Leveraging the Science and Technology of Internet Mapping for Homeland Security

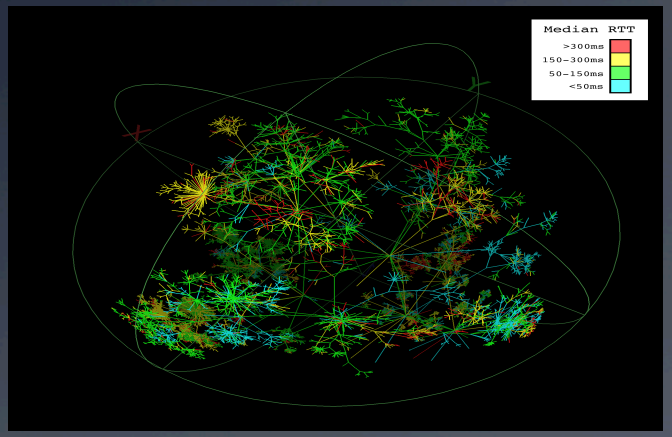

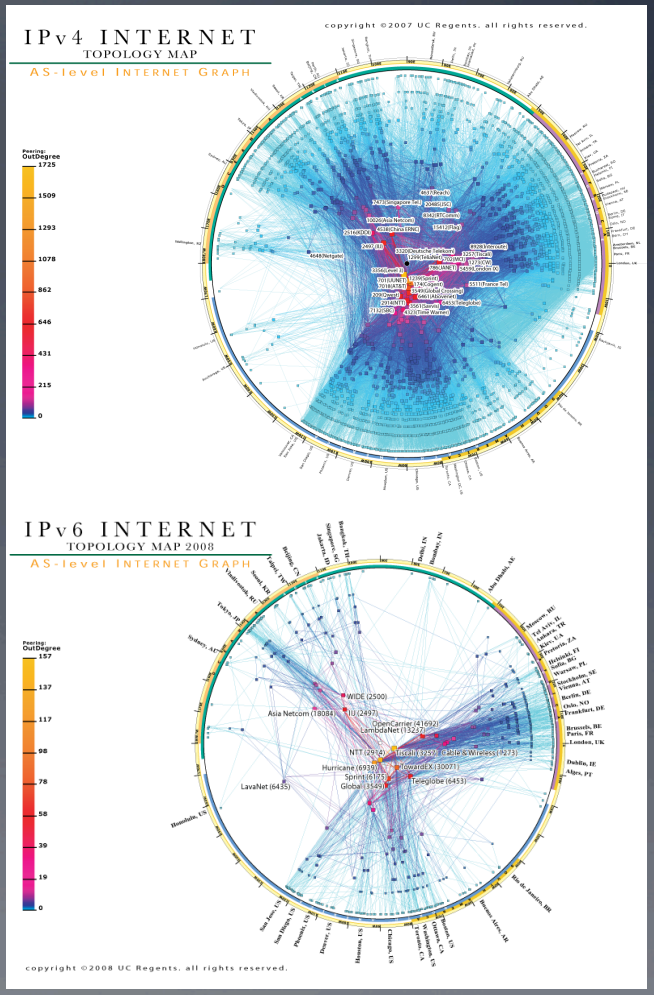

Young Hyun, Ken Keys, Amogh Dhamdhere, Bradley Huffaker, Josh Polterock, Marina Fomenkov, Dima Krioukov, and kc claffy

CAIDA/UCSD
DHS S&T
N66001-08-C-2029
March 2012

http://www.caida.org/

Addressing (Inter)national Security Needs

Objective: to improve DHS' situational awareness and understanding of the structure, dynamics and vulnerabilities of the physical and logical topologies of the global Internet.

Solution: to develop and implement new measurement and data collection technologies and infrastructure.

- *Macroscopic insight into the global Internet infrastructure…*

Technical Approach

- Integrated 6 strategic measurement and analysis capabilities:
  1. New architecture for continuous topology measurements (Archipelago, or "Ark")
  2. Topology analysis techniques, e.g. IP alias resolution
  3. Dual router- and AS-level graphs
  4. AS taxonomy and relationships
  5. Geolocation of IP resources
  6. Graph visualization

http://www.caida.org/funding/cybersecurity/

Benefits to DHS S&T

- Improve critical national capabilities:
  - situational awareness for homeland cyber security purposes
  - Internet measurement, analysis, and inference techniques
  - topology mapping: annotated AS+router graphs
  - geolocation technology assessment
- Address network science crisis:
  - flexibility in measurement methods
  - spend less time on non-research activities
  - rapid prototyping, high-level programming model
- Empirical basis for federal communications policy

Archipelago (Ark)

- Launched 12 Sept 2007 w/ 8 monitors
- 59 active IPv4 probers (March 2012)
  - 17 in US
- 28 active IPv6 probers
- Support for meta-data management
- Collaborators run vetted measurements on security-hardened platform
- Publish statistics and analysis of views from individual monitors

http://www.caida.org/projects/ark/

Ark Infrastructure

- **Archipelago** provides:
  - a powerful, globally distributed measurement infrastructure connected via the Internet to a central server at CAIDA
  - resource coordination using the Marinda tuple space
  - scalable system management
  - versatile and efficient measurement methods
  - flexible scheduling, data transfer, indexing, and archival

An environment for easy development and rapid prototyping of experiments.

Topology Data Architecture

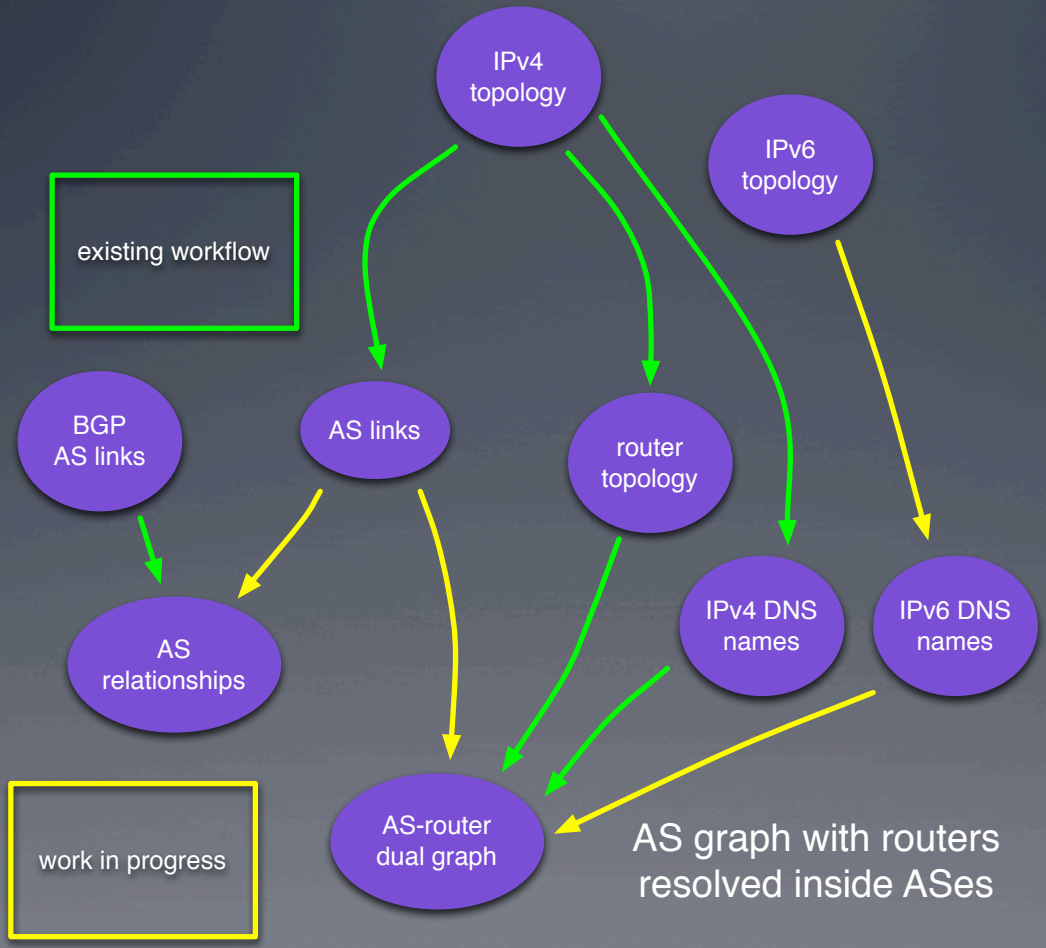

Archipelago Monitor Statistics

- Per-monitor analysis of IPv4 topology data
  http://www.caida.org/projects/ark/statistics/
- Statistics aggregated across all monitors
  - AS path length distributions
  - Integrated RTTs
- Statistics from each monitor
  - Median RTT per country and US state (geographic map)
  - AS hop dispersion graphs (by AS hop and IP hop)
  - IP hop dispersion graphs
  - Distribution of path lengths (IP and AS)
  - RTT distribution (CCDF and quartiles vs hop distance)
  - RTT vs geographic distance

AS Dispersion by AS Hop

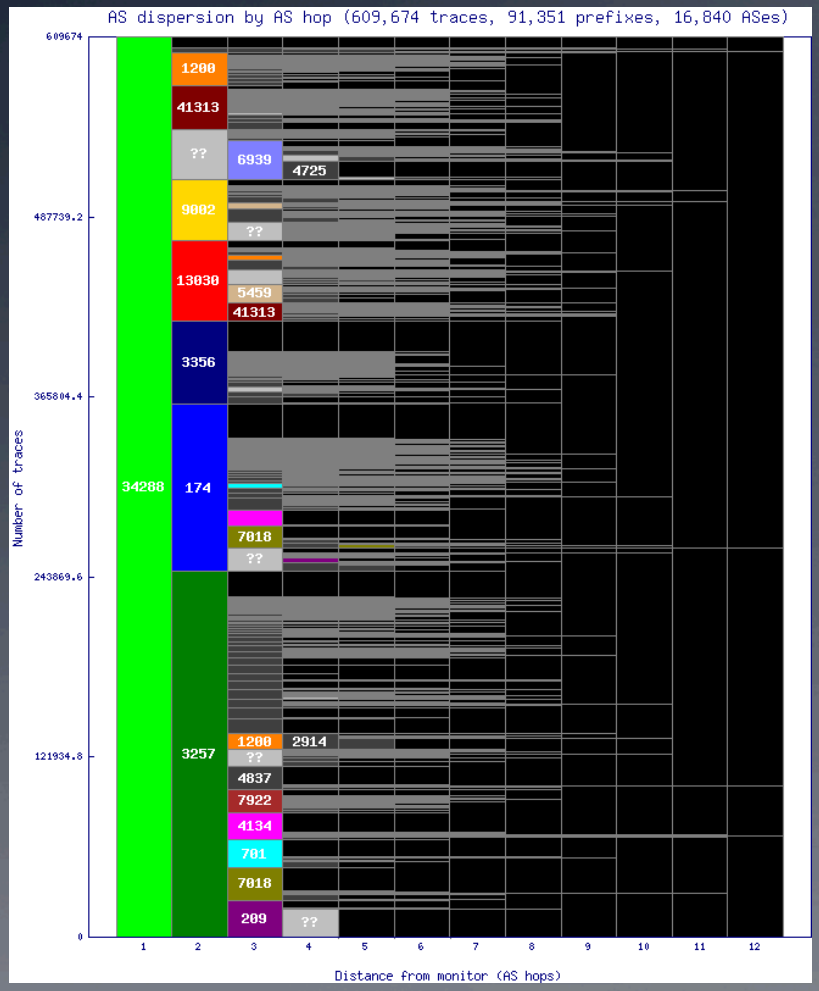

34288	AS34288 EDU-ZG-CH - Public Schools in the Canton of Zug
3257	TINET-BACKBONE Tinet SpA
174	COGENT Cogent/PSI
3356	LEVEL3 Level 3 Communications
13030	INIT7 Init7 Global Backbone
41313	NOVATEL-AS "Novatel" Eood
9002	RETN-AS ReTN.net Autonomous System
7018	ATT-INTERNET4 - AT&T Services, Inc.
1200	AMS-IX1 Amsterdam Internet Exchange B.V.
4134	CHINANET-BACKBONE No.31,Jin-rong Street
209	ASN-QWEST - Qwest Communications Company, LLC
6939	HURRICANE - Hurricane Electric, Inc.
701	UUNET - MCI Communications Services, Inc. d/b/a Verizon Business
5459	LINX-AS London Internet Exchange Ltd.
7922	COMCAST-7922 - Comcast Cable Communications, Inc.
4837	CHINA169-BACKBONE CNCGROUP China169 Backbone
2914	NTT-COMMUNICATIONS-2914 - NTT America, Inc.
4725	ODN SOFTBANK TELECOM Corp.

Monitor: Kantonsschule Zug (zrh2-ch)

AS Dispersion by IP Hop

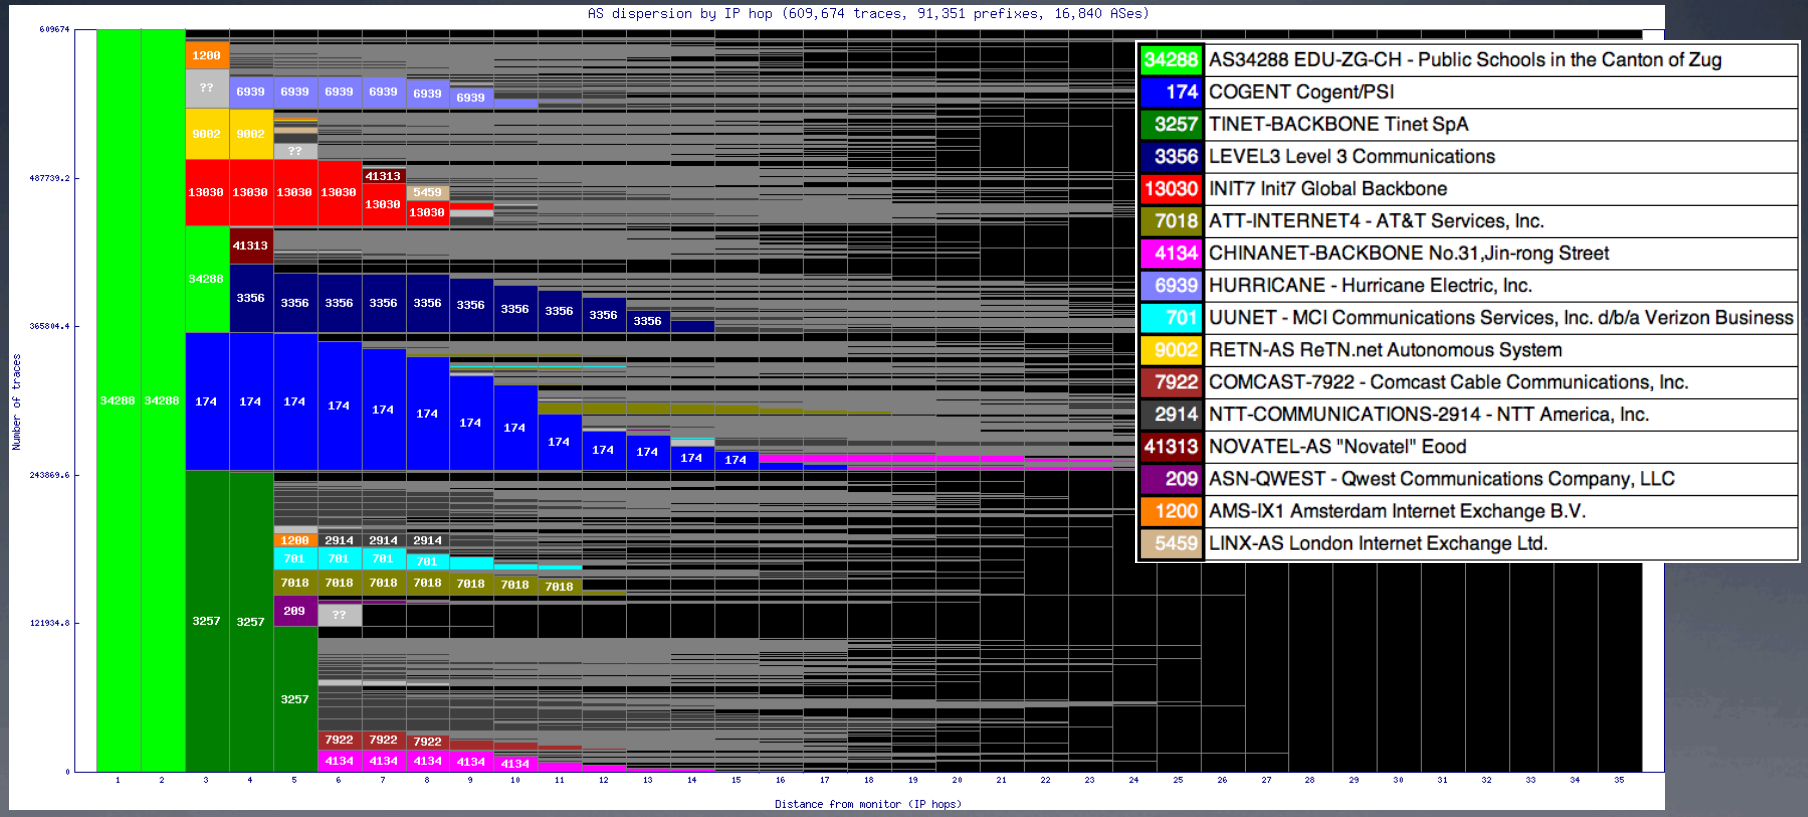

Monitor: Kantonsschule Zug (zrh2-ch)

AS Dispersion by IP Hop: shows load balancing

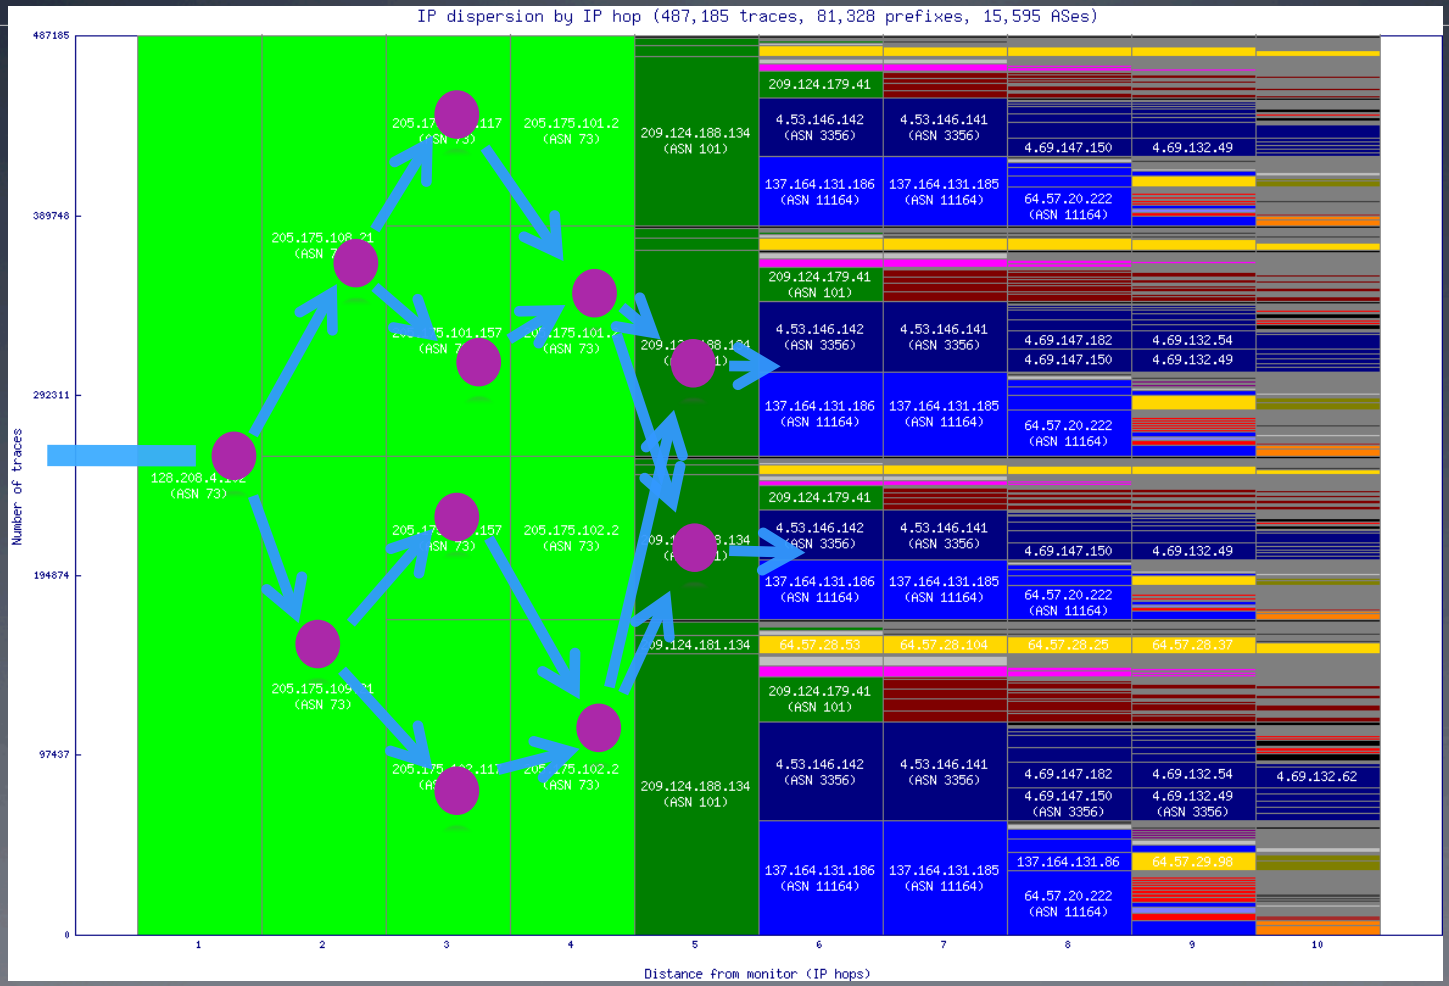

Monitor: University of Washington (sea-us)

Median RTT to Destination Countries

- RTT plotted by country
  - Geolocate destinations with Netacuity (MaxMind Lite for public release)
  - Color each country by median RTT destinations

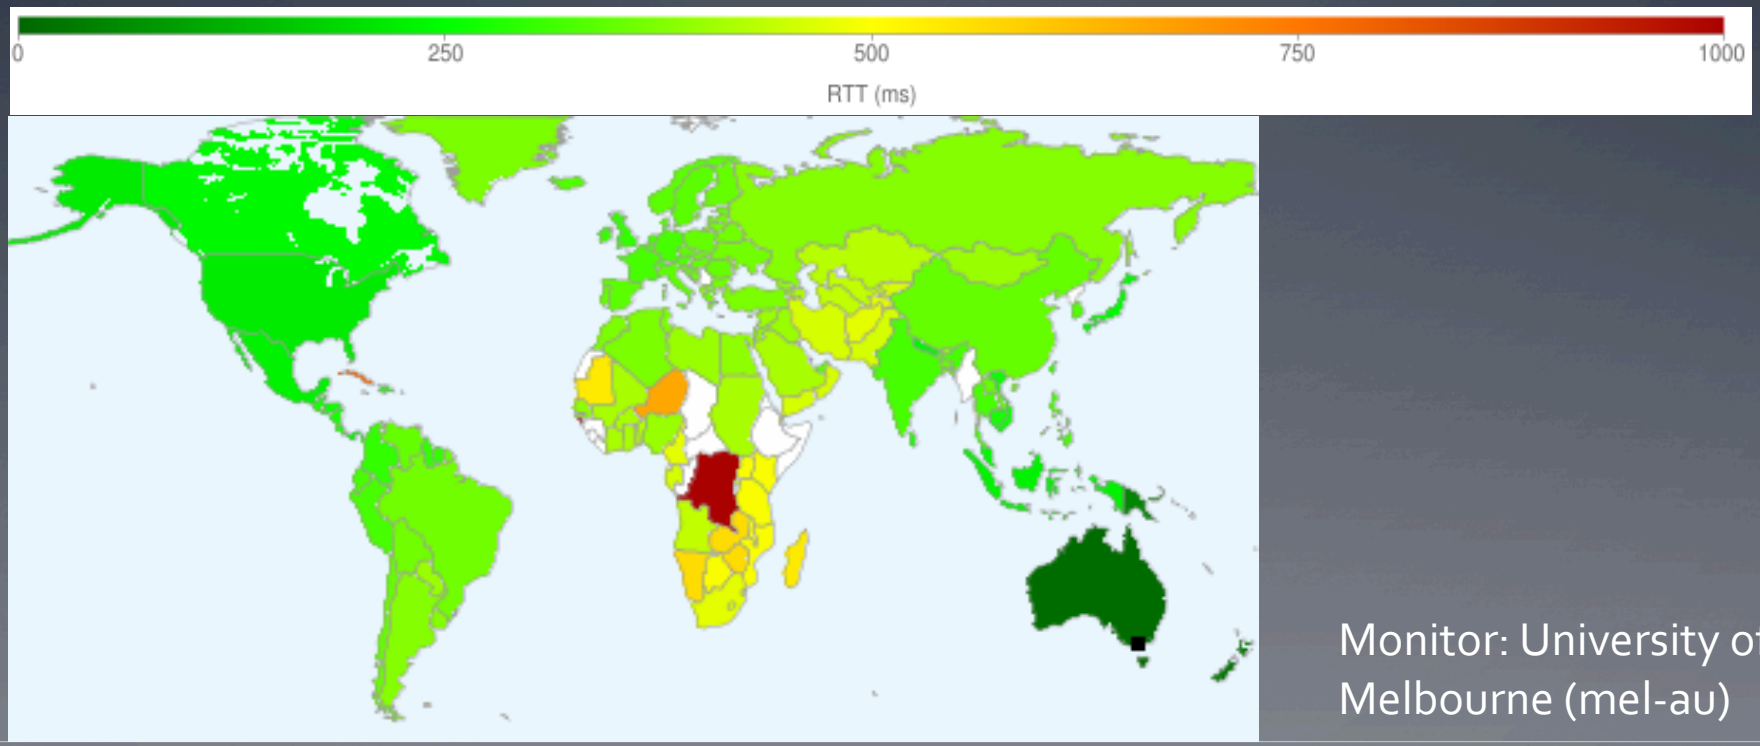

Monitor: University of Melbourne (mel-au)

Median RTT to Destination Countries

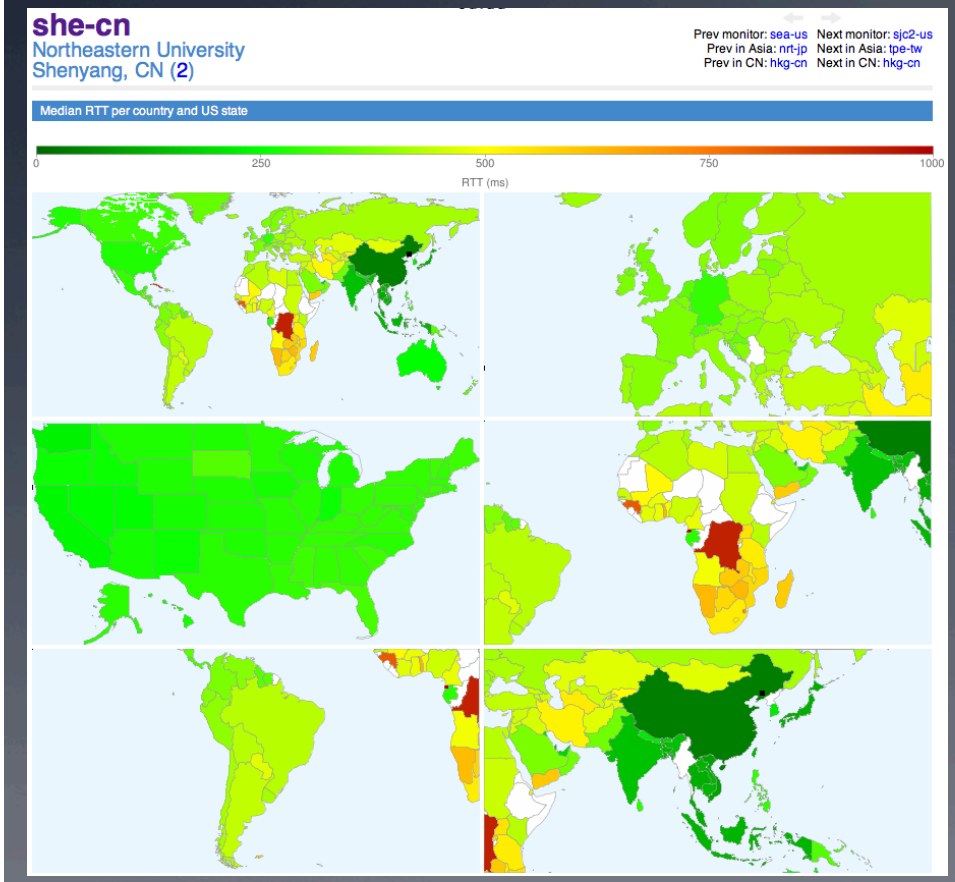

Monitor: Northeastern University (she-cn)

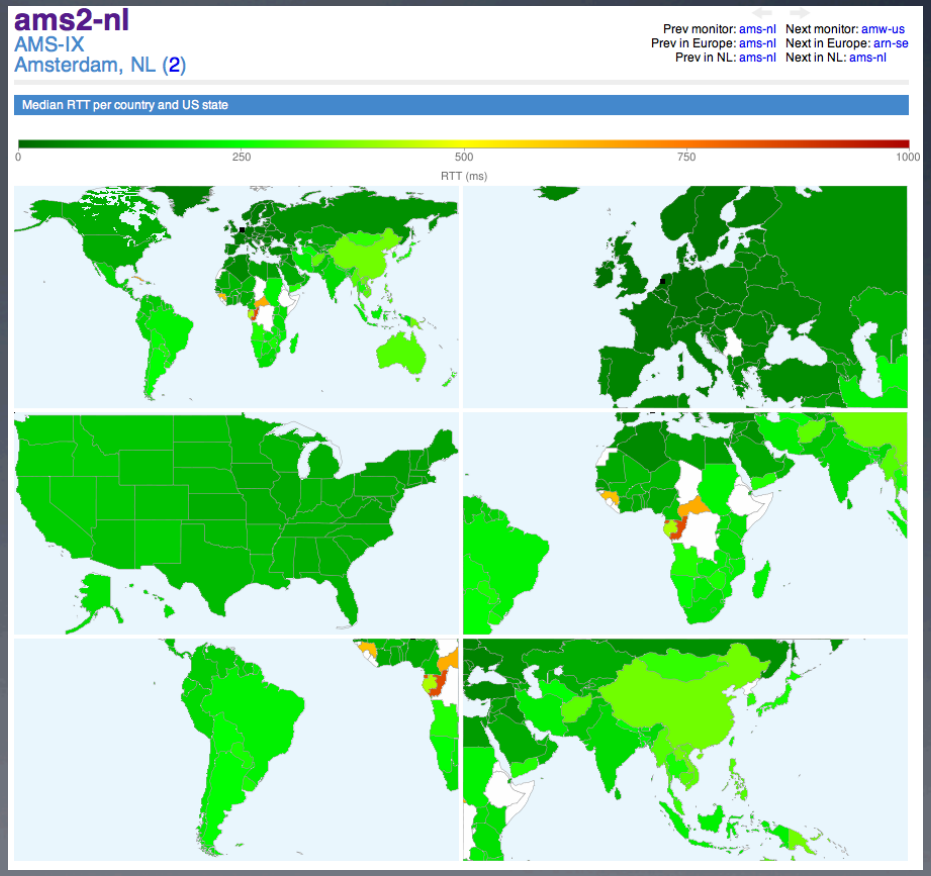

Monitor: AMS-IX Amsterdam, NL (ams2-nl)

Median RTT to Destination Countries

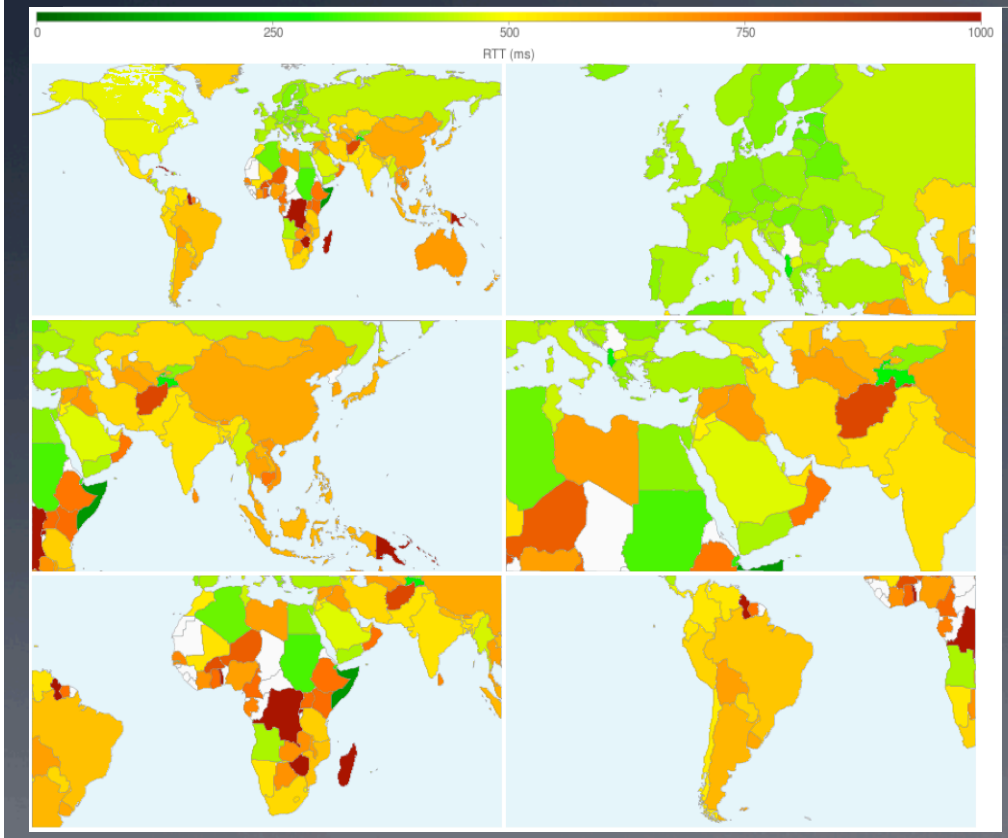

Sept 2010. Prior to new fiber down the west coast of Africa.

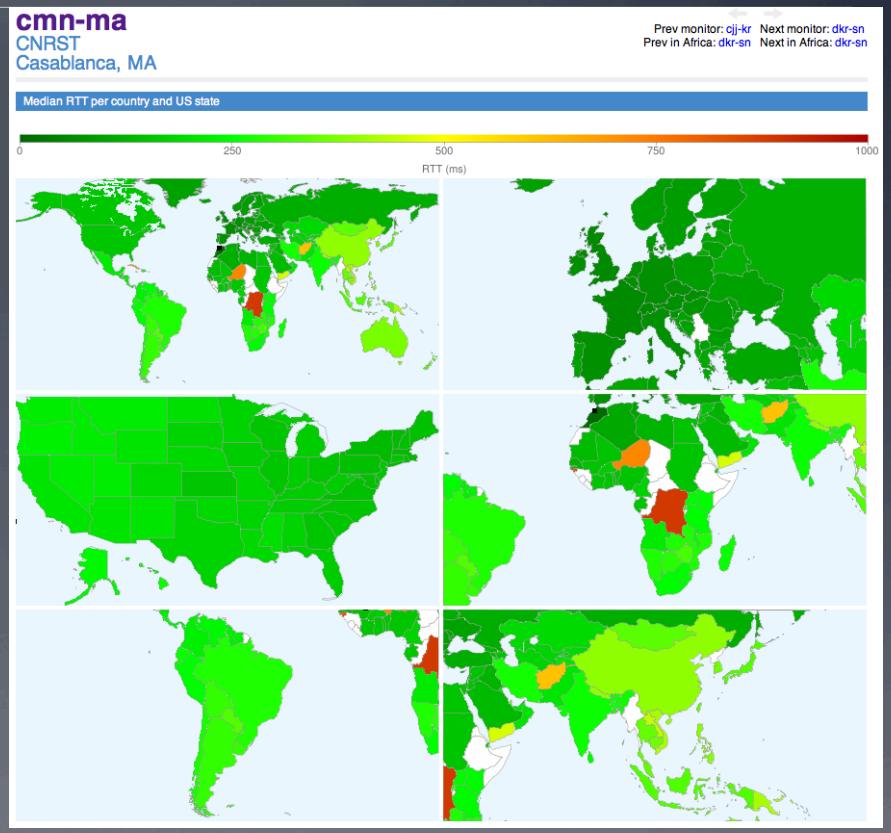

Oct 2011. After new fiber down the west coast of Africa.

Monitor: CNRST Casablanca, Morocco (cmn-ma)

IP Path Dispersion (by IP Hop)

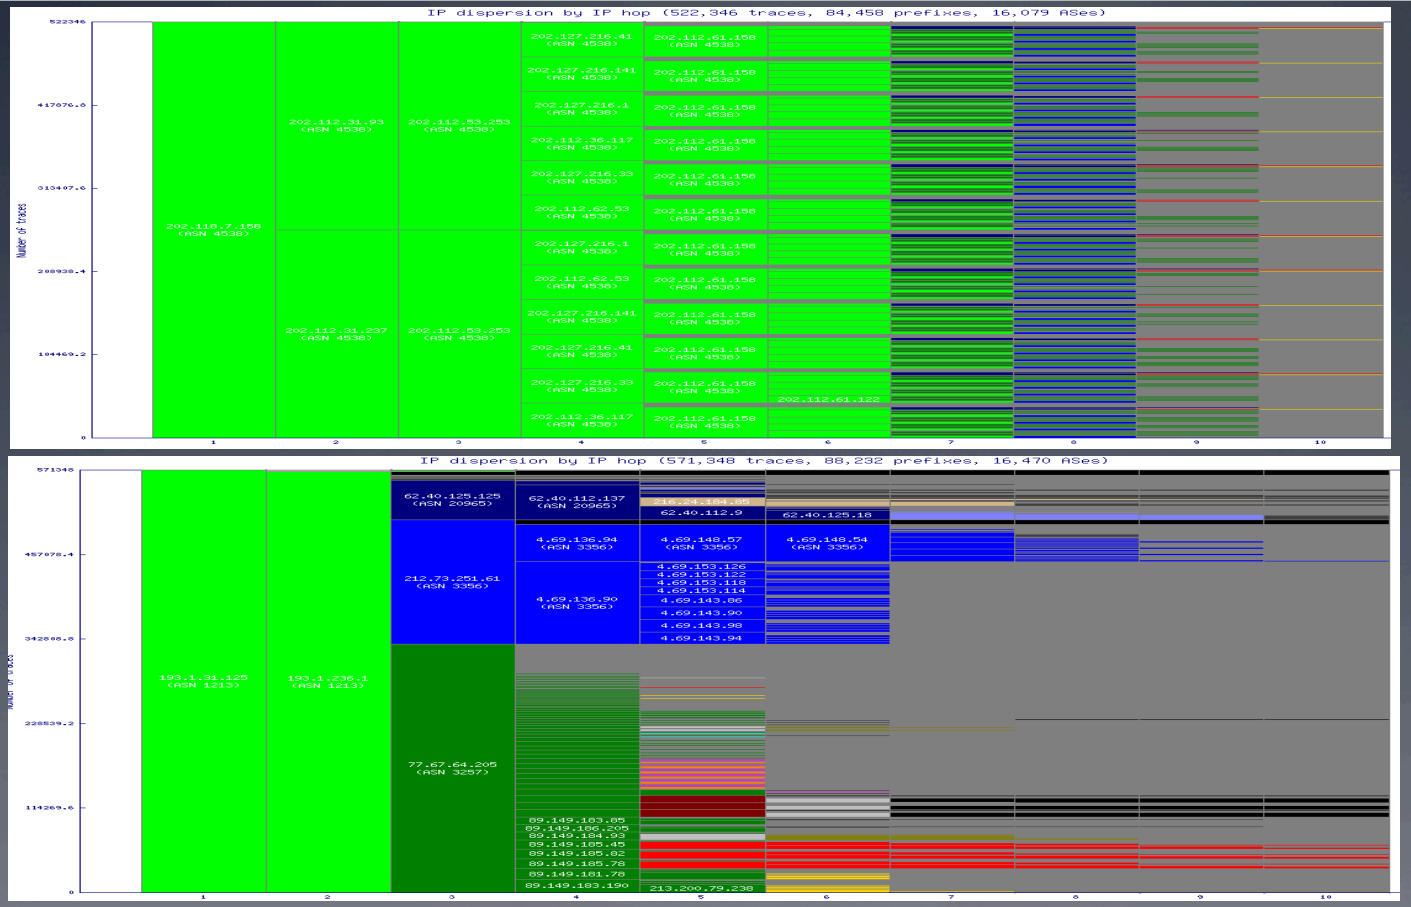

Chinese monitor (top): shows IP load balancing over many hops;
Irish monitor (bottom): shows fewer IP hops to other ASes.

Ark Topology Measurement

- Ark continuously gathers the largest set of IPv4 and IPv6 topology data made available to academic researchers and government agencies.
- From Sep 2007 through Feb 2012, we have collected more than 15.6 Billion traces (6.2 TB uncompressed, 1.9 TB compressed).

Topology Measurements over Time

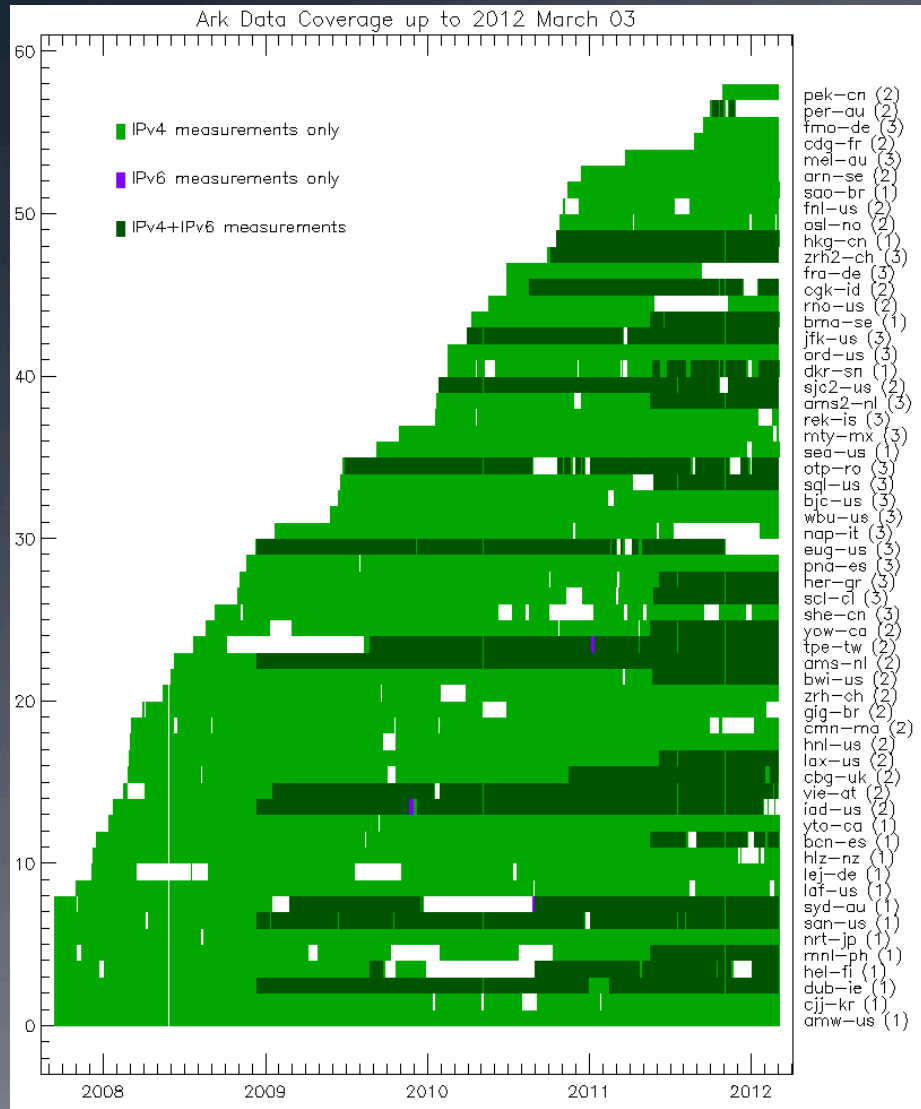

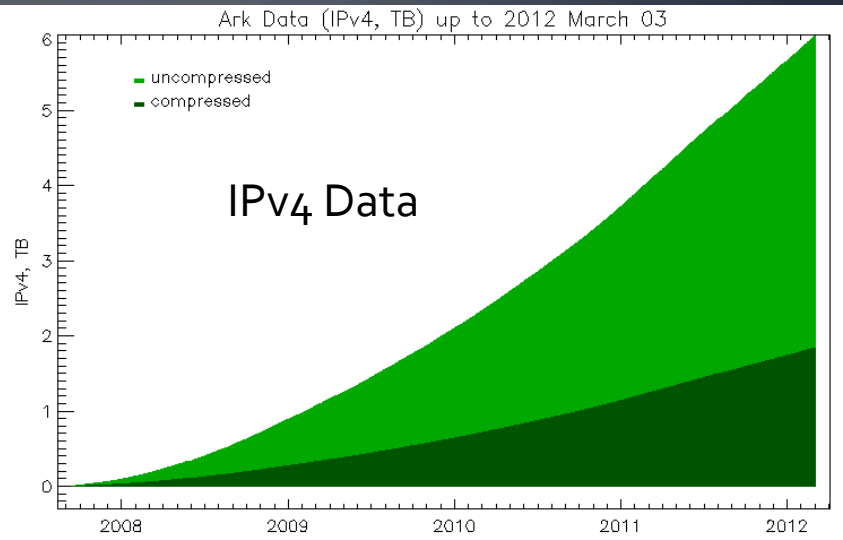

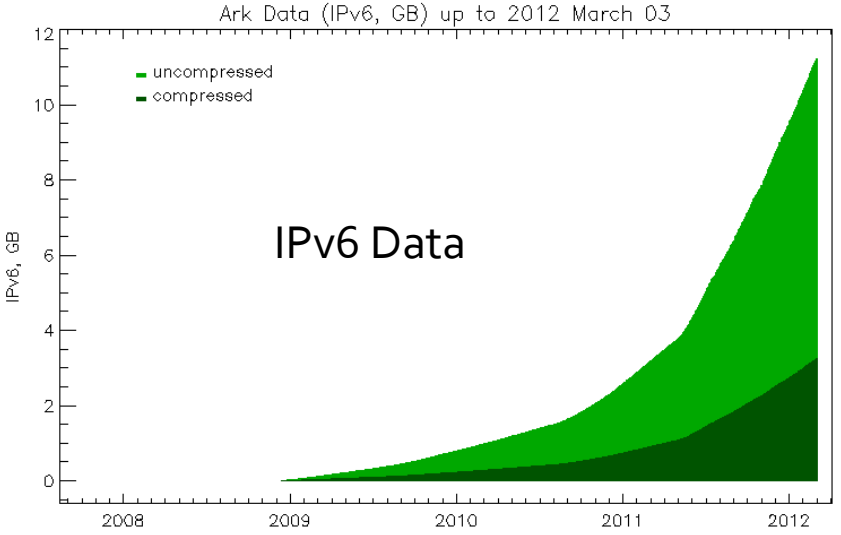

Topology Datasets

1. ***IPv4 Routed /24***: topology probes to each /24, continuously
2. ***IPv4 Routed /24 DNS Names***: DNS annotations, also capture raw DNS query/response traffic
3. ***IPv6 Topology***: topology probes to each routed IPv6 prefix
4. ***Internet Topology Data Kit (ITDK)***: curated IPv4 data
5. ***IPv4 Routed /24 AS Links***: AS adjacencies
6. ***AS Relationships***: inferred AS business relationships
7. ***AS Rank***: inferred AS ranking

http://www.caida.org/data/

IPv4 Routed /24 Topology

- ongoing large-scale topology measurements
- ICMP Paris traceroute to every routed /24 (9.7 million)
  - 57.9% of total IPv4 space (per Feb 2012 Route Views)
  - 7% increase since August 2010
  - probing rate = 100 probes per second
- running *scamper* probing tool
- dynamically assign measurements to teams of monitors
  - 3 teams active, 18-21 members/team
  - a cycle through every routed /24 takes 2-3 days
  - each /24 is probed once per cycle

IPv6 Topology

- ongoing large-scale topology measurements
- Ark monitors continuously probe BGP-announced prefixes /48 or shorter
  - 7,371 prefixes as of February 2012
- Each monitor probes a single random destination in each prefix using *scamper*

Internet Topology Data Kit Process

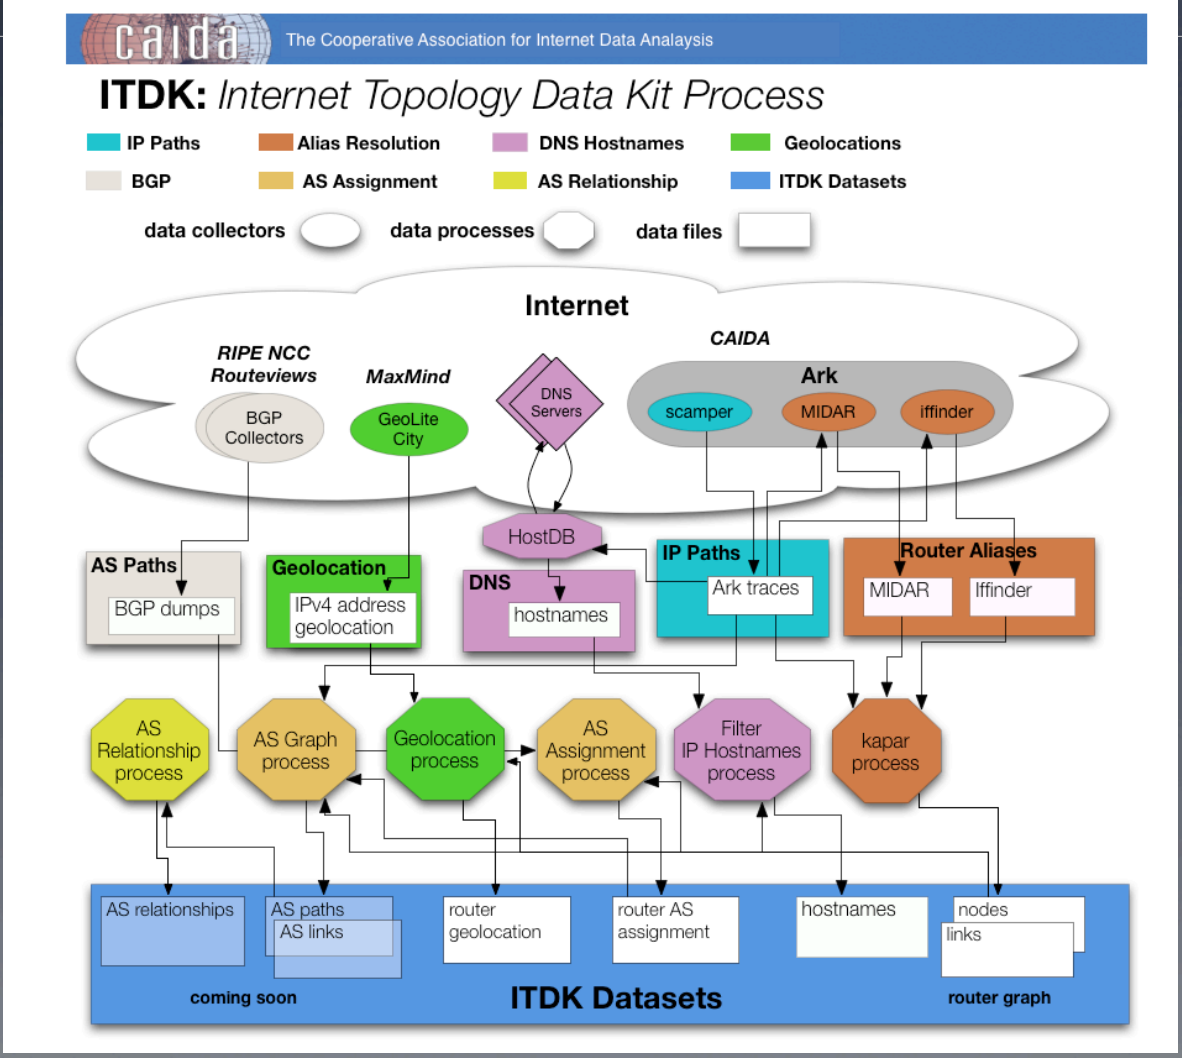

Internet Topology Data Kit (ITDK)

- Derived from two weeks of traceroute data probing IPv4 addresses
- Two router-level topologies
  1) Optimized for accuracy: *MIDAR+iffinder*
     highest confidence aliases with low false positives.
  2) Optimized for completeness: *MIDAR+iffinder+kapar*
     more alias coverage, more false positives (inflating routers)
- Data files: routers, links, router-to-AS mappings, geographic location of each router, DNS lookups of observed IP addresses.

Internet-scale Alias Resolution

- Goal: collapse observed interfaces into routers
- Earlier efforts at CAIDA: *iffinder*, *kapar* (APAR++)
- Most recent approach: *MIDAR* (inspired by RadarGun)
  - Two interfaces on same router respond in similar way
  - IP ID values in responses used as fingerprints to find aliases
    - IP ID is a 16-bit value in the IP header normally used for packet fragmentation and reassembly
    - Two interfaces on same router probed closely in time will return similar IP ID values: over time, similar time-series velocity.
- Architecture paper under peer review for TON 2012

Alias Resolution Misconceptions

// Unfortunately, faithfully mapping interface IP addresses to routers is a difficult open problem known as the IP alias resolution problem [51, 28], and despite continued research efforts (e.g., [48, 9]), it has remained a source of significant errors. While the generic problem is illustrated in Figure 2, its impact on inferring the (known) router-level topology of an actual network (i.e., Abilene/Internet2) is highlighted in Figure 3 -- the inability to solve the alias resolution problem renders in this case the inferred topology irrelevant and produces statistics (e.g., node degree distribution) that have little in common with their actual counterparts…In view of these key limitations of traceroute, it should be obvious that starting with the Pansiot and Grad data set, traceroute-based measurements cannot be taken at face value and are of no or little use for inferring the Internet's router-level topology. //

"Mathematics and the Internet: A Source of Enormous Confusion and Great Potential", http://www.ams.org/notices/200905/rtx090500586p.pdf

MIDAR Approach

- Monotonic ID-based Alias Resolution (MIDAR) is our extension of the RadarGun approach:
  - Monotonic Bounds Test: for two addresses to be aliases, their combined IP ID time series must be monotonic
  - Sliding window for scalable probing
  - 4 probing methods: TCP, UDP, ICMP, "indirect" (TTL expired)
  - Multiple monitors

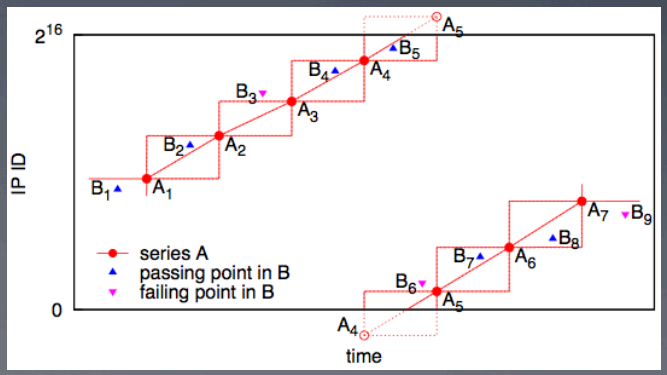

MIDAR Elimination Stage

- Potential alias set found in Discovery stage
- Testing pair-wise not scalable, necessary, or always possible.
- Instead probe subsets [colors in graph], such that most addresses belong to only 1 subset
- Probe a subset in parallel
- Efficiently covers all pairs
- Reduces chance of rate limiting

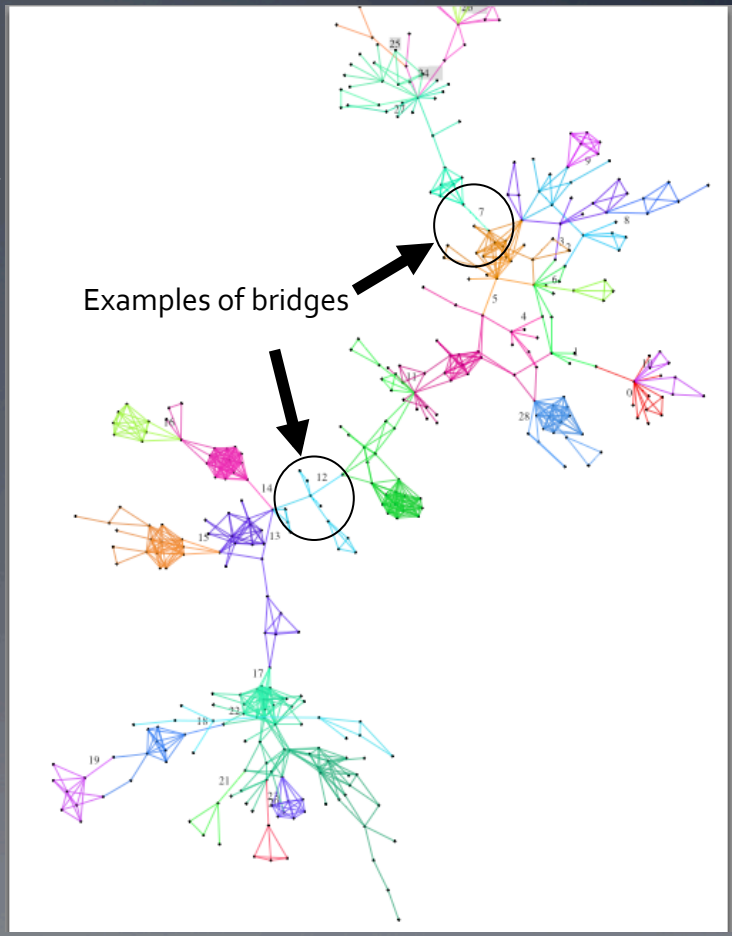

MIDAR Results

	2010-01	2010-04	2010-07	2011-04	2011-10
Input address Monotonic address Possible pairs	1.12 M 0.99 M 486 G	1.50 M 1.20 M 724 G	1.90 M 1.44 M 1038 G	2.32 M 1.87 M 1754 G	2.19 M 1.83 M 1676 G
Shared pairs after Discovery stage	1.63 M	4.00 M	5.49 M	6.83 M	7.00 M
Final Results •Shared pairs •Routers •Addresses on routers	 0.433 M 69 k 189 k	 1.36 M 108 k 383 k	 1.67 M 121 k 426 k	 2.49 M 125 k 413 k	 2.68 M 118 k 403 k

- We have continually improved MIDAR over time:
  - increasing input size of the graph; and
  - improving accuracy and effectiveness of methods.

AS Rank

- based on inferred economics of AS business relationships
- uses data from global routing tables
- orders by "customer cone": number of IP prefixes advertised by each AS, by its customer ASes, by their customer ASes, and so on

http://as-rank.caida.org/

AS Rank: screen shot

AS Ranking | Information for a single AS | Background | Data Sources | Help

Change data set: 2011.01.16 Change dataset

AS Ranking Help

The top AS Rankings are displayed below.

For information on a single specific AS, enter its number or name:

Find this AS

Table **shows** 10 of 36878 ASes, **sorted by** number of ASes in customer cone update view

rank	AS number	AS name	customer cone: Number of ASes	Number of IPv4 prefixes	Number of IPv4 addresses	Percentages of all ASes	Percentages of all IPv4 prefixes	Percentages of all IPv4 addresses	AS degree
1	3356	LEVEL3 Level 3 Commu	35,753	347,759	2,369,727,713	96%	97%	97%	2904
2	6939	Hurricane Electric,	33,621	326,502	2,223,730,966	91%	91%	91%	1959
3	3549	Global Crossing Ltd.	33,427	325,647	2,221,842,470	90%	91%	91%	1496
4	6461	Metromedia Fiber Net	30,524	297,848	2,080,222,108	82%	83%	85%	841
5	3257	Tinet SpA	29,989	293,640	2,032,322,896	81%	82%	83%	886
6	1239	Sprint	28,636	287,039	2,010,397,654	77%	80%	82%	1183
7	2914	NTT America, Inc.	28,501	284,683	1,974,794,207	77%	79%	81%	718
8	174	Cogent/PSI	27,722	273,242	1,794,711,050	75%	76%	73%	2972
9	1299	TeliaNet Global Netw	27,573	273,081	1,818,208,126	74%	76%	74%	630
10	7018	AT&T Services, Inc.	27,375	289,124	1,970,281,363	74%	80%	81%	2365

data sources

country	BGP	2011.01.16, 2011.01.17, 2011.01.18, 2011.01.19, 2011.01.20	routeviews eqix, isc, kixp, linx, routeviews2
	database	2011.03.30	netacuity
	delegation	2011.03.23	AFRINIC, APNIC, ARIN, IANA, LACNIC, RIPENCC
	whois	2010.08.31	AFRINIC, APNIC, ARIN, LACNIC, RIPE
name	autnum.txt	2011.03.23	www.potaroo.net
	whois	2010.08.31	AFRINIC, APNIC, ARIN, LACNIC, RIPE
topology	BGP	2011.01.16, 2011.01.17, 2011.01.18, 2011.01.19, 2011.01.20	ripe rrc00, rrc01, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc11, rrc12, rrc13, rrc14, rrc15, rrc16
			routeviews eqix, isc, kixp, linx, routeviews2

AS Rank (cont)

- Tabular views of inferred ISP info, rank, degree, customer cone size, customers, peers, and providers.

ISP info →

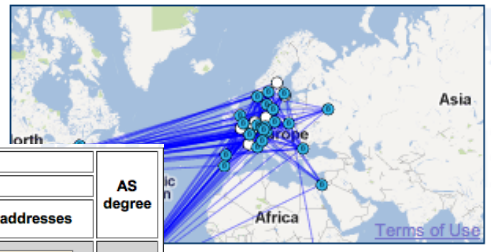

AS number:	1299
AS name:	TeliaNet Global Network
rank:	9
customer cone size:	27573
degree:	630

rank	AS number	AS name	customer cone						AS degree
			Number of			Percentages of all			
			ASes	IPv4 prefixes	IPv4 addresses	ASes	IPv4 prefixes	IPv4 addresses	
4	6461	Metromedia Fiber Net	30,524	297,848	2,080,222,108	82%	83%	85%	841
5	3257	Tinet SpA	29,989	293,640	2,032,322,896	81%	82%	83%	886
6	1239	Sprint	28,636	287,039	2,010,397,654	77%	80%	82%	1183
7	2914	NTT America, Inc.	28,501	284,683	1,974,794,207	77%	79%	81%	718
8	174	Cogent/PSI	27,722	273,242	1,794,711,050	75%	76%	73%	2972
9	1299	TeliaNet Global Netw	27,573	273,081	1,818,208,126	74%	76%	74%	630
10	7018	AT&T Services, Inc.	27,375	289,124	1,970,281,363	74%	80%	81%	2365
11	3320	Deutsche Telekom AG	27,114	274,055	1,854,661,572	73%	76%	76%	535
12	6453	TATA Communications	26,018	268,534	1,776,070,663	70%	75%	73%	569
13	701	MCI Communications S	25,632	252,604	1,702,064,736	69%	70%	70%	1946

← Ranking

rank	neighbor AS	neighbor name	type
3	3549	Global Crossing Ltd.	↔ peer
4	6461	Metromedia Fiber Net	↑ provider
5	3257	Tinet SpA	↑ provider
6	1239	Sprint	↔ peer
7	2914	NTT America, Inc.	↔ peer
8	174	Cogent/PSI	↔ peer
10	7018	AT&T Services, Inc.	↔ peer
11	3320	Deutsche Telekom AG	↔ peer
12	6453	TATA Communications	↔ peer
13	701	MCI Communications S	↔ peer

Customers, providers, and peers →

AS Rank Validation

- Interface to provide corrections to relationships

rank	neighbor AS	neighbor name	type	correction
3	3549	Global Crossing Ltd.	↔ peer	provider
4	6461	Metromedia Fiber Net	↑ provider	
5	3257	Tinet SpA	↑ provider	peer
6	1239	Sprint	↔ peer	
7	2914	NTT America, Inc.	↔ peer	
8	174	Cogent/PSI	↔ peer	
10	7018	AT&T Services, Inc.	↔ peer	
11	3320	Deutsche Telekom AG	↔ peer	
12	6453	TATA Communications	↔ peer	
13	701	MCI Communications S	↔ peer	

Disclaimer: We show these corrections as examples of the interface not as actual corrections received by TeliaNet Global Network.

Geolocation Tools Comparison

- Service evaluation criteria
  - What geographic granularity does it provide?
    - Continent, country, state/prefecture, city, zip code
  - What Internet identifier granularity does it support?
    - IP address, network prefix, Autonomous System (AS)
  - Does accuracy vary by region or type of network?
- We evaluated: Digital Envoy's Netacuity, MaxMind (Free and commercial), IP2Location, Ipligence, and HostIP.info. Quova and Akamai remain unwilling to participate.
- Results generally agreed on IP-address-to-country mappings
  - MaxMind Lite and GeoIP had the highest level of agreement (99.1%)
  - IPligence had the lowest level (94.3%)
  - Finer granularity harder to evaluate
  - Netacuity and MaxMind GeoIP performed "best" in our testing

Integrated Visualization of Topological Connectivity

- AS –level graph

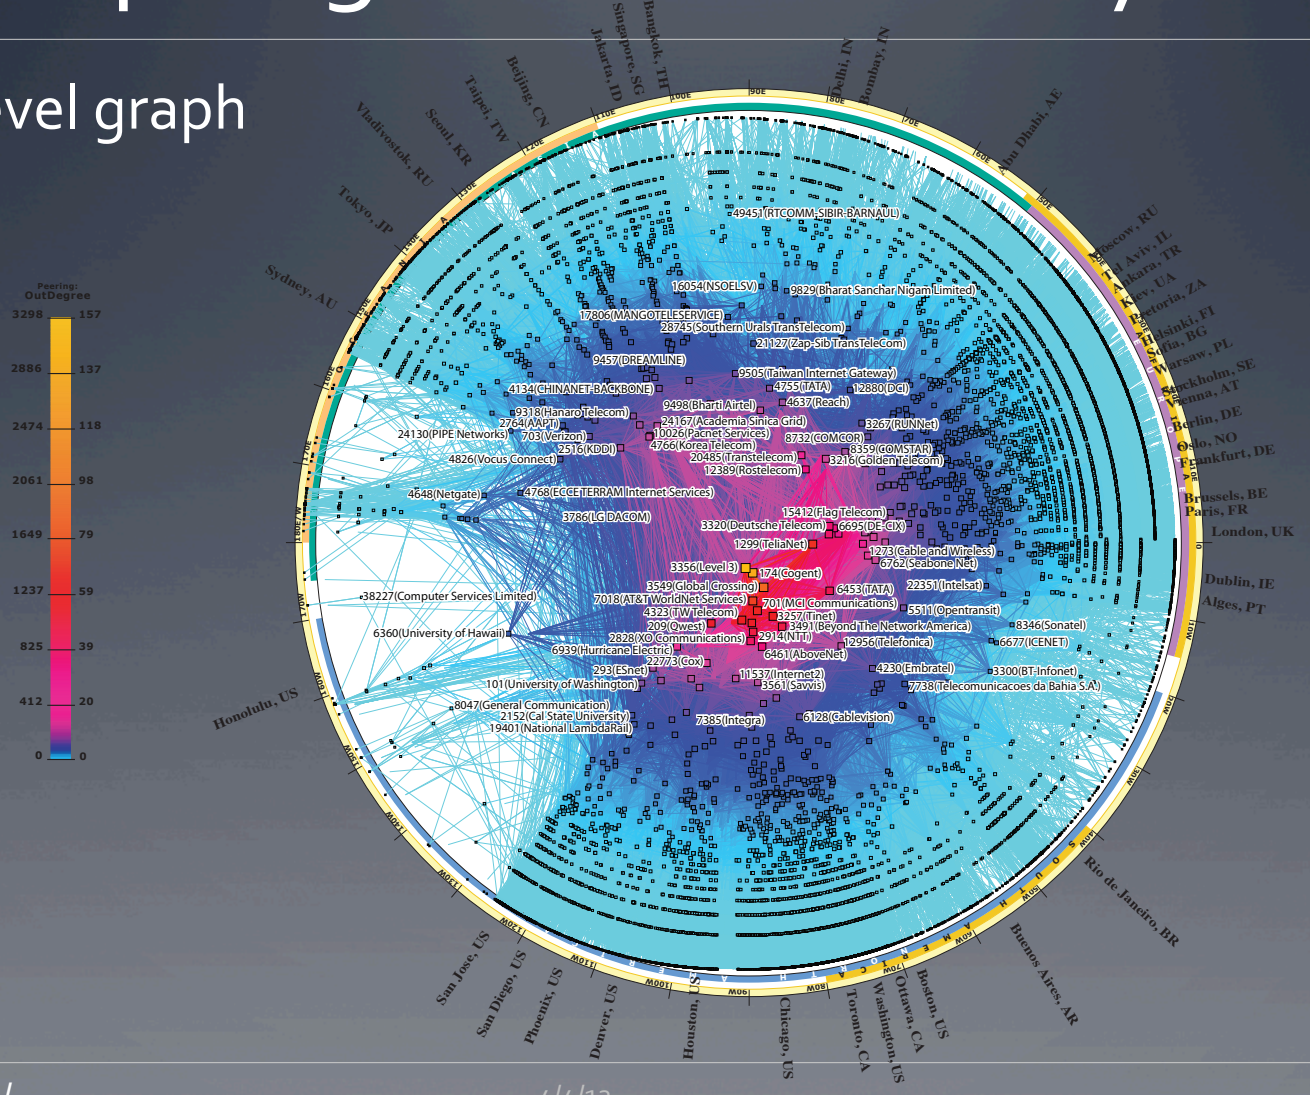

Integrated Visualization of Topological Connectivity

- Router-level graph (prototype)

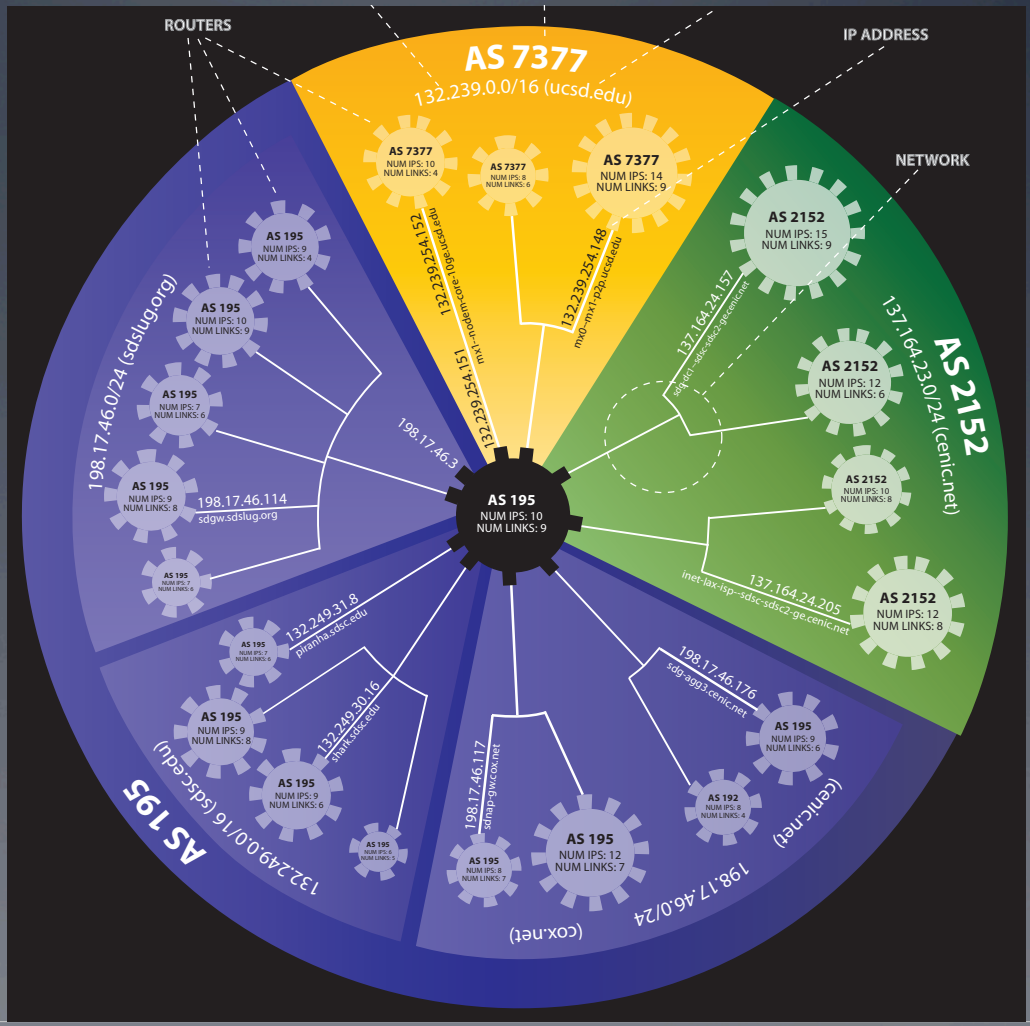

Internet Topology Data Comparison

- Topology maps needed to analyze or model Internet structure
  - many studies use single, inconsistent, incomplete, or undocumented sources which can undermine integrity of analysis results
  - objective: enabling informed selection of topology datasets
- Approach: systematic comparison of best available data sources
  - characterizing the Internet topology at three granularities: IP address (interface), router, Autonomous System (AS)
  - most comprehensive study: sources, metrics, methods, results
    URL: http://www.caida.org/research/topology/topo_comparison/

Published Experiments Using Ark

1) "Traceroute Probe Method and Forward IP Path Inference", IMC'08.
2) "Understanding the efficacy of deployed internet source address validation filtering", IMC'09.
3) "Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers", PAM 2010.
4) "The ISMA 2010 AIMS-2 Workshop on Active Internet Measurement Report", ACM SIGCOMM Computer Communication Review (CCR), Sep 2010.
5) "Measured impact of crooked traceroute", CCR, Jan 2011.
6) "The ISMA 2011 AIMS-3 Workshop on Active Internet Measurement Report", ACM SIGCOMM Computer Communication Review (CCR), July 2011.

Published Experiments Using Ark

7) "Geocompare: a comparison of public and commercial geolocation databases", Network Mapping and Measurement Conference, May 2011.
8) "Twelve Years in the Evolution of the Internet Ecosystem", IEEE/ACM Transactions on Networking, Sep 2011.
9) "Analysis of Country-wide Internet Outages Caused by Censorship", IMC Nov 2011.
10) "Efficient Internet Topology Discovery Techniques", Masters Thesis, U. Waikato, Alistair King, 2010.
11) "Sustaining the Internet with Hyperbolic Mapping", Nature Communications, Oct 2010.
12) "Hyperbolic Geometry of Complex Networks", Physical Review E, Oct 2010.

- **Another 107 articles in Google scholar cite or use data from Ark as of 02 sept 2011.

Scheduled, Planned Activities

- Deploy 1-2 monitors/month to measure IPv4 and IPv6 topology
- Continue to release and refine ITDK
- Publish alias resolution study and release three versions of code
- Annotated router-level graph visualization and database support
- Topology on demand measurements
- AIMS 2012 Workshop report -> CCR
- Explore coupling of data plane and (DHS-funded BGPmon) control plane data
- AS Rank documentation, validation and improved algorithms and interface
- Develop a web-based interface to topo-on-demand service

BAA Number: Cyber Security BAA 07-09 Title: Science and Technology of Internet Topology Mapping	Offeror Name: Kimberly Claffy Date: 06/26/07
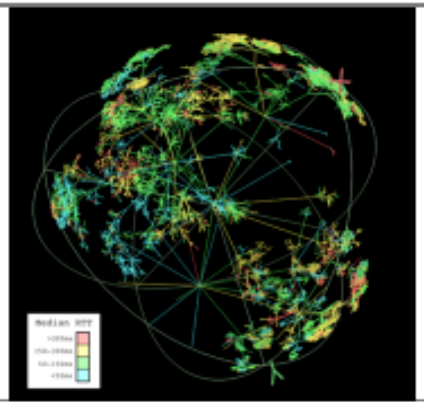 Walrus visualizations of round-trip time measurements made by CAIDA's macroscopic Internet topology monitor located in Herndon, VA, USA.	Internet Topology Mapping: 1. Operational infrastructure to support continuous Internet topology mapping. 2. Periodic active probing of 100% of BGP prefixes announced 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