Leveraging the Science and Technology of Internet Mapping for Homeland Security

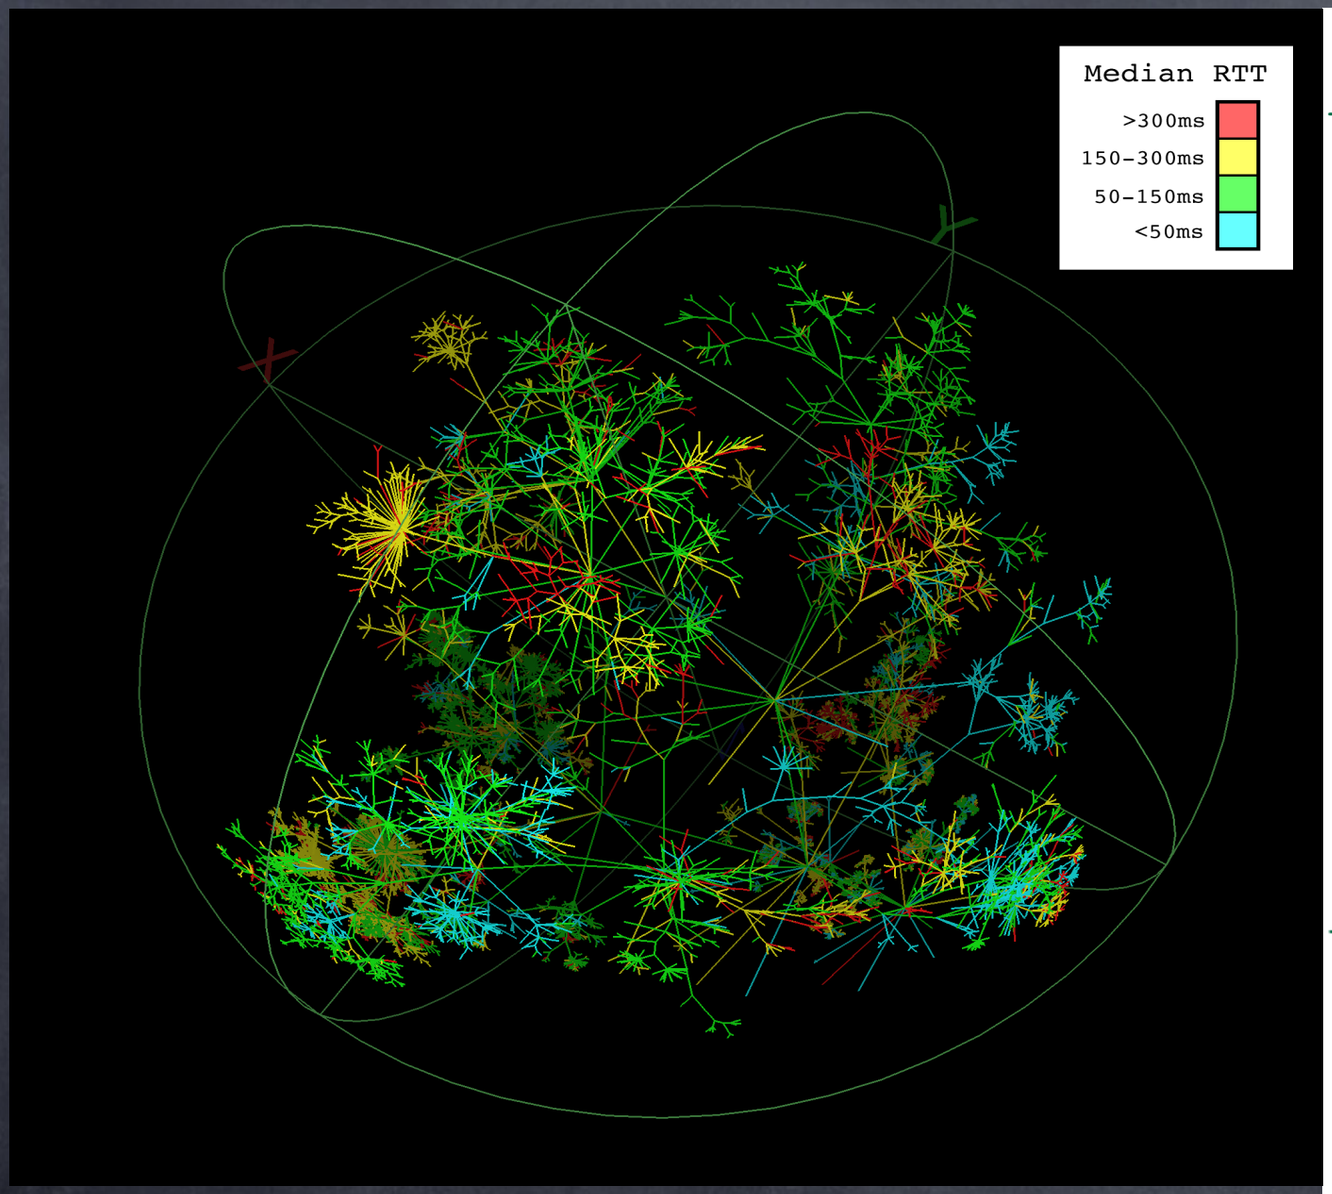

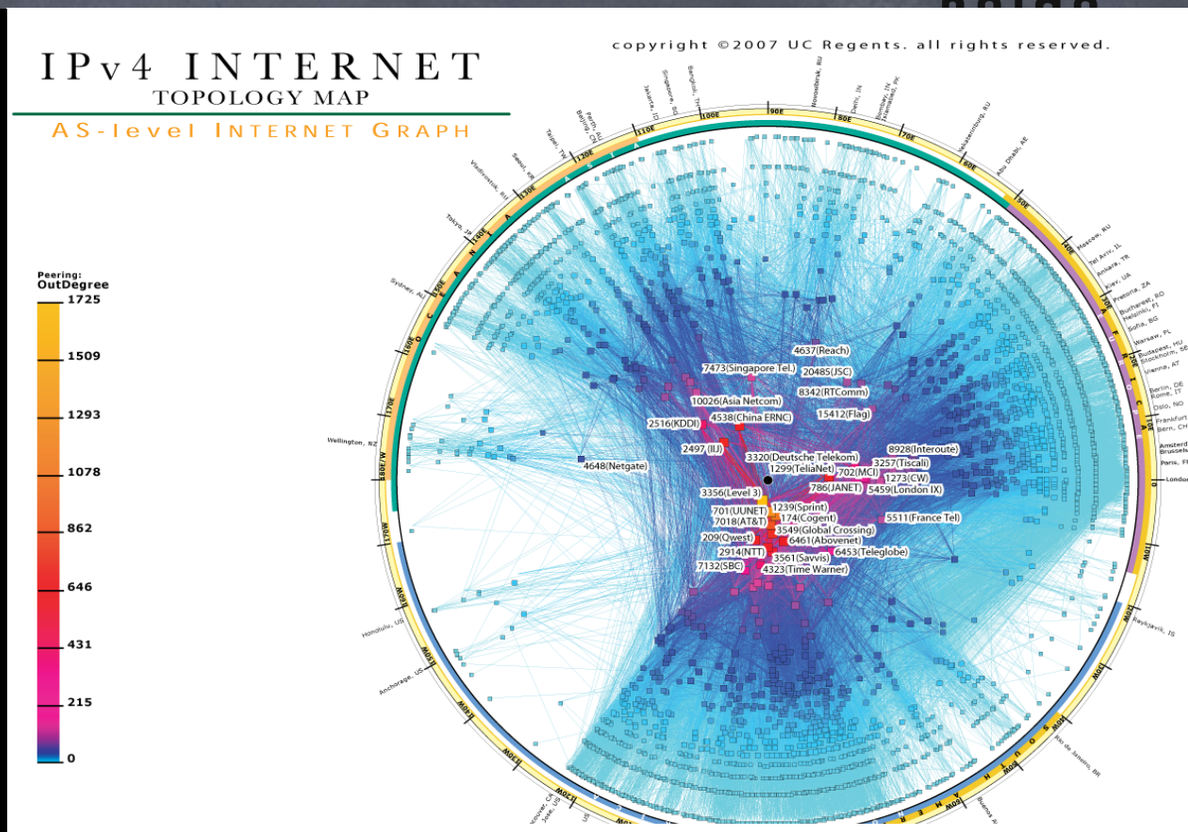

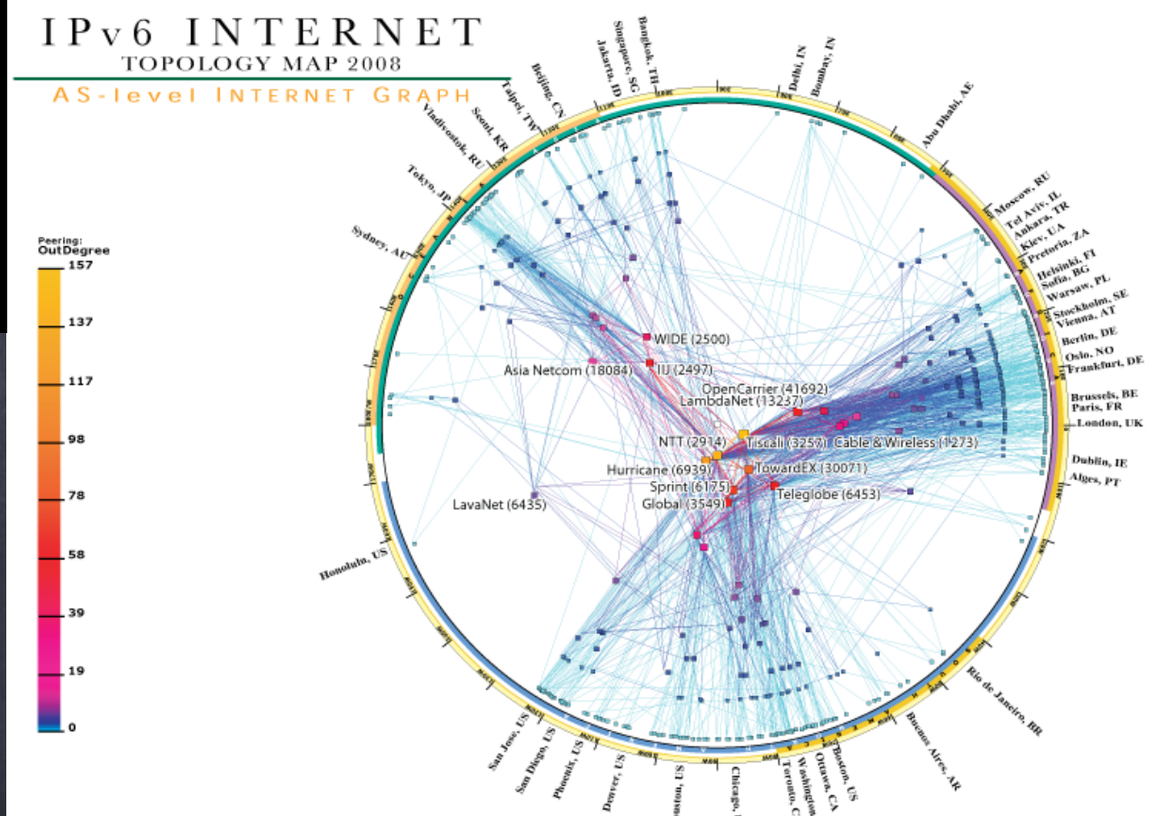

kc claffy

CAIDA
DHS – PI meeting
SRI Menlo Park, CA
10 Sept 2009

Addressing (Inter)national Security Need

To develop and implement new measurement and data collection technologies and infrastructure to improve DHS' situational awareness and understanding of the structure, dynamics and vulnerabilities of the physical and logical topologies of the global Internet.

Macroscopic insight into what we have built...

Where are we (going)?

. Telephone system: 140+ years of history, including regulated data collection requirements (and profits). and a precisely defined system.

·Data networks: 40 years old, ad hoc/hack, tossed to private sector before mature, with no govt support for research or metrics (or profit), ill-defined system.

·Current academic projects either lack sustainability or ability to dedicate resources

·USG spending $350M to build U.S. broadband map.. neither of which has yet been defined

Technical Approach

- Integrate 6 strategic measurement and analysis capabilities:
  - new architecture for continuous topology measurements (Archipelago, or “Ark”),
  - IP alias resolution techniques,
  - dual router- and AS-level graphs,
  - AS taxonomy and relationships,
  - geolocation of IP resources, and
  - graph visualization.

http://www.caida.org/funding/cybersecurity/
http://www.caida.org/projects/ark/
http://www.caida.org/projects/ark/statistics/

New architecture: Ark

- CAIDA's new measurement infrastructure
- Built on decade of from SIGCOMM to MOMA
- Launch 12 Sept 2007
- 37 active probers
- 10 are IPv6-capable

- collaborators can run vetted measurements on security-hardened platform
- publish analyses of views from individual monitors
- support for meta-data mgt, analysis, and infoviz

Nugget of CAIDA's Internet mapping

- **Archipelago** provides a unique enabling infrastructure, featuring the Miranda tuple space, that allows researchers to quickly design, implement, and easily coordinate the execution of experiments across a widely distributed set of dedicated resources (monitors). Ark coordination facilities also enable ease of data transfer, indexing, and archival.

"operating system" for Internet measurement

Benefits to S&T

- Improve critical national capabilities:
  - situational awareness for homeland security purposes
  - internet measurement, analysis and inference techniques
  - topology mapping: annotated AS+router graph (2010)
  - geolocation technology assessment (2010)
  - empirical basis for federal communications policy

- Address network science crisis
  - scalability in system management, monitor deployment, measurement efficiency, resource utilization
  - flexibility in measurement methods
  - let researchers spend less time on non-research

Benefits to Internet researchers

- Ease of experiment design, implementation, and coordination
- Dedicated resources (monitors)
- No restrictive intellectual property
- Multiple levels of trust and access
- Overcome constraints of other platforms that:
  - do not provide dedicated resources
  - cannot guarantee the veracity of the collected data
  - lack fine granularity access control

Insights previously enabled

- Incongruity between topology and routing system
  - topology evolving away from what routing system needs
  - radical implication for future of the Internet (IP)
- Concentration of ISP ownership (as-rank.caida.org)
  - Inform communications, Internet policy
- Incongruity between topology and routing data
  - still no guaranteed way to capture Internet topology
  - but some methods are better than others, e.g., ICMP

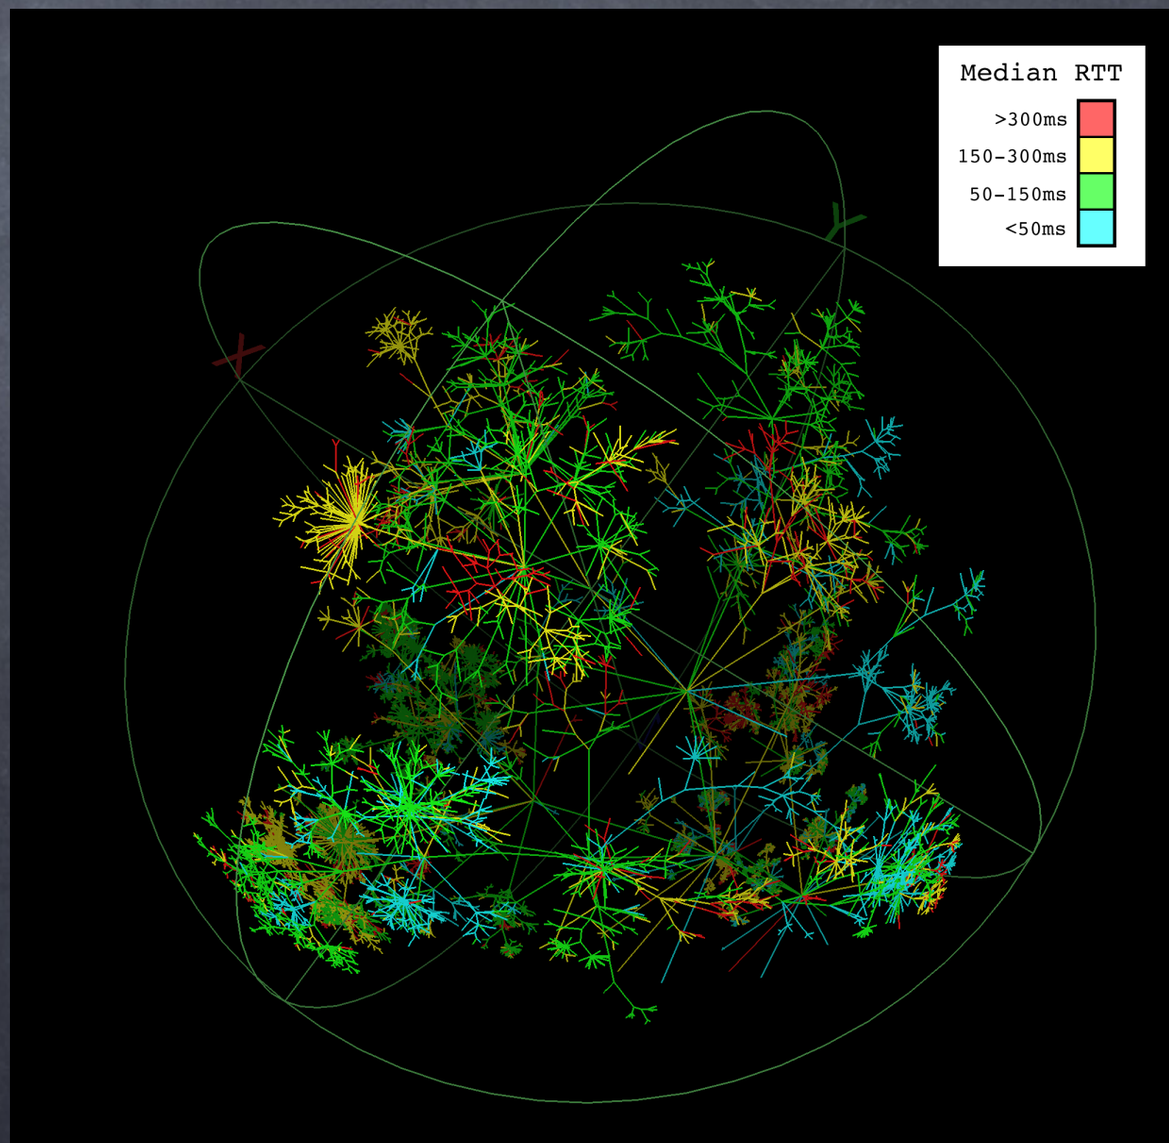

Methodology insights enabled

- What probing method discovers most topology? (IMC2008)
- Do per-flow load balancers implement different forwarding policies for TCP and UDP? (in process)
- IP alias resolution techniques (CCR2009)
- compare performance and accuracy of known alias resolution techniques used at Internet scale
- develop enhancements (kapar, radargun++)
- combine techniques (iffinder, kapar, ally)
- produce most accurate complete IP-to-router mapping
- (while others still saying it's impossible, AMS2009)
- daunting challenge remains validation (not tech problem)

2009 technical accomplishments

- 37 monitors now active, 10 with IPv6
- IPv4 topology data
  - 1.7TB Ark data served thru PREDICT and data.caida.org
  - 23M total traces (paths) in most recent cycle (636)
    - IPv6 topology data
      - AIMS workshop (Feb 2010) {w/PREDICT}
- IP alias resolution recommendations (Oct.)
- Data for IP-to-router resolution (Dec.)
- Ark-based AS-level and router-level graph (Jan.)
- Ark-based dual AS-router graph (June)

probing methodology

- kapar optimizations: incorporate TTL, IP-ID vel., "router links", needed for router-level graphs
- RadarGun optimizations and experiments to complement kapar and iffinder
- Obtained ground truth data from Cogent, Comcast, and HEAnet on router topology
- Determine appropriate measurement windows

On-demand probing functionality

- probe all addresses found for a given hostname
- resolve hostnames on remote Ark boxes
- mesh probing between specified Ark monitors
- plot geographical map of paths (plot-latlong++, NetAcuity)

“Ground truth” vs NetAcuity

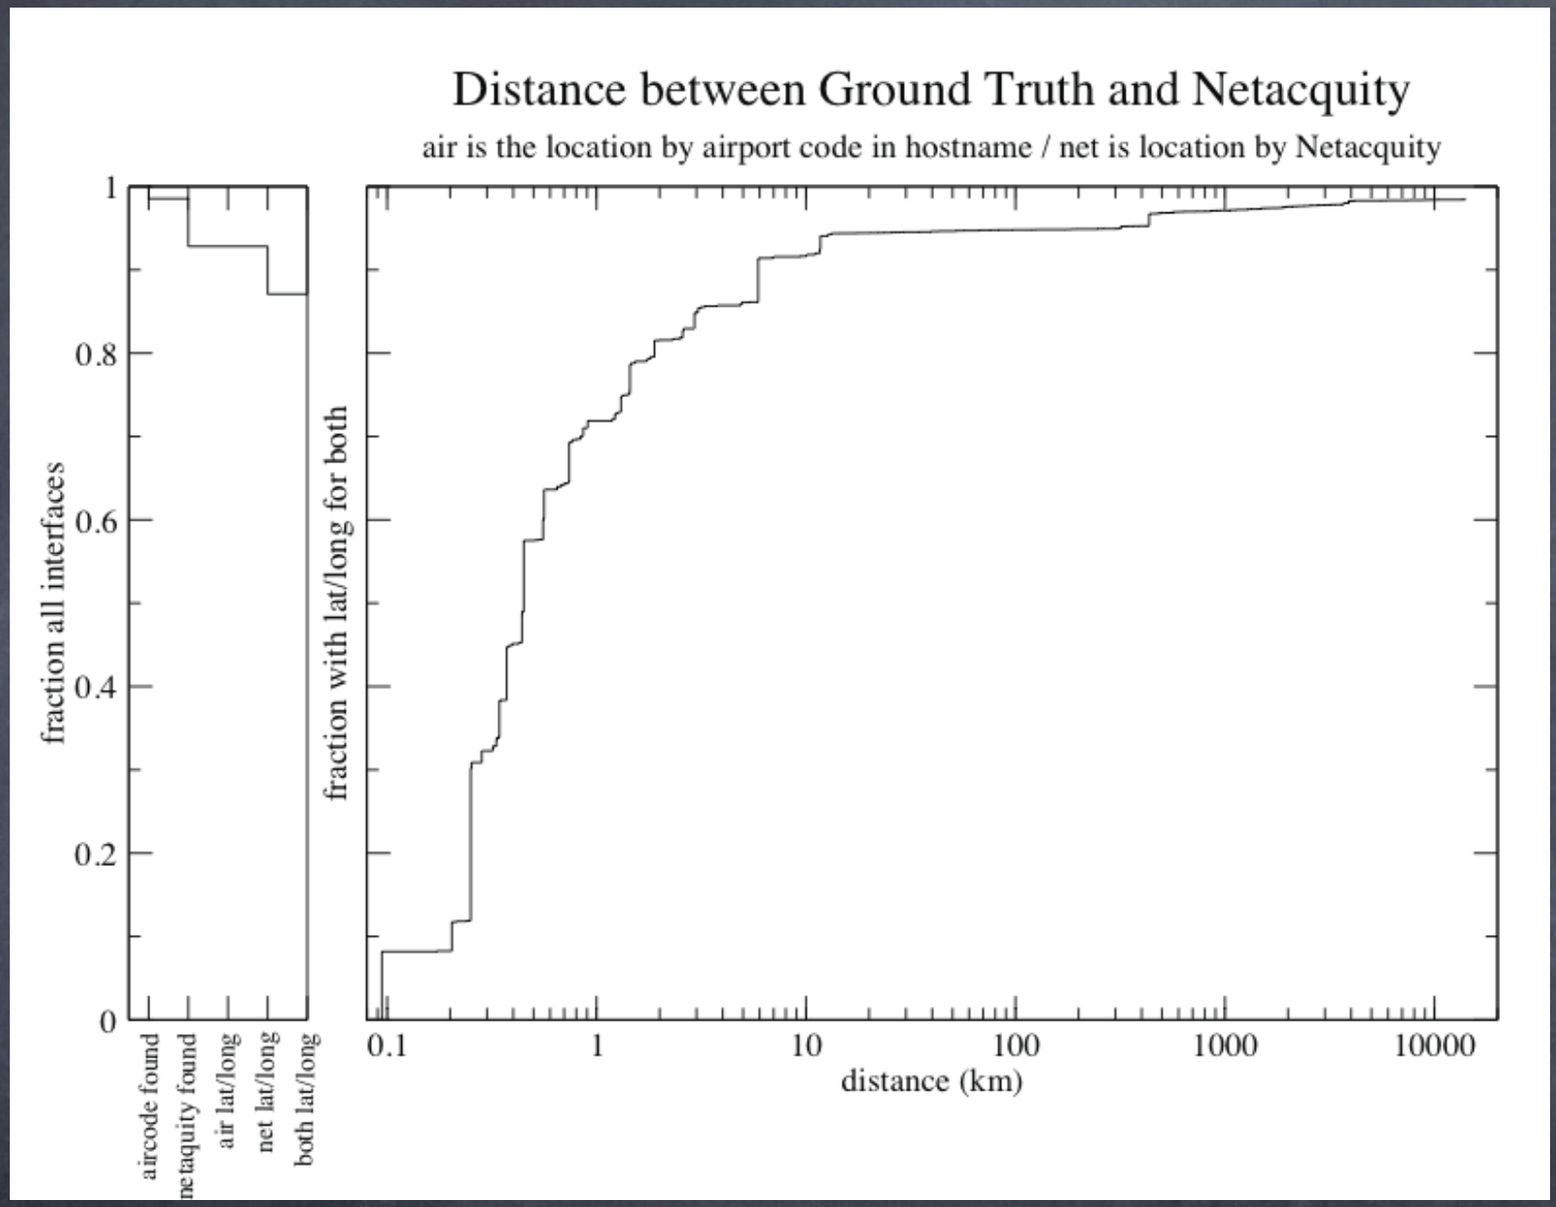

- 88% of interfaces geolocated w/ FQDN & NetAcuity
  - 70% within 1(km) (same city)
  - 90% within 10 (km) (same metro area)
- Planning “bake-off” workshop next year

AIMS workshop

- First Active Internet Measurement Systems workshop (AIMS2009), 12-13 Feb 2009
- Introduce Ark infrastructure capabilities
- Engage research community regarding future needs and interests in collaborations
- Breakouts on research, operational issues
- Final report published in CCR
- http://www.caida.org/workshops/isma0902/
- AIMS 2010 in Feb 2010

Spoofer project

- collaborated with Rob Beverly to support MIT spoofer analysis project
- deploy traffic listeners at (30) Ark monitors
  - extended from one receiver to 30
  - collect UDP probes from test clients
- forwarded to MIT server for analysis
- Report aggregated stats on ‘spoofable’ networks (giving haven to attackers)
- Test your network hygiene!

http://spoofer.csail.mit.edu

In progress: validation

We have complete ground truth data from a tier 1 ISP: backbone interfaces mapped to routers.

In probing every routed /24, we see only 25% (5938 of 23492) of interfaces on ISP's backbone

For more comprehensive validation, we propose intensive probing experiment (w/ ISP's permission)

- two methods of direct probing (of possible aliases) plus extracting from topology probes (maximizing coverage requires preperatory study)

In progress: validation

- Identified direct customers and SFI peers of ISP as reported by our AS-relationship algorithm (Aug09)
  - ~2500 customers (~35000 prefixes, 1.7M /24s)
  - 28 SFI peers (~15000 prefixes, 800k /24s)
- Next: probe this set of /24s from all probing sources.
- Enable view of multiple entry points into ISP
- Then drill down into how many monitors and probing destinations capture how much topology/aliases
- Rocketfuel revisited (revalidated..)
- See next PI meeting..

Routers of unusually high degree

- Some routers w/ degree (inferred) in thousands
- Some of these have been confirmed to be Honeyfarms, simulated computers used to attract hackers
- Another group appears to be the result of faulty alias resolution
- Ongoing attempts to contact ISPs for explanations

Statistical netview from monitors

- General statistics on path lengths and RTTs
- Aggregated views for all monitors

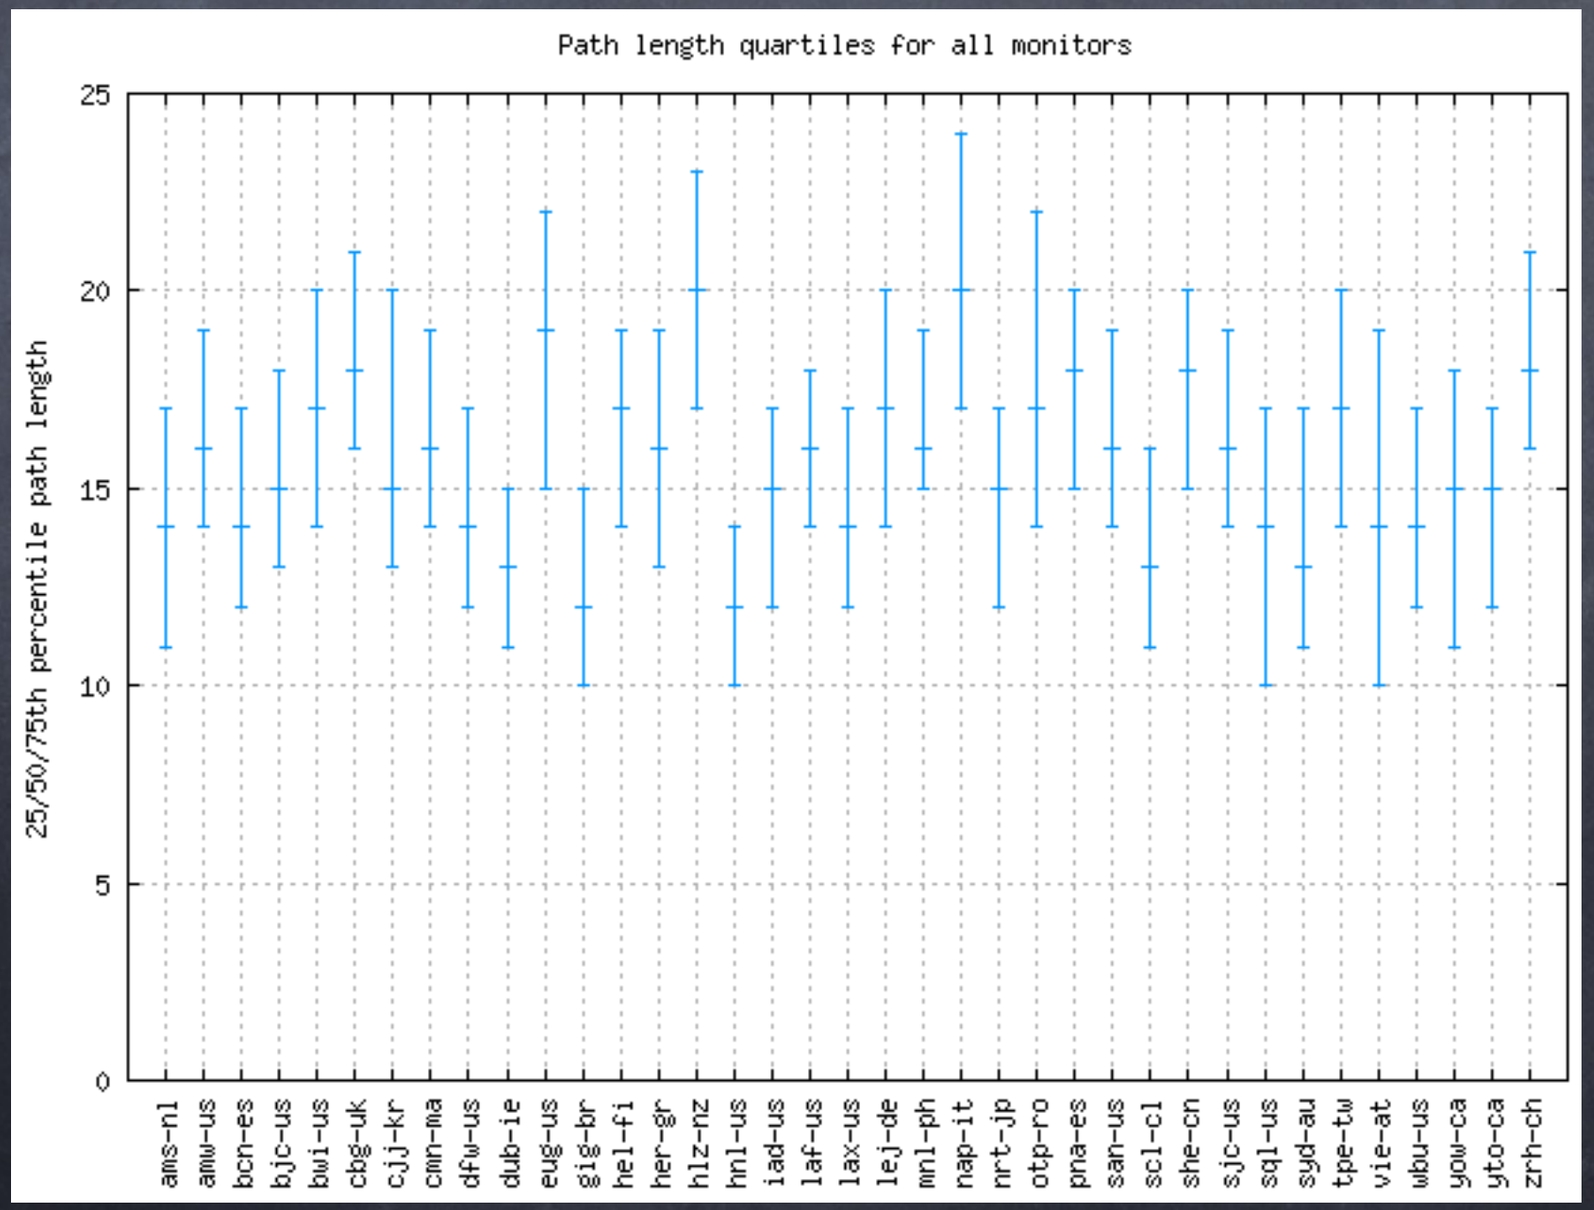

View from monitor: AS/IP dispersion

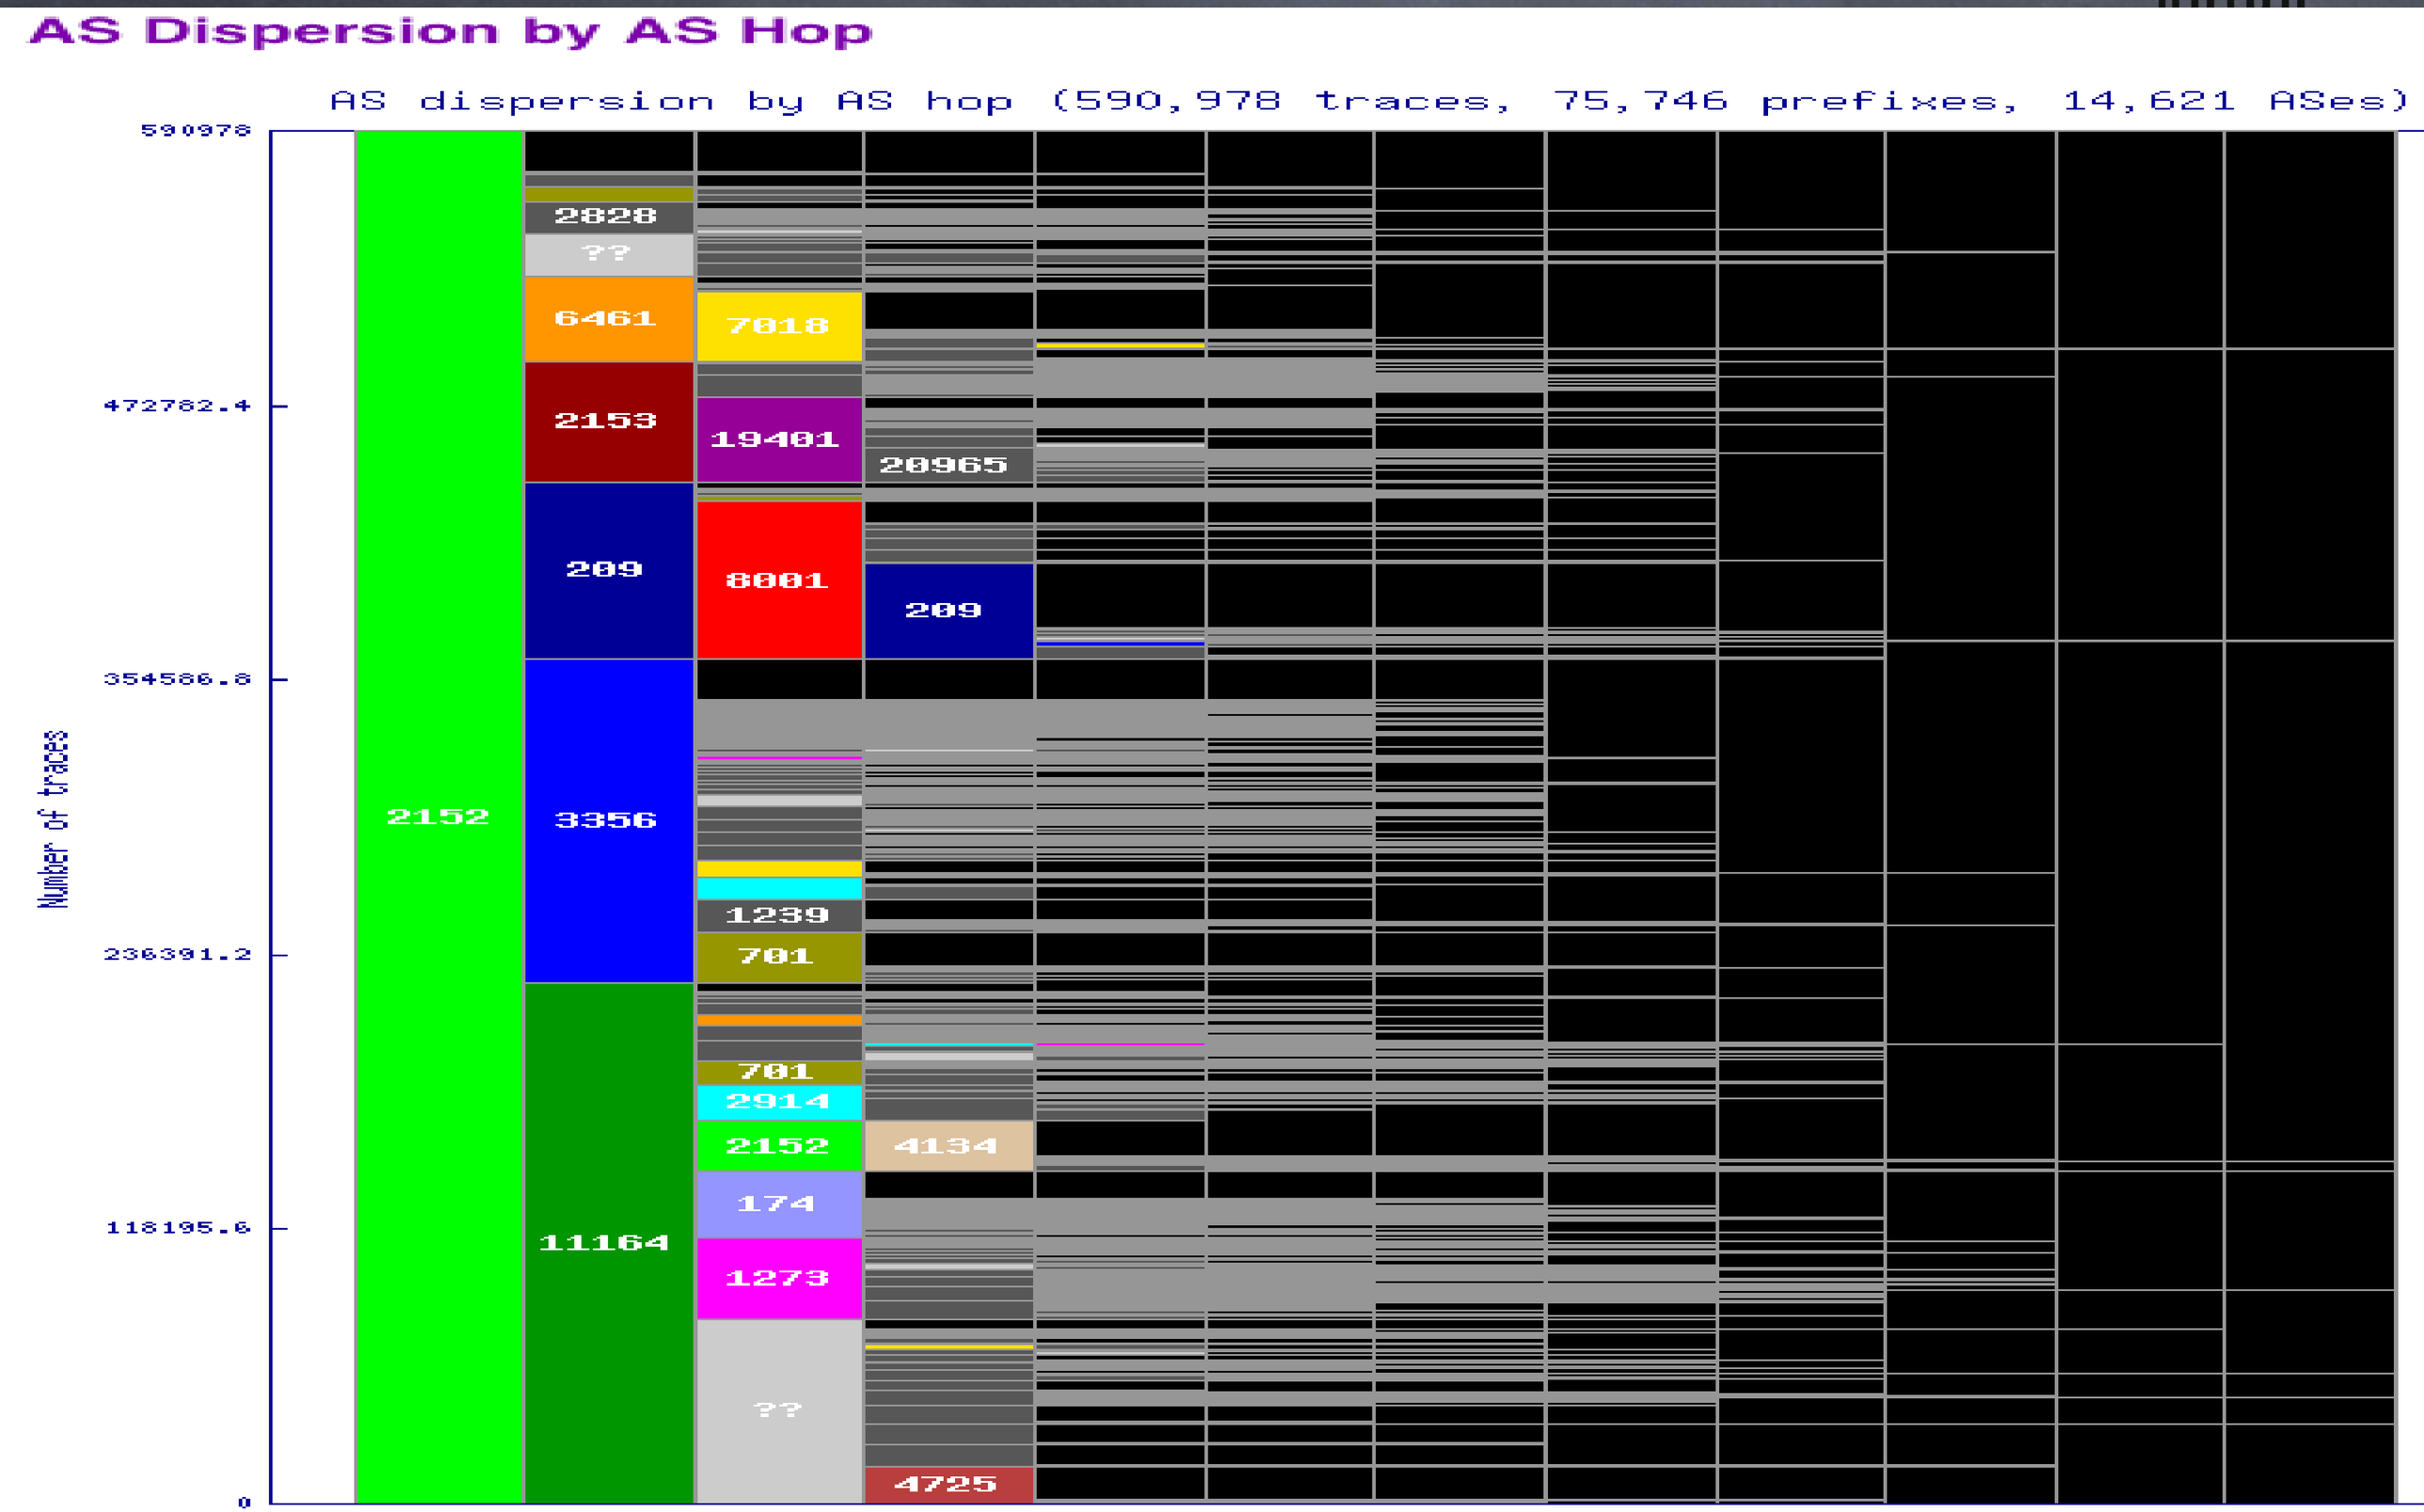

technical accomplishments: views from monitors

Chinese monitor (top) shows IP load balancing over many hops; Chilean monitor (bottom) many fewer IP hops to other ASes.

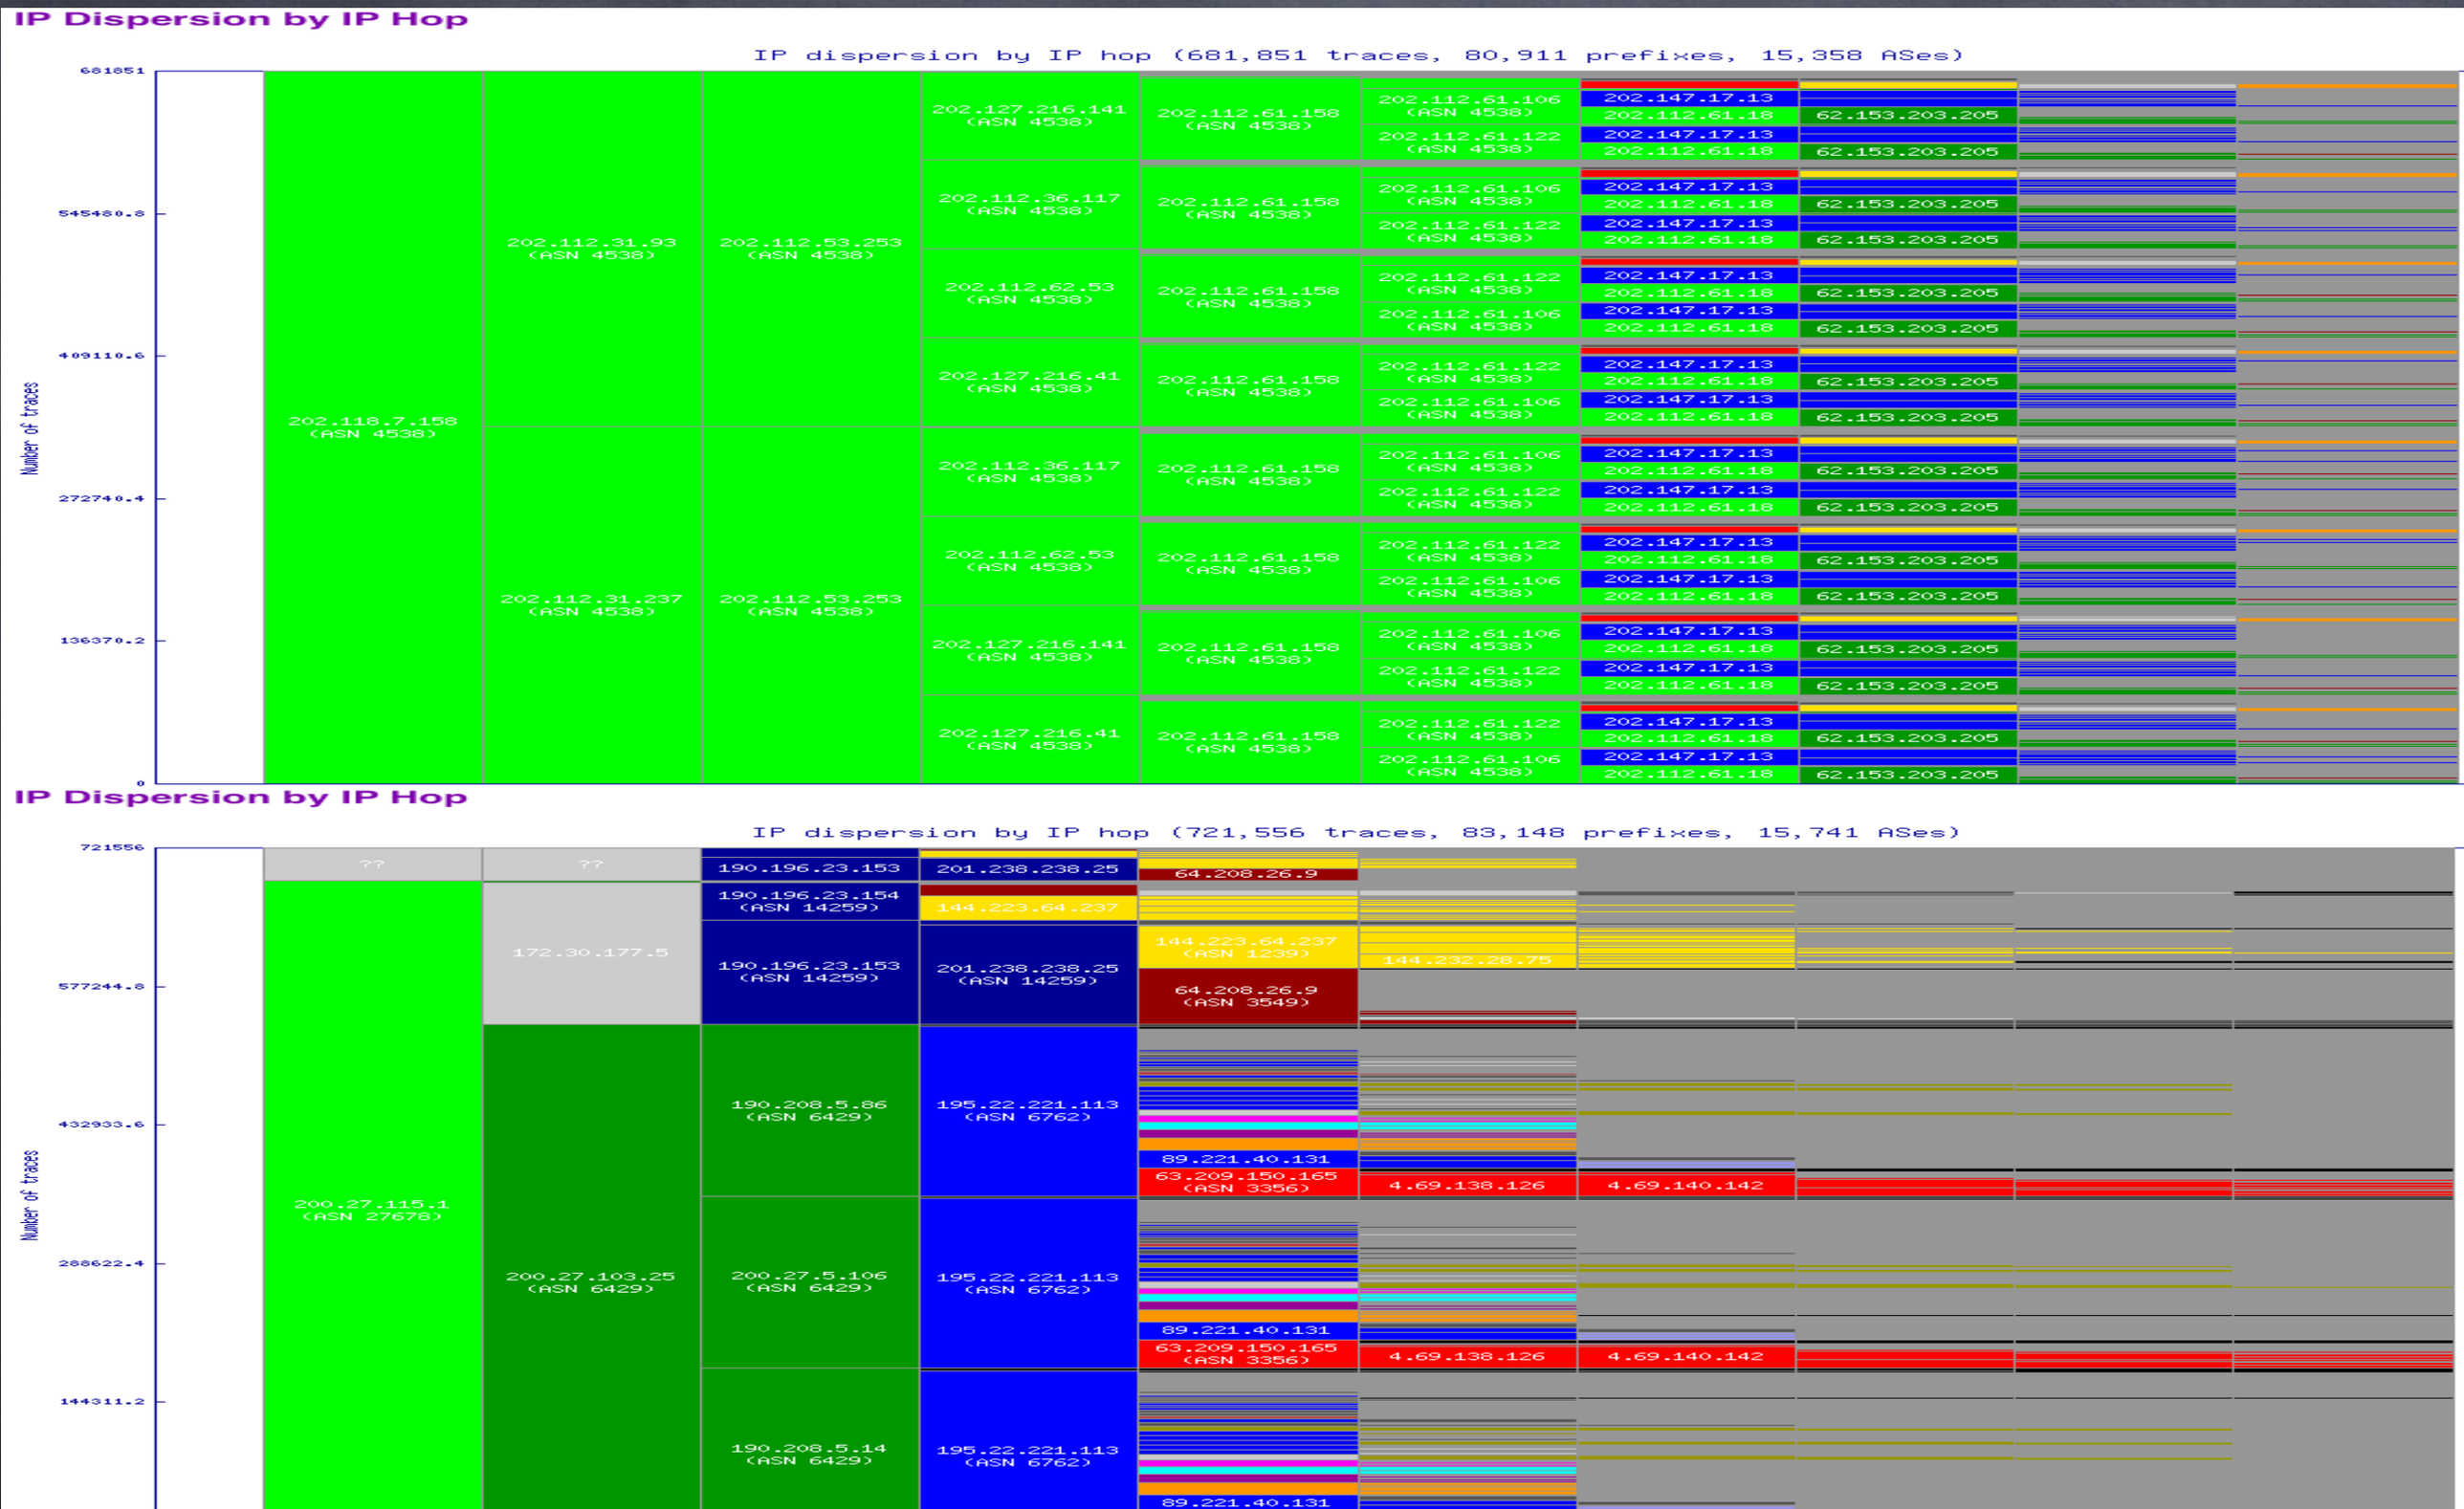

In progress: dual router+AS-level graphs

- Map traceroute data to AS-level
- conceptually simple, well known
- use Route Views BGP tables
- discard and filter ~5% of links in the process
- AS sets, multi-origin & private ASes, indirect links

- Two distinct topologies: AS and router- level
- Need to merge into a dual graph
- assign routers to ASes

- Will evaluate multiple techniques
- *dK-series*, powerful methodology for topo analysis
- validation, validation, validation

digression

Speaking of validation..

Language on data collection that made it thru sausage-making into NOFA for $7B of broadband grants (check back in 2010):

Awardees receiving Last Mile or Middle Mile Broadband Infrastructure grants must report, for each specific BTOP project, on the following:

i. The terms of any interconnection agreements entered into during the reporting period;
ii. Traffic exchange relationships (e.g., peering) and terms;
iii. Broadband equipment purchases;
iv. Total & peak utilization of access links;
v. Total & peak utilization on interconnection links to other networks;
vi. IP address utilization & IPv6 implementation;
vii. Any changes or updates to network management practices;

digression

Speaking of engaging data..

- Lot of people lately saying “policy is the problem”
- But you don't see solicitations for policy research...
  - Landscape of empirical Internet science
  - Prepare for re-write of privacy legislation for Internet age
  - Educating researchers & lawyers

PREDICT-supported papers:

“An Internet Data Sharing Framework For Balancing Privacy and Utility” (6-page intro. Companion paper in progress)

http://www.caida.org/publications/papers/2009/engaging_data/

“Ten Things Lawyers Should Know About Internet Research”

http://www.caida.org/publications/papers/2008/lawyers_top_ten/

Schedule, Planned activities

- 1-2 monitors/month
- IPv4, IPv6 topology data
- Characterize load-balancing behavior
- Try other approaches to dual-graph construction
- Continue alias resolution study, improve tools
- Visualization
- Validation against ground truth
- AIMS 2010
- Submit proposal for BGP data coupling to Ark

BAA Number: Cyber Security BAA 07-09
Title: Science and Technology of Internet Topology Mapping

Offeror Name: Kimberly Claffy
Date: 06/26/07

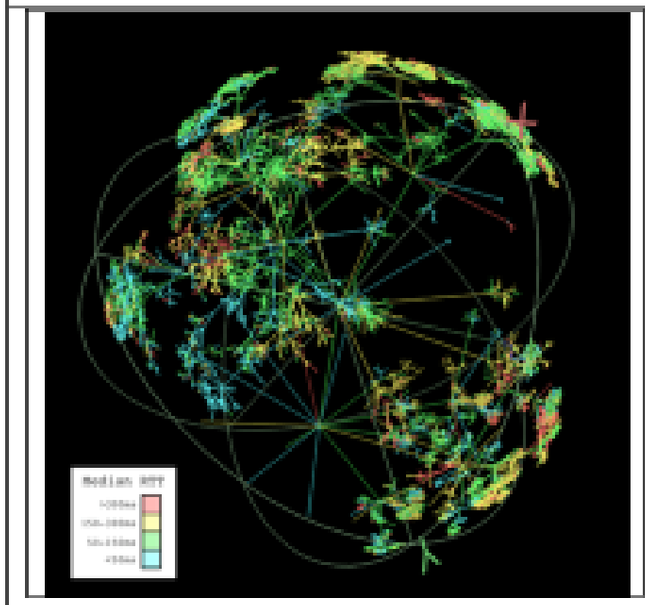

Walrus visualizations of round-trip time measurements made by CAIDA's macroscopic Internet topology monitor located in Herndon, VA, USA.

Internet Topology Mapping:

1. Operational infrastructure to support continuous Internet topology mapping.
2. Periodic active probing of 100% of BGP prefixes announced in publicly available routing tables.
3. ISP relationship inference with accuracy up to 98%.
4. Topologies at the router and AS granularity annotated with AS relationships, AS types, geolocations, latencies, etc.
5. Empirically grounded quantified understanding of robustness, reliability, scalability and other characteristics of the Internet topology as critical infrastructure.
6. Improved annotated topology maps will enhance modeling and monitoring capabilities to help identify threats and predict cascading impacts of damage scenarios.
7. Visualization capabilities will provide powerful interface for use by DHS and other national security personnel.

Technical Approach:

1. Expand current deployment of new distributed platform for continuous measurement of Internet topology, performance, state, and other characteristics.
2. Use and improve IP alias resolution techniques to identify common routers to which IP interfaces belong.
3. Further test and improve performance of software to convert IP technology data into router-level and AS-level graphs.
4. Utilize CAIDA's AS relationship and AS taxonomy inference techniques and data infrastructure to annotate AS graphs with AS types and relationships.
5. Apply and evaluate publicly available geolocation tools for use in annotating topologies with geographic data.
6. Use CAIDA's or other visualization capabilities to depict structure and vulnerability-related characteristics of observed annotated Internet topologies.

Schedule, Deliverables, Contact Info:

1. Current: new active measurement architecture: design complete; prototype implementation being tested.
2. Year 1:
   a. establish on-going IPv4 topology measurements using the new infrastructure;
   b. release software for calculation and exhaustive analysis of topology characteristics.
3. Year 2:
   a. weekly updates of router topology with IP aliases resolved using best available techniques;
   b. weekly updates of AS/router graphs annotated with inferred AS relationships and types.
4. Year 3:
   a. topology annotated with latencies and geolocations;
   b. annotated AS/router topology visualizations.
5. POC: Jennifer Ford, UCSD Contracts&Grants, 9500 Gilman Dr. MC 0934, La Jolla, CA 92093-0934 Fax : (858) 534-0280

Other Links

- Archipelago (Ark) network measurement platform
http://www.caida.org/projects/ark/

- Autonomous System Taxonomy Repository
http://www.caida.org/data/active/as_taxonomy/

- Internet Measurement Conference
http://www.imconf.net/