



Geo-Location of PoPs

Noa Zilberman & Yuval Shavitt

Tel Aviv University

February-2010

[Agenda

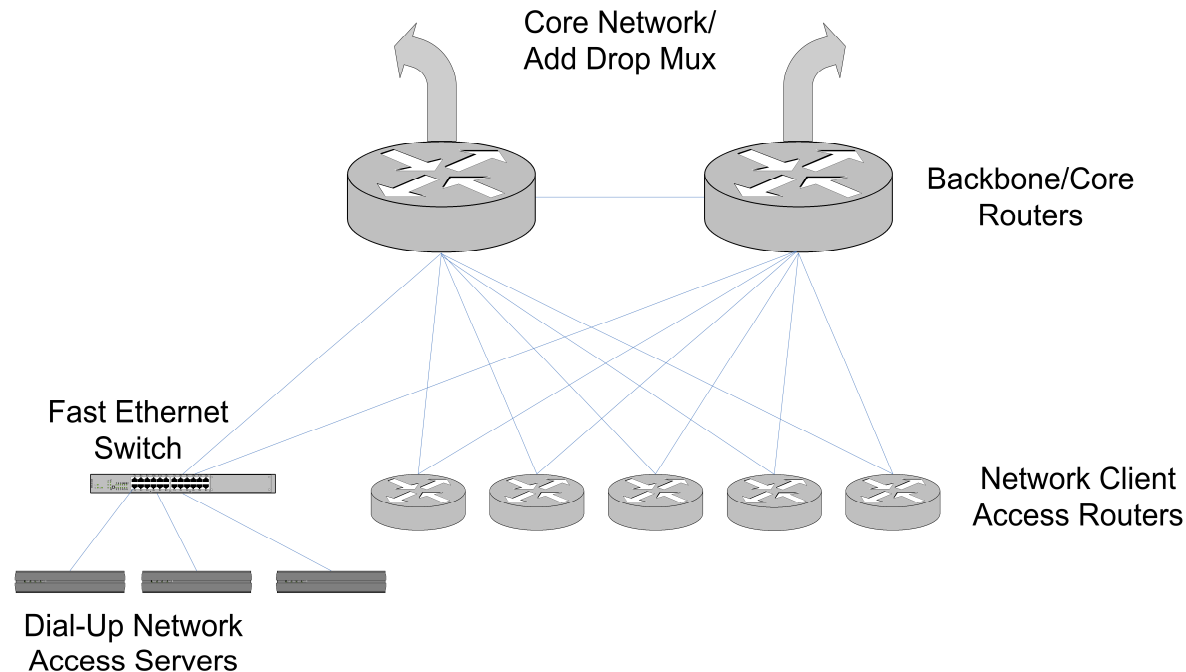


- Background
- PoP Discovery
- PoP Geolocation
- Evaluating Geolocation Databases
- AS Connectivity on PoP Level

[Background





- PoP – Point of Presence - a concentration of routers and other networking devices in a campus from which Internet connectivity is offered to the region.
- DIMES worked so far on either IP or AS level.



[PoP Discovery



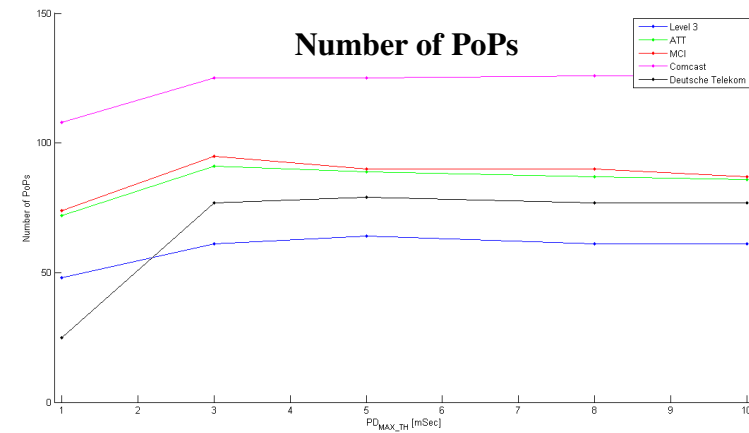
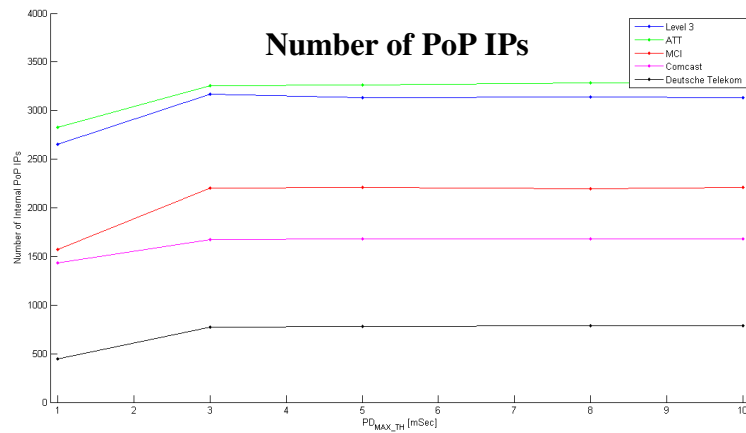
- Use *Link Delay* and *Network Motifs* to identify a PoP:
 - *An earlier work by D. Feldman & Y. Shavitt*
 - Look for edges with small link delay
 - Indicates nodes proximity.
 - Require a minimal number of measurements per link, for delay accuracy.
 - Identify bi-partite motifs in the graph
 - Classify to Parent-Child groups
 - Localization and unification to PoPs




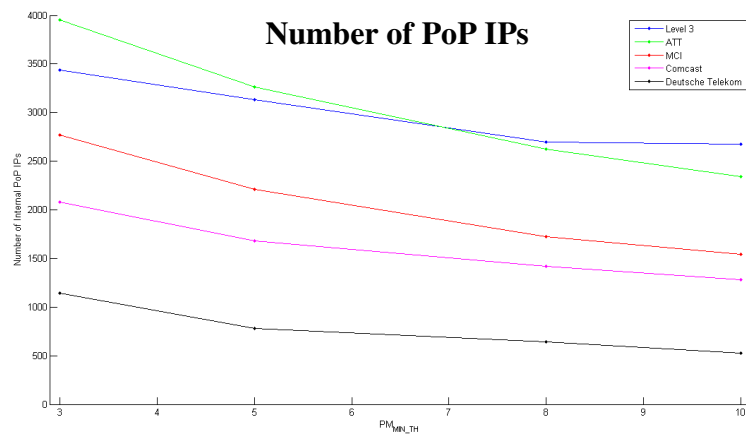
PoP Discovery



Sensitivity to delay threshold:



Sensitivity to number of measurements threshold:



[PoP Discovery



- Running on bi-weekly basis
 - Increased number of discovered PoPs compared to 1 week period.
 - More sensitive to changes than 4 weeks period.
- Using Traceroute measurements
 - 30M-40M measurements per week.
 - 5.5M-6.5M distinct edges discovered.
 - ~1000 agents in over 200 ASes are used for the measurements.
 - 2.5M IP addresses in over 26,000 ASes are being targeted.
 - Using Median algorithm to estimate distance between nodes.

[PoP Discovery



- Discovered PoPs
 - ~4400 discovered PoPs.
 - Over 50K IPs within discovered PoPs.
- Discovered mostly large PoPs and not access PoPs.
- Enhancements
 - Targeting iPlanes's PoP's IP addresses – increased the number of discovered PoPs by less than 20%.
 - Targeted measurements to specific AS doubled the number of discovered PoPs in small ASes.
 - Had some effect in large PoPs but not to that extent.

[PoP Discovery



- Limitation: number of measurements
 - The number of discovered PoPs directly relates to the number of discovered edges
 - DIMES new Agent will more than double the amount of measurements
 - Beta version available this month!
 - We are interested to use traceroute measurements with delay information from other databases to improve PoP discovery.

- We'll be happy to discuss in detail, but lets move to GeoLocation...

[PoP GeoLocation



- We strongly believe that if we identify IPs as belonging to the same PoP - they are in the same geographic proximity.
- Use location information from several geolocation databases to determine PoP's location.
- Location is selected by majority vote.
 - Majority vote uses the location of all IPs within the PoP taken from all geolocation databases.
 - A range of error is given for each PoP location.
 - No more than 100km radius.
 - The location is given as Latitude, Longitude.
 - With some refinements....

[PoP GeoLocation

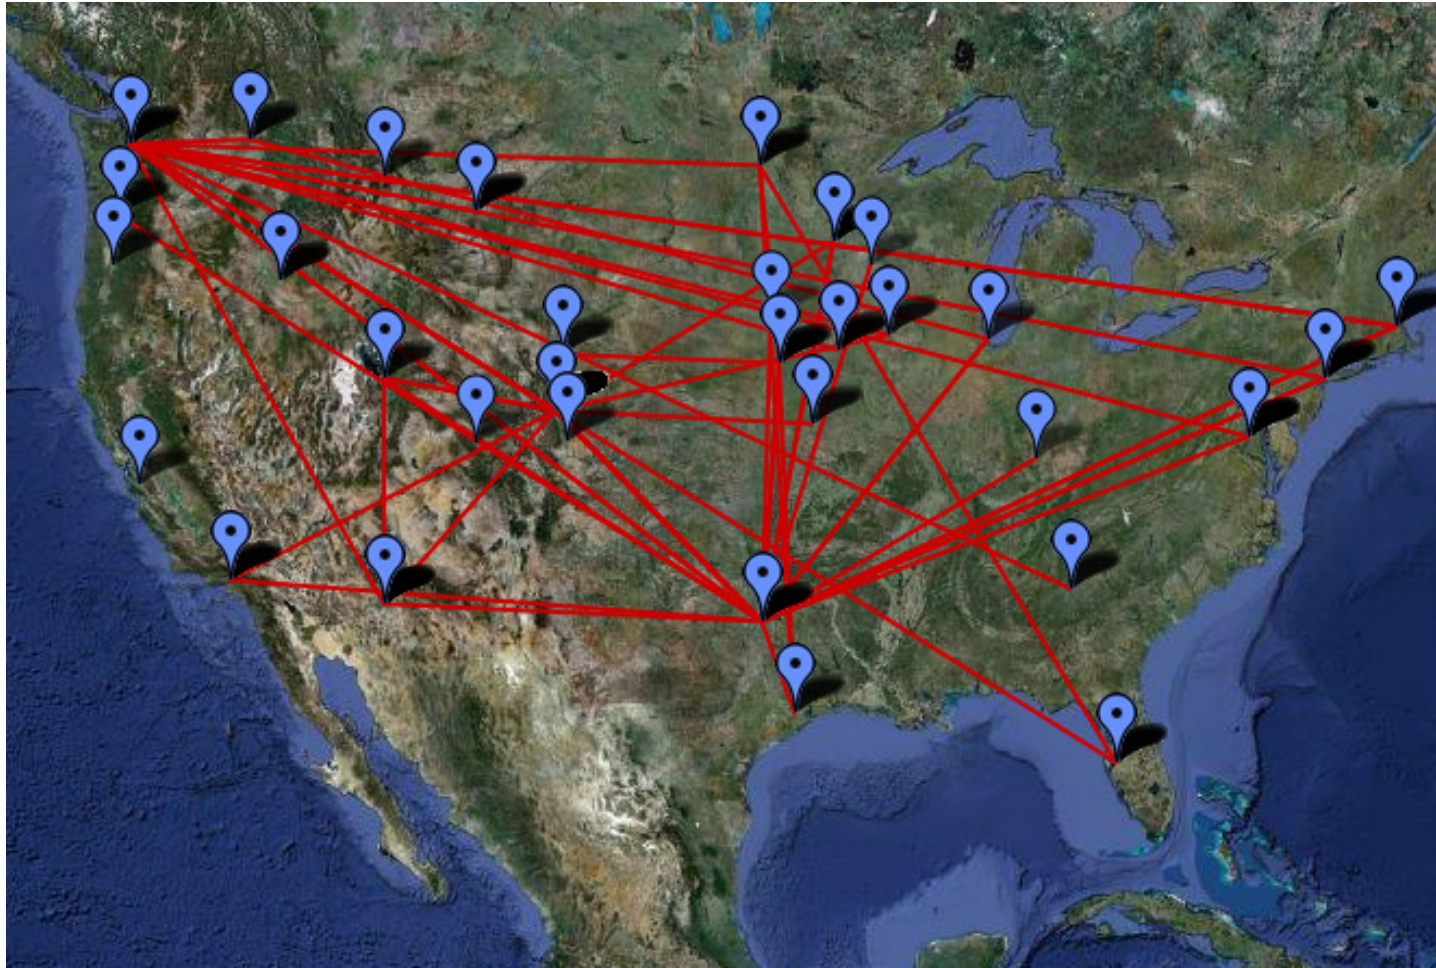


- Used commercial GeoLocation Databases:
 - MaxMind GeoIP
 - IPLigence
 - HostIP.info
 - IP2Location
- Quova was not used, though it is supposed to be more accurate
 - Budget limitations
- DNS was used for limited testing

World PoPs Map



[Qwest US PoPs Map



[PoP GeoLocation - Validation



- Compared generated PoP maps to published ISP PoP maps:
 - Sprint, Qwest, Global crossing, British Telecom, ATT etc.
 - PoPs are correctly located

- Compared against Universities locations
 - Selected 50 PoPs belonging to universities world-wide
 - 49 universities were correctly located by the algorithm
 - University of Pisa was located in Rome
 - Wrong information in MaxMind and Ipligence, HostIP.info was right.

PoP GeoLocation - Results



- 82% of the PoPs have majority vote considering all the IPs in the PoP.
- 12% more have majority vote only when considering nodes with location information.
 - Geolocation databases sometimes lack information on some IP addresses.
- 68% of PoPs are located with 1km range of convergence.
- For only 28% of the PoPs there is over 90% agreement between all location services.
- We fail to locate 5% of PoPs with high accuracy.

Evaluating GeoLocation databases



Missing Location Information

- MaxMind:
 - 12% of IPs
 - 10% of PoPs
 - Informed us that the quality information is on end-user and not router-IP.
- IPLigence:
 - 6.5% of IPs
 - 1% of PoPs
- HostIP.info:
 - 28% of IPs
 - ~33% of PoPs
- IP2Location:
 - 4.2% of IPs
 - 0% of PoPs

[Evaluating GeoLocation databases



Agreement within the same database

- Does nodes within the same PoP have the same location?
 - MaxMind: 72%
 - IPligence: 86%
 - HostIP.info: 77%
 - IP2Location: 74%
- In some cases, the location variance is negligible
 - i.e. considering larger PoP range of convergence can get a higher level of agreement

[Are GeoLocation DB truthful?



Qwest as an example

- 70 PoPs were discovered by the algorithm
- MaxMind assigned the PoPs to 55 different locations
- HostIP.Info assigned the PoPs to 46 different locations
- IP2Location assigned the PoPs to 35 different locations
- IPLigence located the PoPs in only one distinct location;
 - All the PoPs were placed in Denver, where Qwest HQ are located.

- MaxMind had the same problem as IPLigence in their May-2009 DB, but it was fixed in July-2009 DB.

Can GeoLocation DB be trusted?



- Global Crossing
 - A selected PoP, includes 4 IPs, all databases had 100% similarity
 - IP2Location located near Washington DC
 - IPelligence located in Pheonix
 - Distance is ~2500 mile from Washington
 - MaxMind located near Chicago
 - Distance is ~720 mile from Washington
- China Telecom
 - A selected PoP, includes 23 IPs, all databases had over 95% similarity
 - IP2Location located in Beijing
 - IPelligence located in Harbin
 - Distance is ~750 mile from Beijing
 - MaxMind located in Putian
 - Distance is ~1400 mile from Beijing

Keeping Track of DB updates

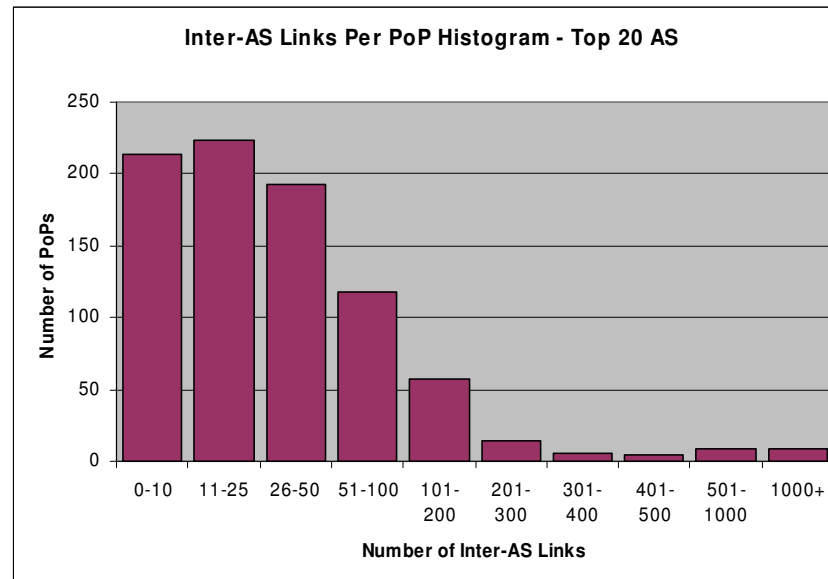


- Databases can significantly change between updates
- IPLigence as an example
 - ~0.6% of the entries changed between consecutive months (Nov/Dec 2009)
 - ~9.5% of the entries changed over 8 months period (April/Dec 2009)
- Other databases behave similarly
 - We have gaps in past databases, so it's hard to compare

AS Connectivity on PoP Level



- PoP level maps can also be used for the analysis of AS-level connectivity.
- Very high connectivity of PoPs within Top-20 measured AS:
 - Median of 22 links per PoP
 - A link is defined as a distinct connection between 2 different ASes
 - Multiple connections between two PoPs are counted only once



AS Connectivity on PoP Level



- Connectivity pairs between Top-10 and Top-20 measured ASes:
 - Average of 35 links between Top-10 AS
 - Median of 26 links between Top-20 AS
 - No case of a single-connection between Top-10 AS
 - Highest connected groups:
 - Comcast-GLBX, Comcast-MCI, Comcast-QWEST, ATT-GLBX, ATT-MCI

