CAIDA's Macroscopic Topology Project

Bradley Huffaker, Daniel Plummer, David Moore, and k claffy
CAIDA, SDSC, UC San Diego {**bradley,djp,david,kc**}**@caida.org**

SAINT 2002, The 2002 Symposium on Applications and the Internet, January 2002

overview

- topology project
- advantages and limitations
- visualization techniques
- analysis results
- conclusions
- future work

topology project

- monitors
  data collection (probe) machines
- skitter
  topology and RTT collection process
- destination lists
  list of IP address probed by a set of monitors
- support tools
  iffinder, NetGeo, skdesttest, dnsstat

topology project : monitors

topology project : skitter tool

- methodology
  - incremental ICMP echo requests
  - collects forward topology from source to destination
  - collects RTT (destination only)
- trace
  single forward IP path & RTT between a monitor and destination
- cycle
  a single run through the destination list

methodology : advantages and limitations

- **advantages**
  - IP topology vs BGP topology
    skitter [IP-level] probed topology yields far richer
    connectivity data than BGP announced [AS-level] topology
  - largest, longest running macroscopic topology project that
    combines topology and RTT collection
- **limitations**
  - Internet scale
    topology much more vast than can be realistically captured
  - path stability
    IP path can change, even while being probed
  - interfaces vs routers
    differences between interfaces and actual routers

topology project : destination list

- DNS clients
  - **goal:** monitor RTT (Round Trip Time) and connectivity
    for clients who send requests to DNS root servers
  - **method:** collect IP addresses sending requests to DNS roots
    and then select one IP address per routable prefix
- IPv4 list
  - **goal:** representative coverage of IPv4 address space
  - **method:** one IP address per /24 in routable prefixes,
    currently 54% of prefixes and 5.7% of /24s
- small
  - **goal:** provide higher granularity sample of topology and RTTs
  - **method:** fewer destinations (< 1000)
- web
  - **goal:** enable long term studies of large proportion of web
  - **method:** oldest list, mostly web (some dns) servers
    also allows comparison with current enviroment

topology project : support tools

- iffinder
  finds which IP interfaces belong to the same router
- NetGeo
  database that uses whois records to map IPs and ASes
  to geographic locations
- skdesttest
  helps construct destination lists: finds a set of IP addresses
  (from a given set of prefixes) that respond to ICMP
- dnsstat
  used in DNS study to passively collect IP addresses
  of DNS root clients

visualization techniques

- AS graph
  macroscopic visualization of complete AS topology
- walrus
  hyperbolic space graph visualization tool
- dispersion graph
  depict dispersion of paths from source to set of destinations
- PlotPath
  visualizes paths from source to one or few destinations

visualization : AS Graph

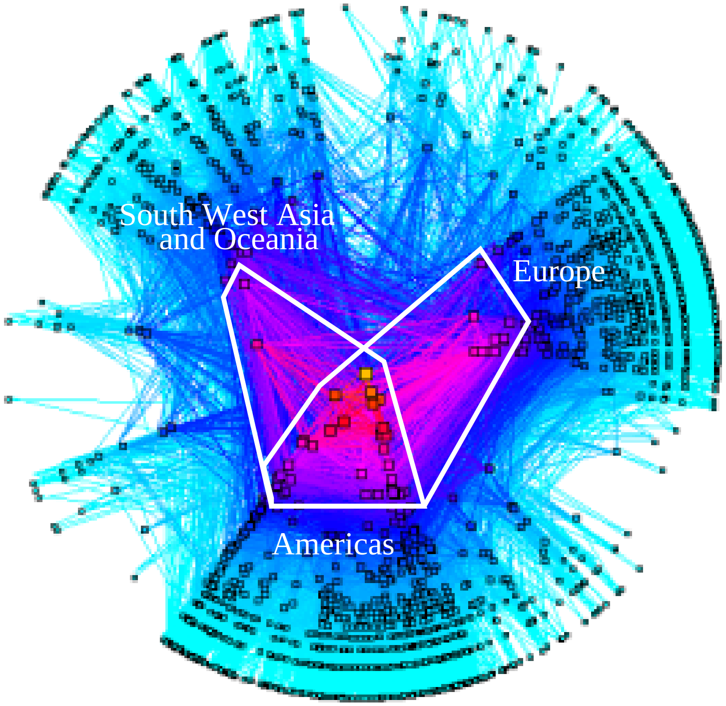

$$r = 1 - log(\frac{out\ degree(AS)+1}{maximum\ outdegree+1})$$

$$\theta = (longitude\ of\ headquaters)$$

- complete AS graph with emphasis on most connected ASes
- depicts geographical ownership of ASes
- illustrates US-centric nature of current Internet topology

visualization : Walrus

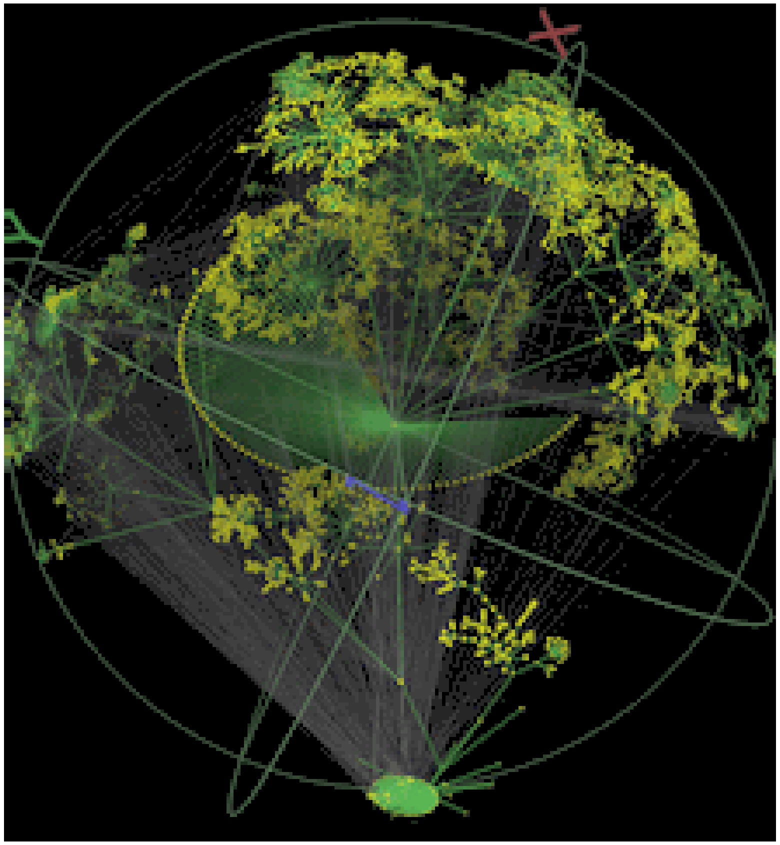

hyperbolic layout graph

visualization:

CAIDA’s San Diego, CA

topology monitor

- uses fisheye-like distortion
- simultaneously displays local detail and global context

visualization : AS dispersion graph

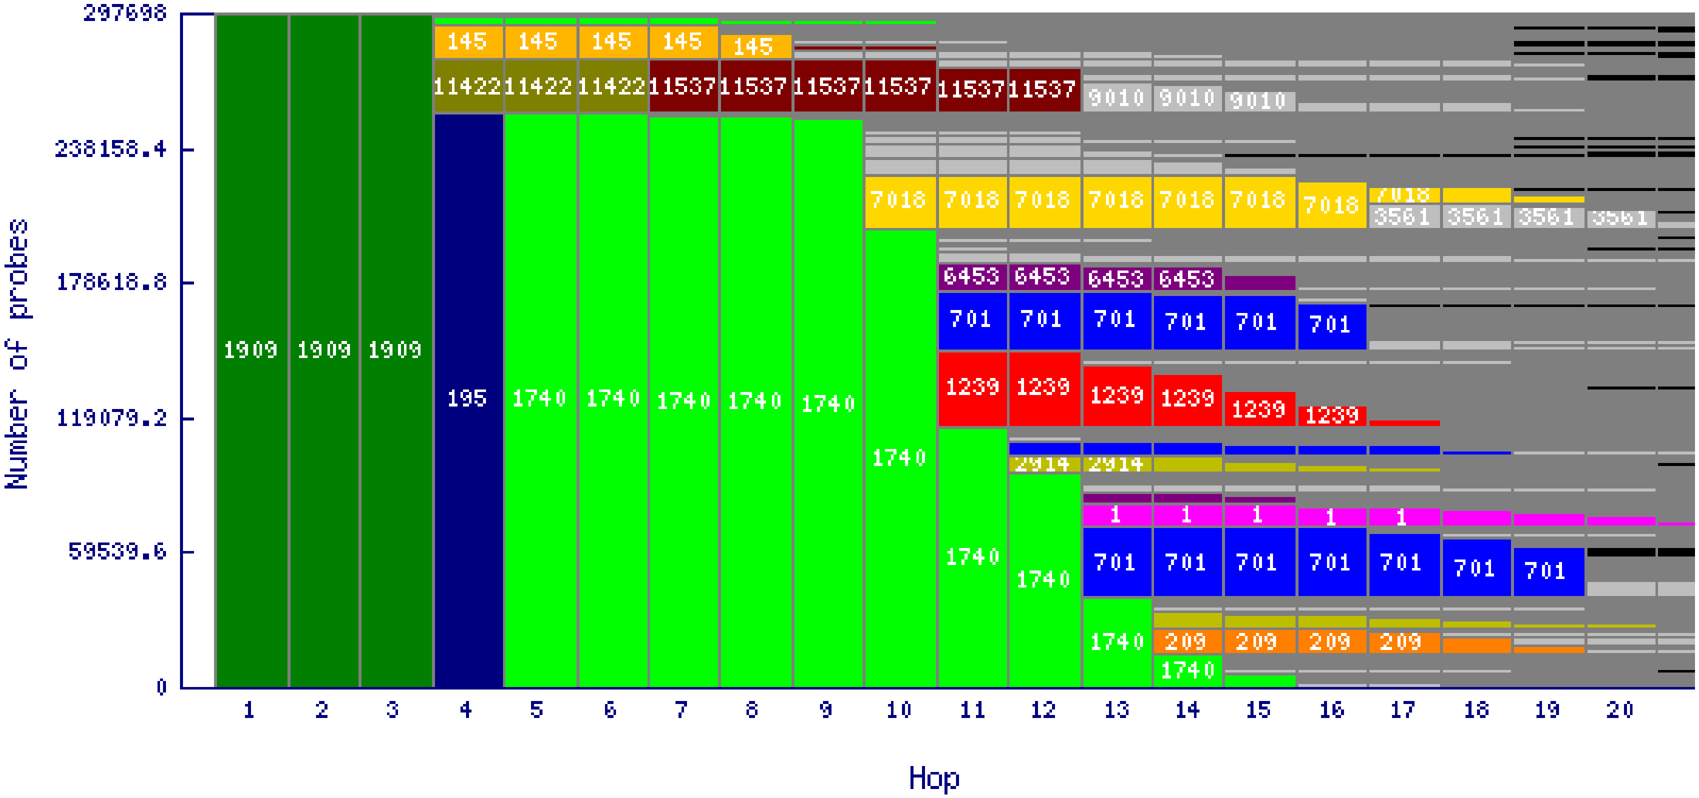

AS dispersion plot for CAIDA's San Diego US topology monitor 24 hours, 27 April 2000

- depicts breakdown of forward paths by percentage from a single source to all destinations
- illustrates ISP diversity, as well as peering points

visualization : PlotPath

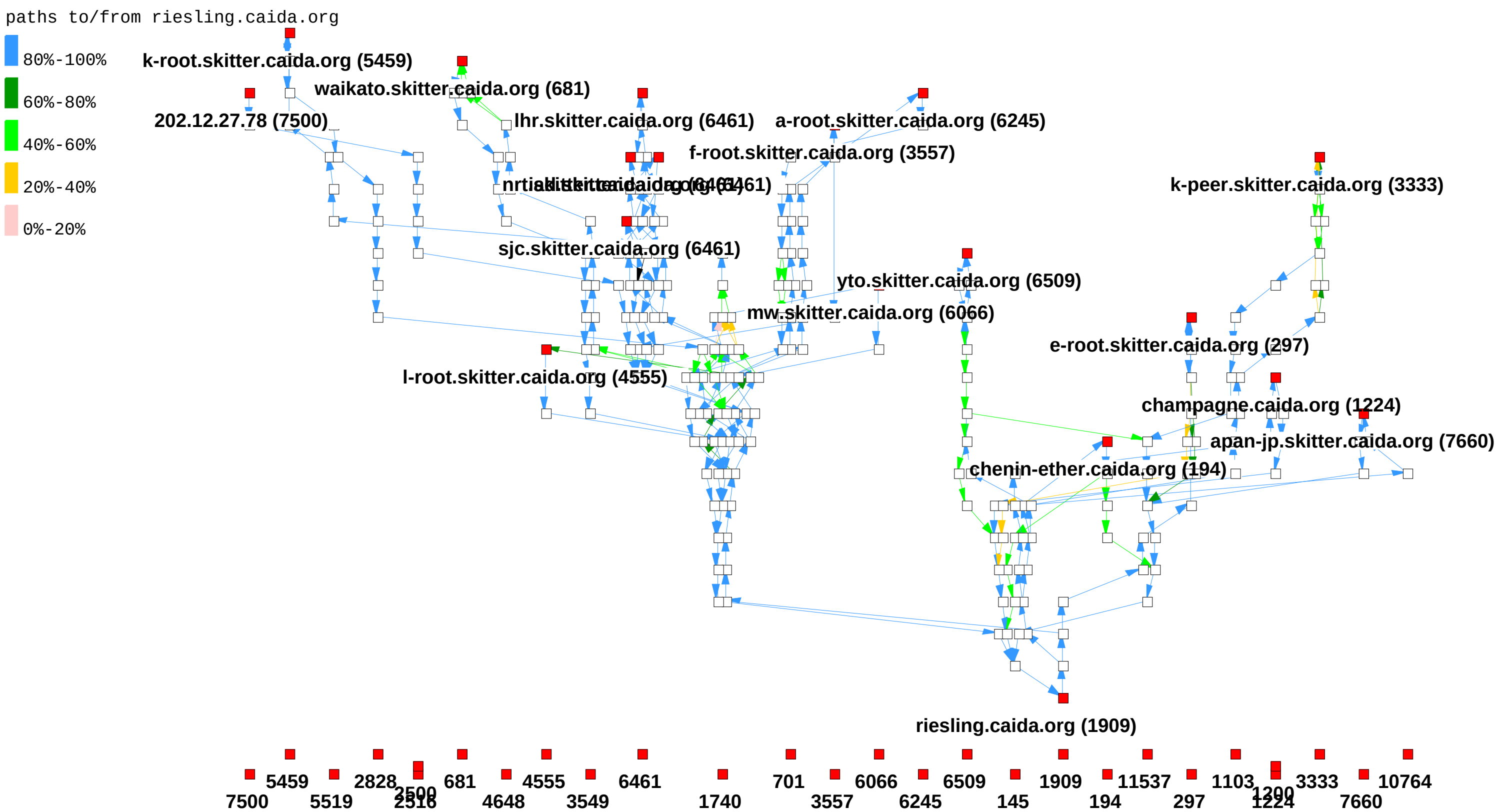

CAIDA's San Diego monitor to all *skitter* monitors, 1-3 Jan 2001

- use forward path info from both monitors to get bidirectional path
- can be used to follow a path between two monitors and observed peering relationships between ASes in these paths

analysis : transit

- **goal:** determine countries providing international transit
- **method:** using skitter traces in October 1998, we counted the percentage of paths that neither started nor ended within a country, but passed through it
- **result:** The United States was primary provider of international transit for Pacific Rim in October 1998.

	US	CA	AU	JP	NZ	EUR	UK	SEA
% transit	71.5	13.3	2.8	1.2	0.9	0.7	0.7	0.3

- EUR : European countries, except for the United Kingdom

- SEA (South-East Asia): Brunei, Indonesia, Papua New Guinea, Philippines, Singapore, Thailand and Viet Nam

- C_H: China and Hong Kong

- SWA (South-West-America): Chile and Peru

analysis : geopolitical classification of large RTTs

- **goal:** find geopolitical areas that were poorly served
  by the DNS (Domain Name Server) root system.
- **method:** We defined 'poor service' in terms of large RTTs
  relative to other clients served at the same time.
  LLDs (Large Latency Destinations) are those in the
  90th percentile of the overall RTT distribution in over
  50% of probing cycles.

analysis : geopolitical classification of large RTTs

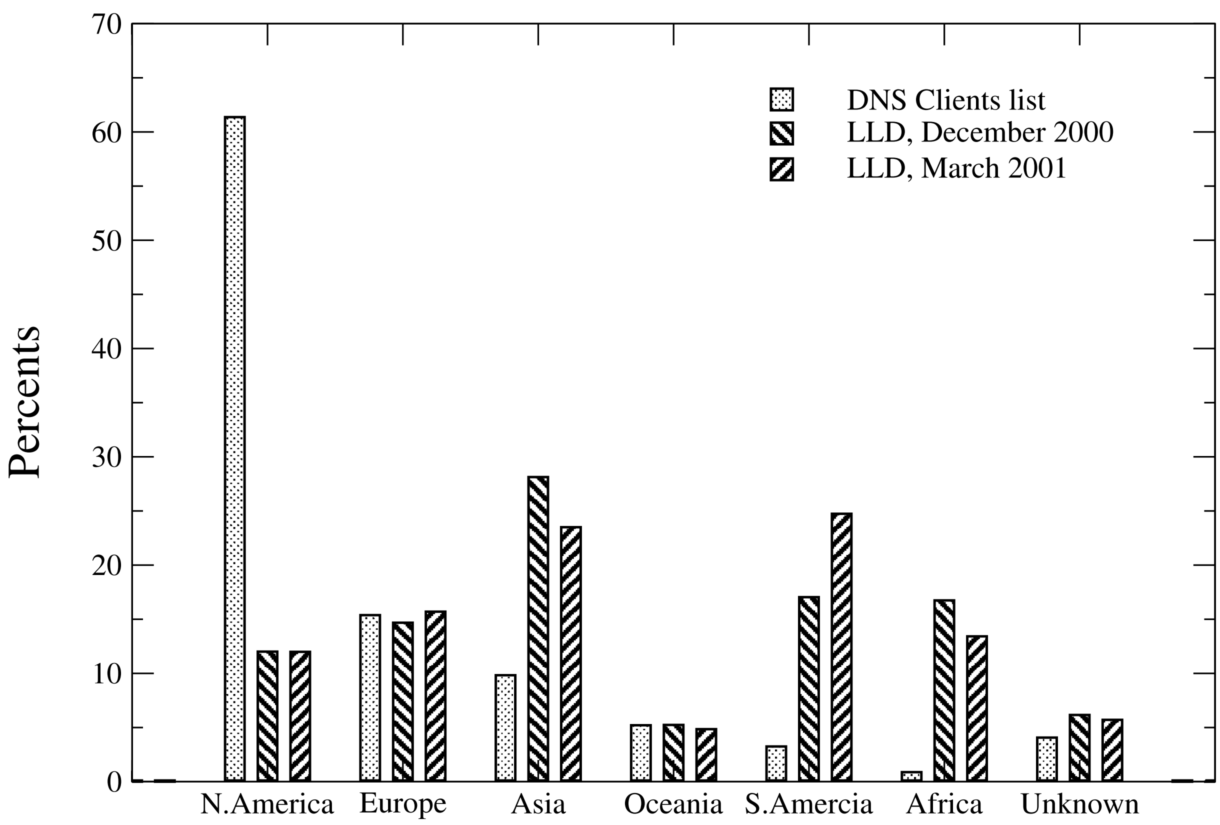

- **results:**
  - DNS client list dominated by American clients
  - LLDs, in contrast, dominated by Asia, South America, Africa

analysis : comparison of multiple distance metrics

- **goal:** compare different metrics of distance between source and destination, in terms of their predictive utility in selecting highest performing server
- **method:** success is the percentage of trials in which a metric successfully selected the destination with the lowest RTT
- **results:**
  - RTT median unsurprising: provides over 90% success rate
  - geographic distance provides 75% success for all US monitors and two non-US monitors; other three non-US monitors had success rates around 60%
  - IP path length is a little better than chance (60%)
  - AS path length is only as good as chance (50%)

analysis : comparing different RTT metrics

- **goal:** compare median, average, and a single sample point in terms of predictive power.
  also see how success rate changed with more information.
- **method:** success rate for each metric compared against each data point using a different number of previously collected points.

analysis : comparing different RTT metrics

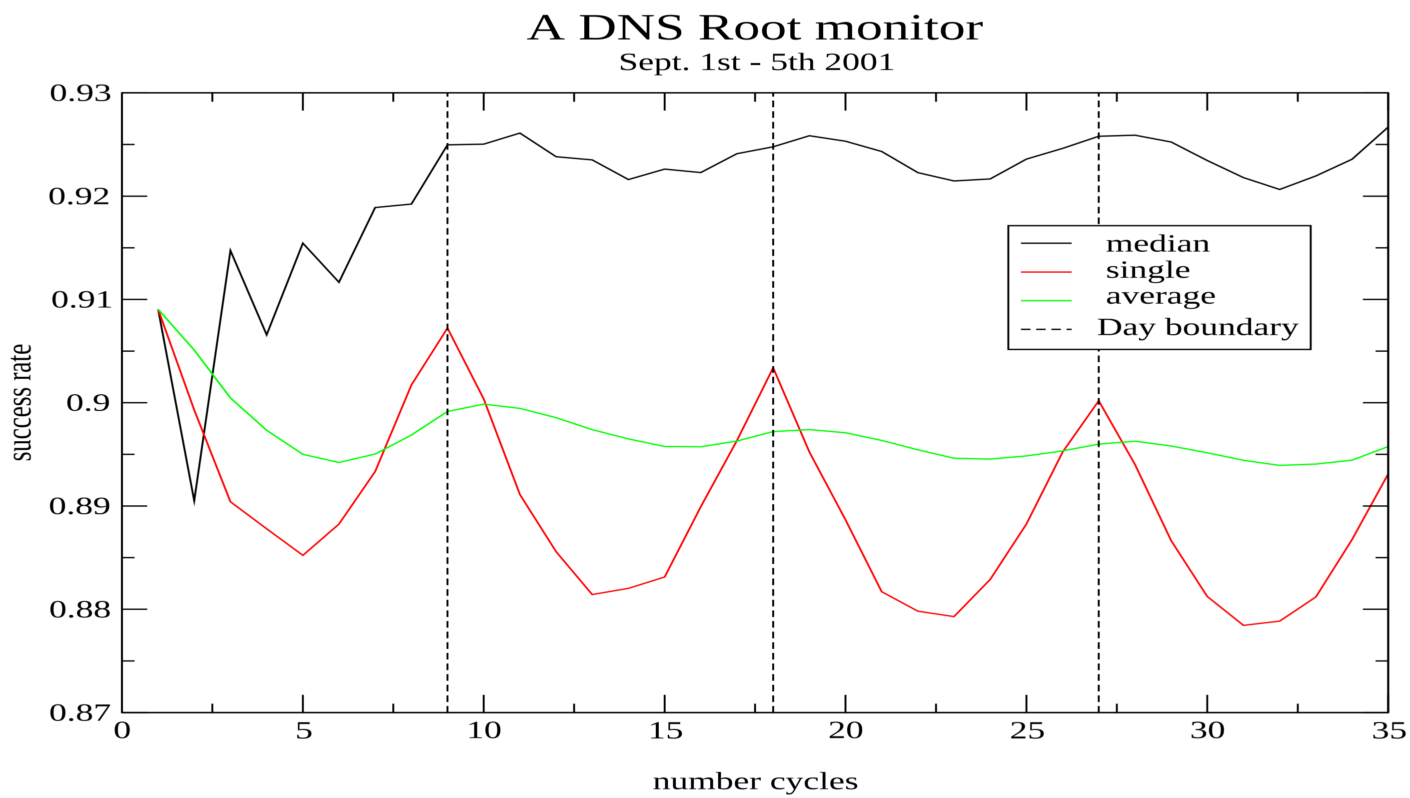

- **results:**
  - median enjoys greatest success, but its success rate decreases after the initial 24 hours
  - taking a single RTT value near the current time of day is more effective than averaging all values in between

conclusions

- CAIDA's probing infrastructure provides a richer model of Internet topology than one based on BGP tables.
- In our October 1998 data sample, the United States provided the majority of transit for international paths.
- The majority of root DNS clients with large RTTs are disproportionately found in Asia, South America, and Africa.
- Median RTT and geographic distance from source are the most powerful metrics for predicting low RTT.
- To optimize success rates when calculating median RTT only data points from the last 24 hours should be used.

future work

- Find the degree to which forward probed IP topology data captures actual Internet topology.
- Characterize the type and frequency of load balancing within the Internet.
- Compare probed IP paths (from src to dst) with corresponding shortest IP paths derived from a macroscopic topology graph (derived from probing from many sources)
- Increase the number and diversity of topology monitors.

availablity of CAIDA’s macroscopic topology data

- The *skitter* data set is available, upon request
  to academic researchers under AUP
- Over thirty different research groups are using *skitter* data

http://www.caida.org/tools/measurement/skitter/research.xml