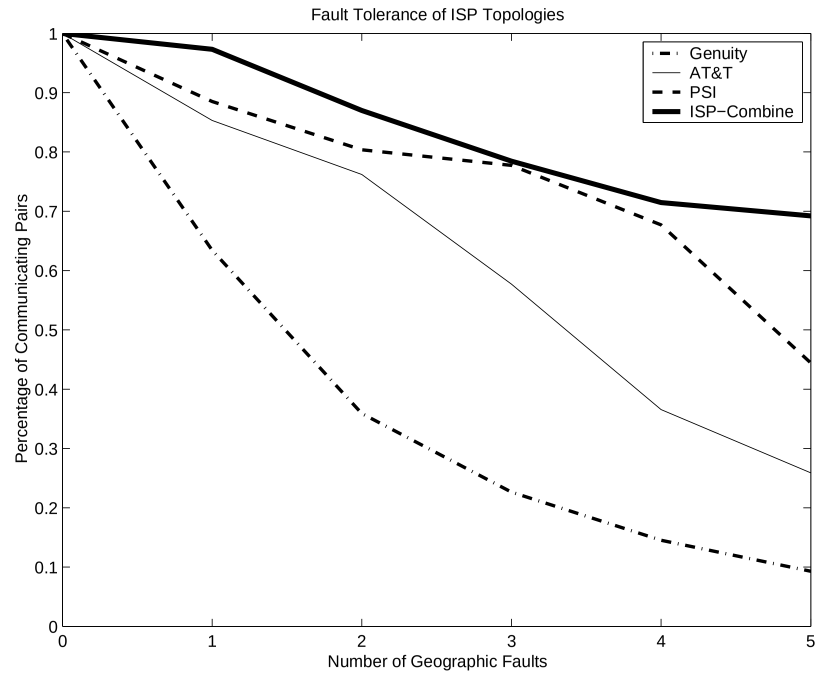

- PSINet has the best fault tolerance properties. Degree distribution of PSINet very balanced(very few leaves).
- Genuity(BBNPlanet) had very low fault tolerance.(60% of the nodes are leaves)

Summary: Geographic Fault Tolerance

- **Many ISPs do not have a good fault tolerant structure. This is true even for the combined topology of 9 major ISPs.**
- **ISPs based on *Ring* topology normally have better geographic fault tolerance(tolerance to 1 failure). KPNQwest is one example.**
- **In some cases it may not be economically viable to build fault tolerant topologies.**