Leveraging the Science and Technology of Internet Mapping for Homeland Security

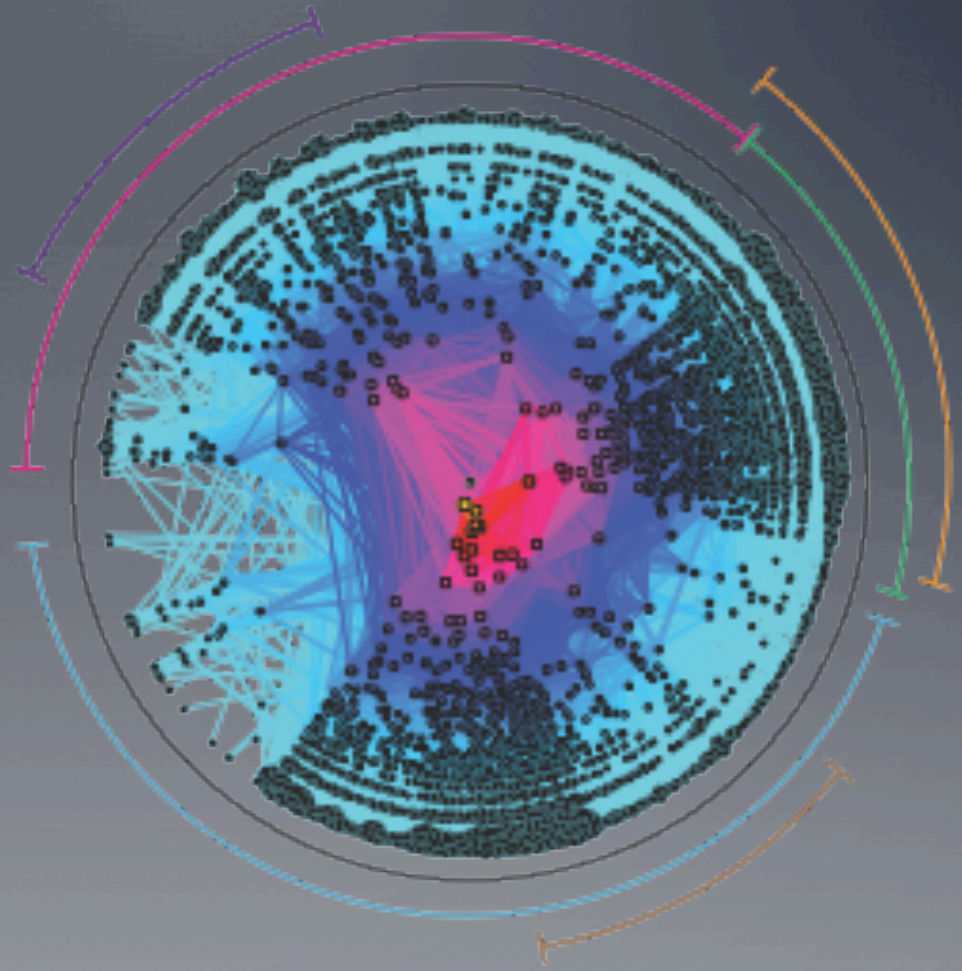

Young Hyun, Ken Keys, Amogh Dhamdhere, Bradley Huffaker, Josh Polterock, Marina Fomenkov, Dima Krioukov, Matthew Luckie, and kc claffy

CAIDA/UCSD
DHS S&T
N66001-08-C-2029
9 Oct 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 six strategic Internet 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/

Technical Transfer Approach

- Integrated strategic measurement & analysis capabilities:
  1. Ark Measurement platform: software, data, access
  2. Topology analysis: software, data kits, papers
  3. Dual router- and AS-level graphs: software, viz
  4. AS taxonomy and relationships: published algorithms, interactive web service (AS Rank)
  5. Geolocation of IP resources: comparison report
  6. Graph visualization: part of AS Rank web service

[all software GPL or UCSD license (no patents); UCSD supports commercial license.]

Benefits to DHS S&T

- Improve critical national capabilities:
  - situational awareness for homeland cybersecurity 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

Archipelago (Ark)

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

Ark monitor locations

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.

Data from Infrastructure

- IPv4 topology data
  - 10.1 TB data served by PREDICT, data.caida.org
  - Sep 2007 to June 2012 (58 months)
    - 17 B traceroutes; 1850+K cycles
  - Per month:~431M traceroutes; ~175 GB/month
  - Key input to, e.g., AS links and alias resolution
  - Each team collects traces from 10.1 million /24s
- IPv6 topology data
- Supporting software: mper, Marinda, MIDAR, kapar

Archipelago Monitor

- 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

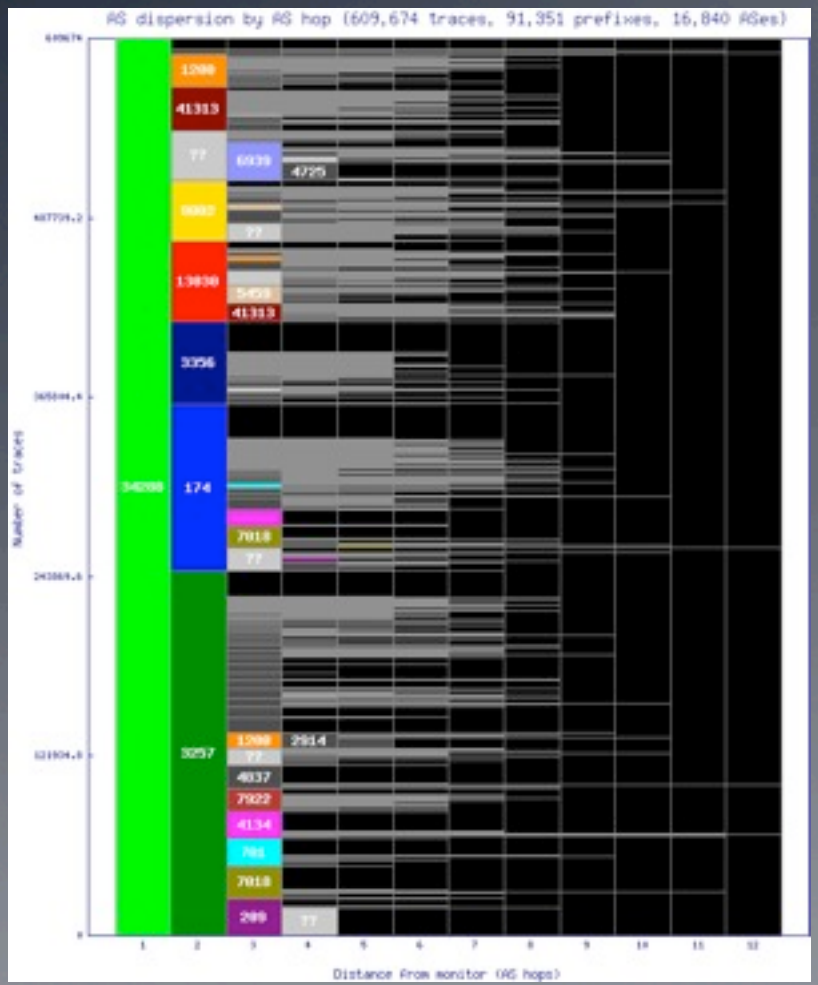

Kantonsschule Zug (zrh2-ch)

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.

http://www.caida.org/

AS Dispersion by IP Hop

Kantonsschule Zug (zrh2-ch)

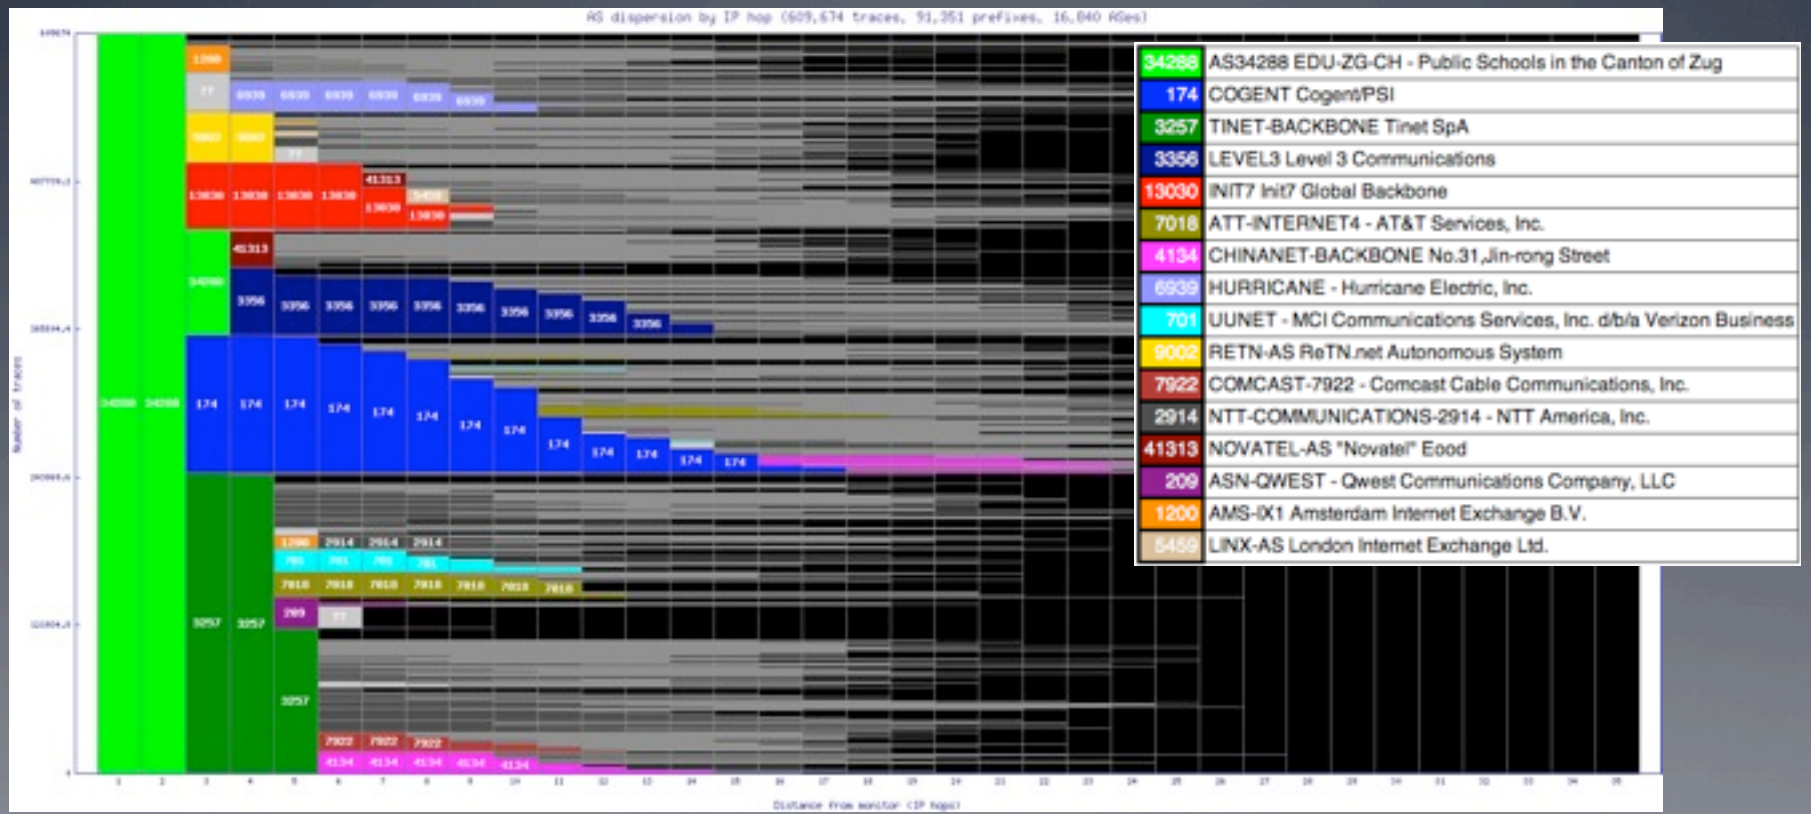

AS Dispersion by IP Hop: shows load balancing

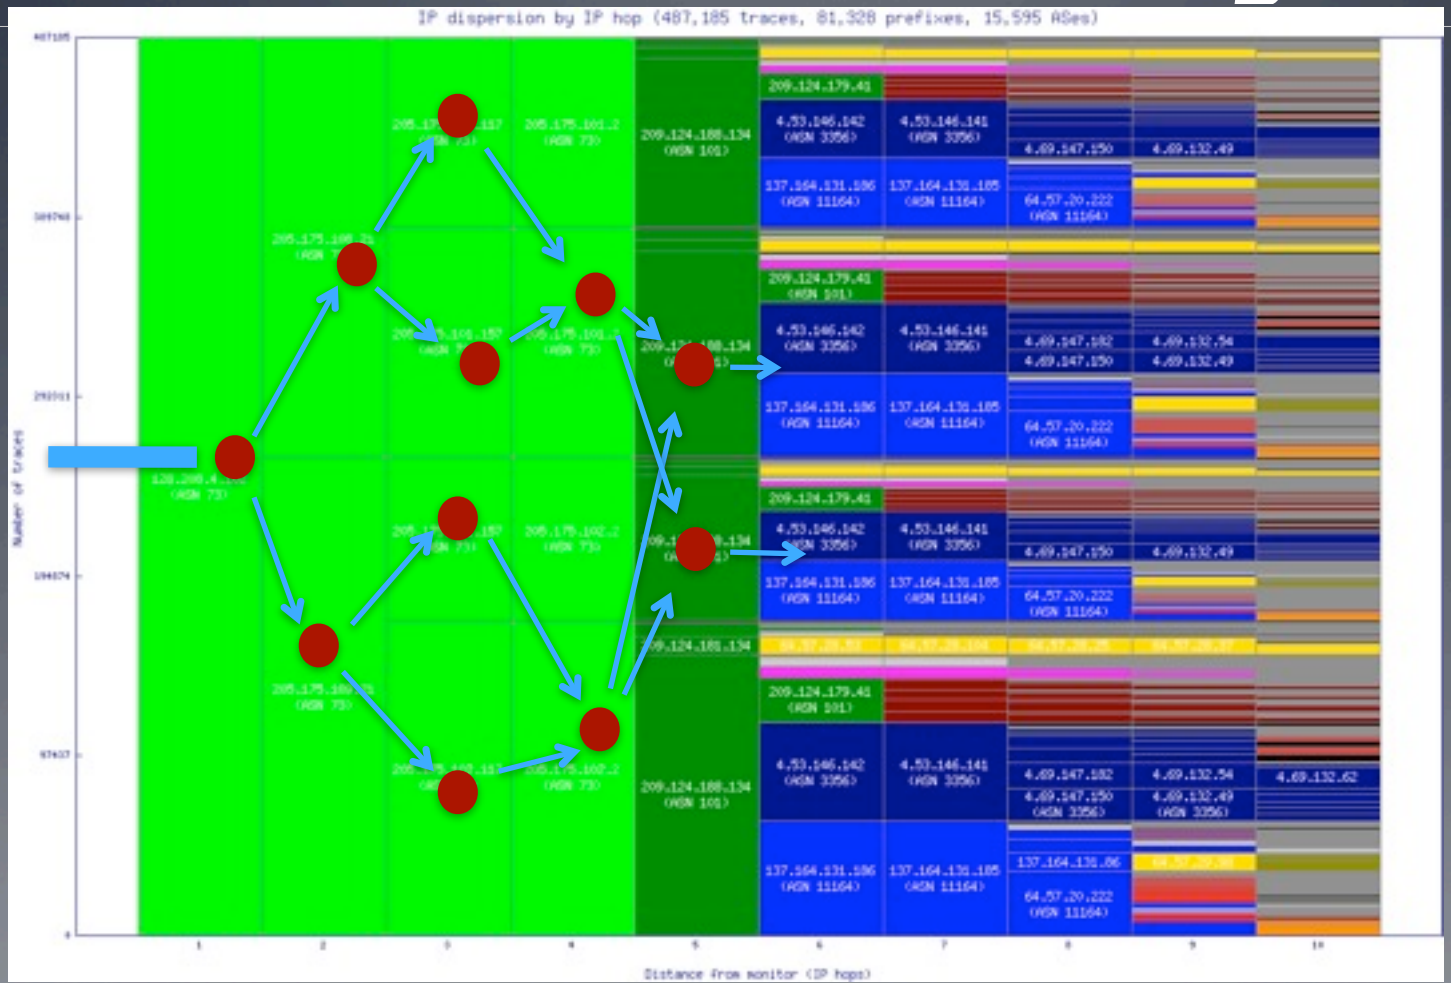

Median RTT to Destination Countries

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

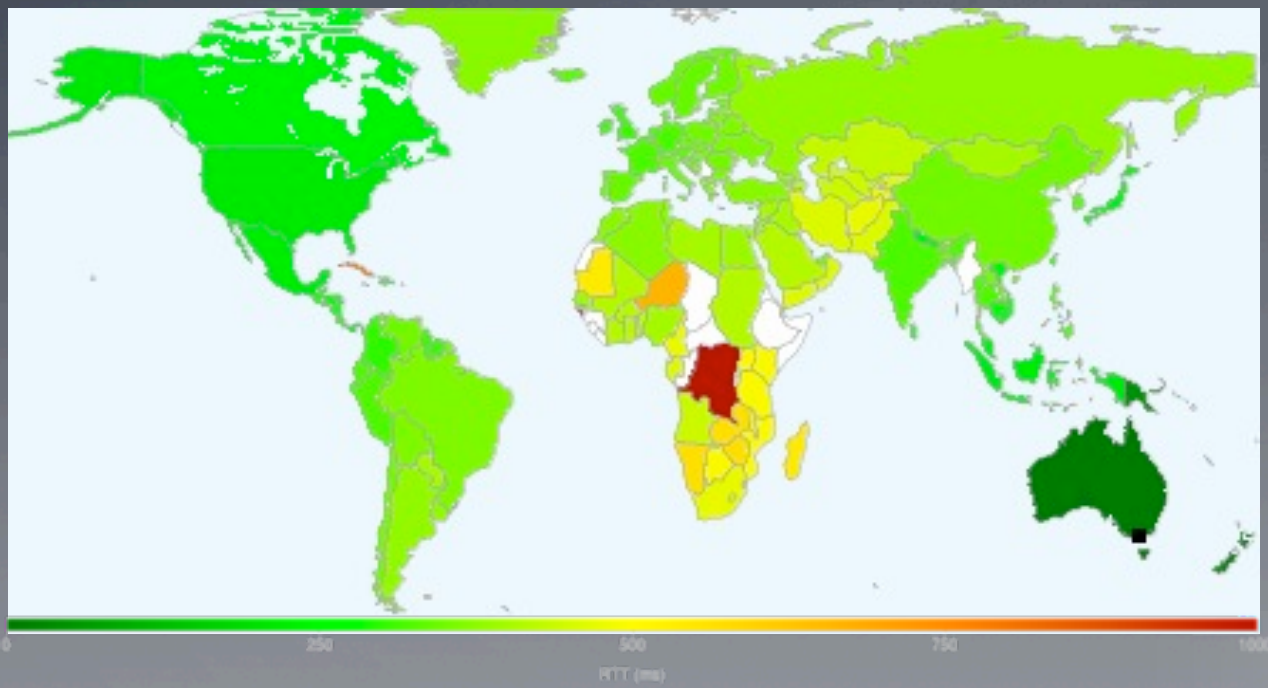

University of Melbourne
(mel-au)

http://www.caida.org/

Median RTT to Destination Countries

Northeastern University (she-cn)

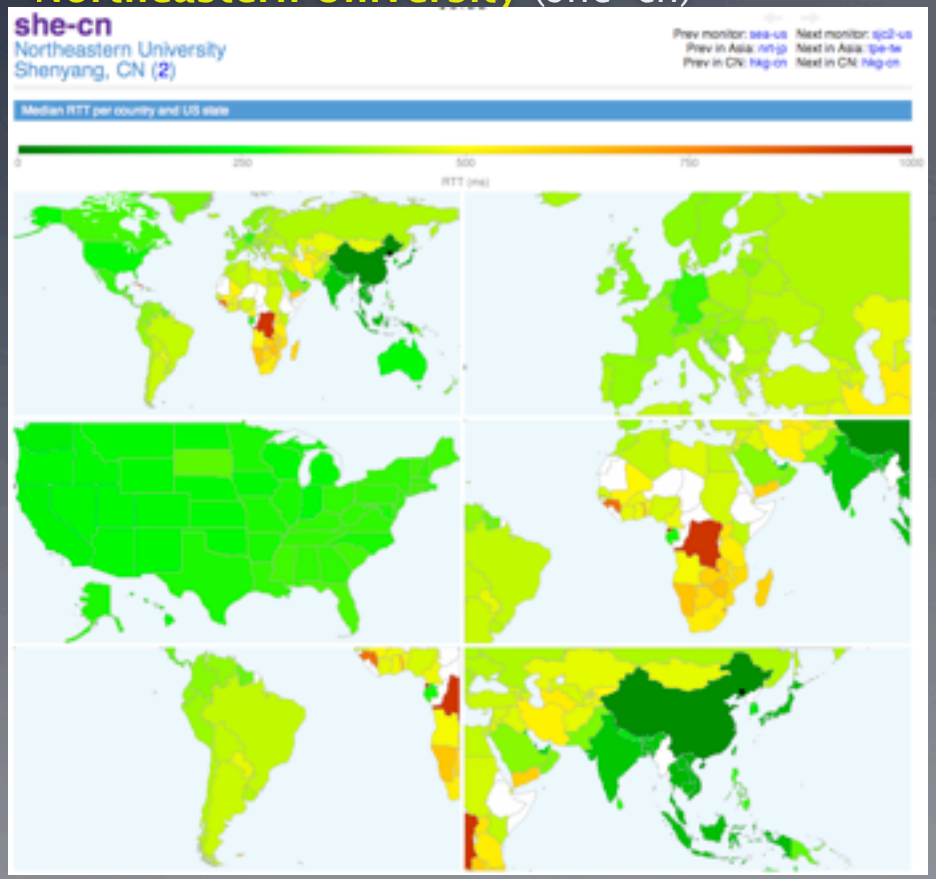

AMS-IX Amsterdam, NL (ams2-nl)

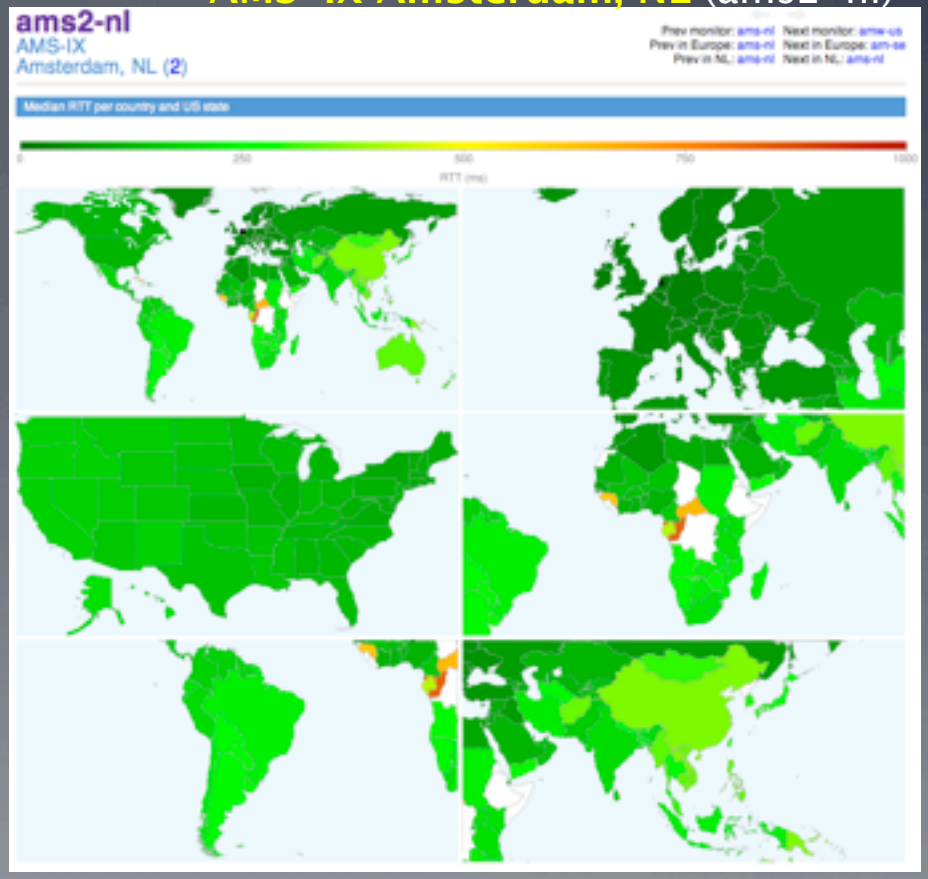

http://www.caida.org/

Median RTT to Destination Countries

Sept 2010. Prior to new west coast Africa fiber

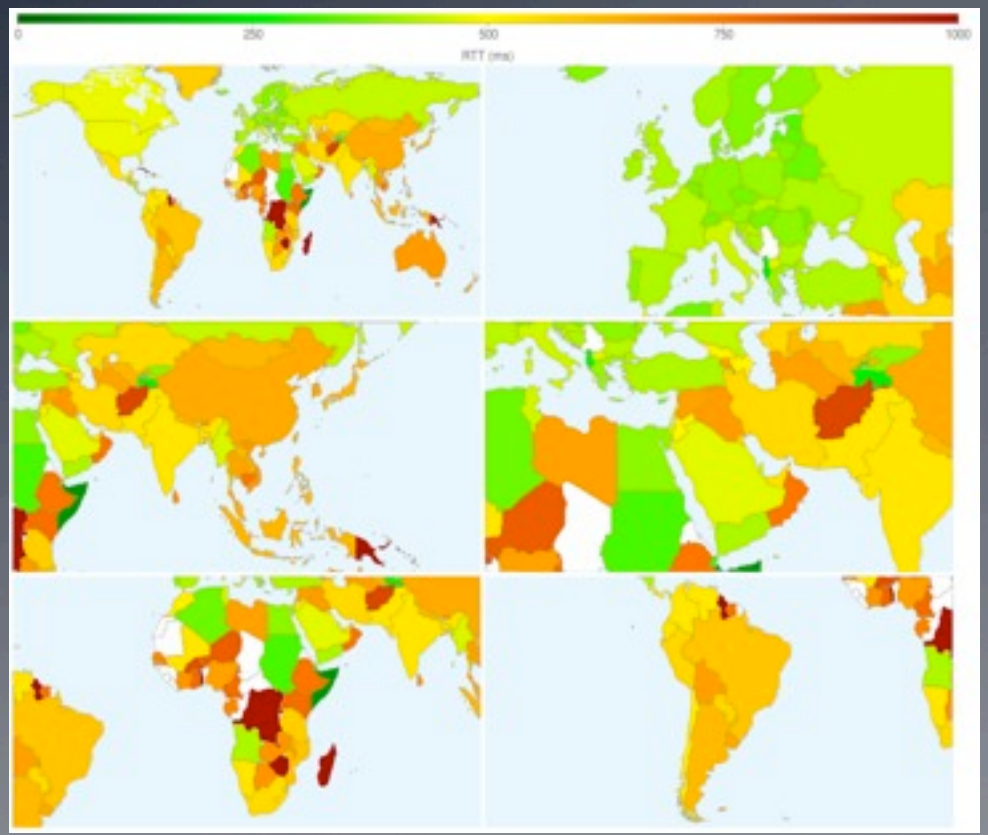

Oct 2011. After new fiber deployed

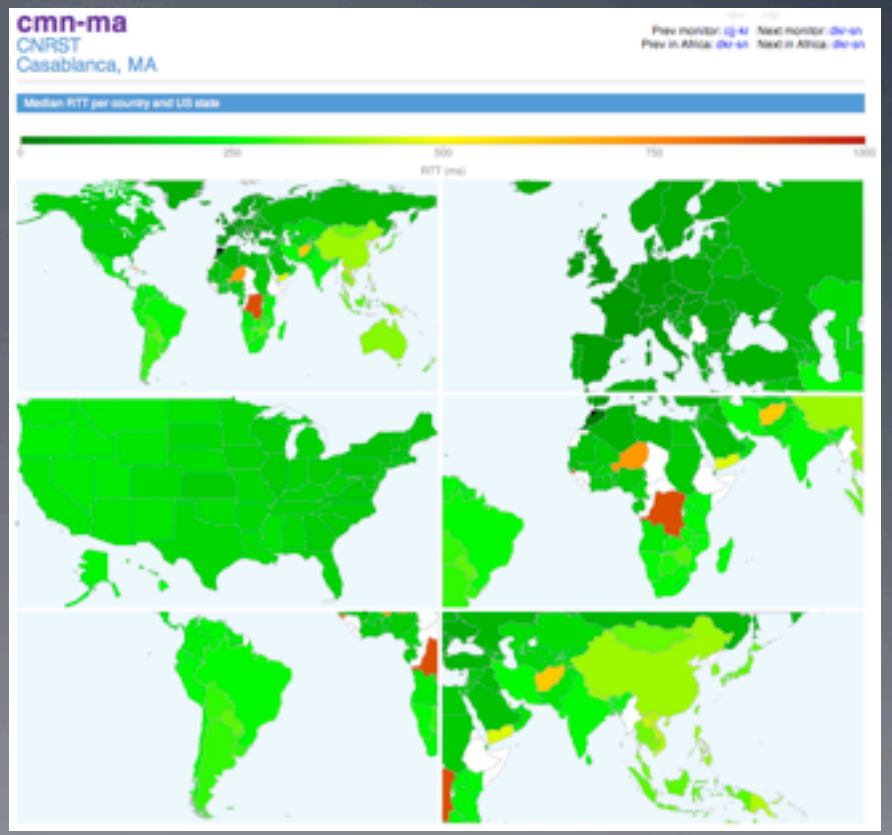

CNRST Casablanca, Morocco (cmn-ma)

http://www.caida.org/

IP Path Dispersion (by IP Hop)

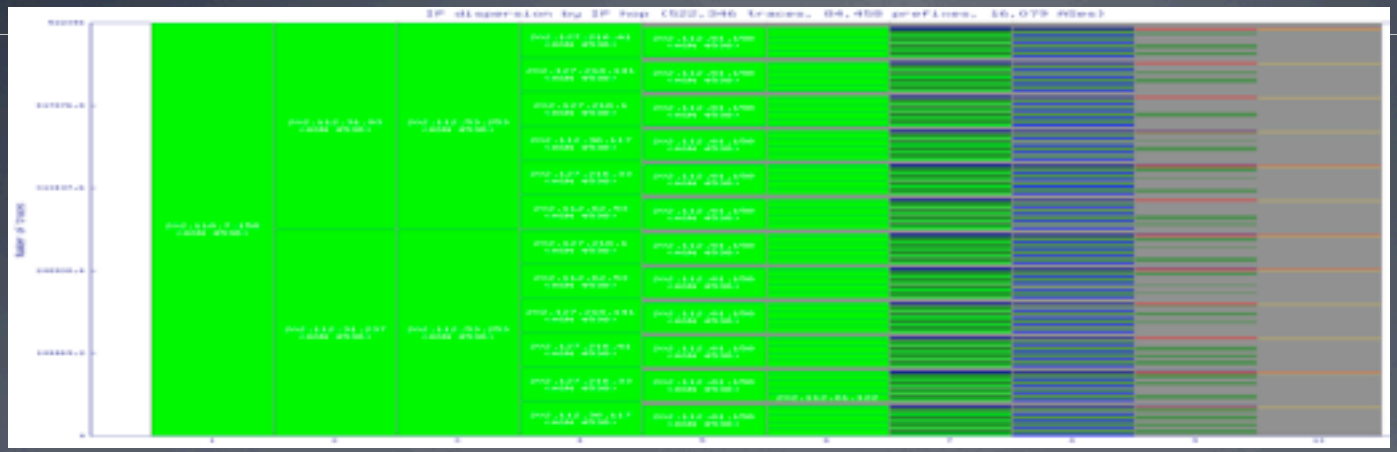

Chinese monitor: shows IP load balancing over many hops.

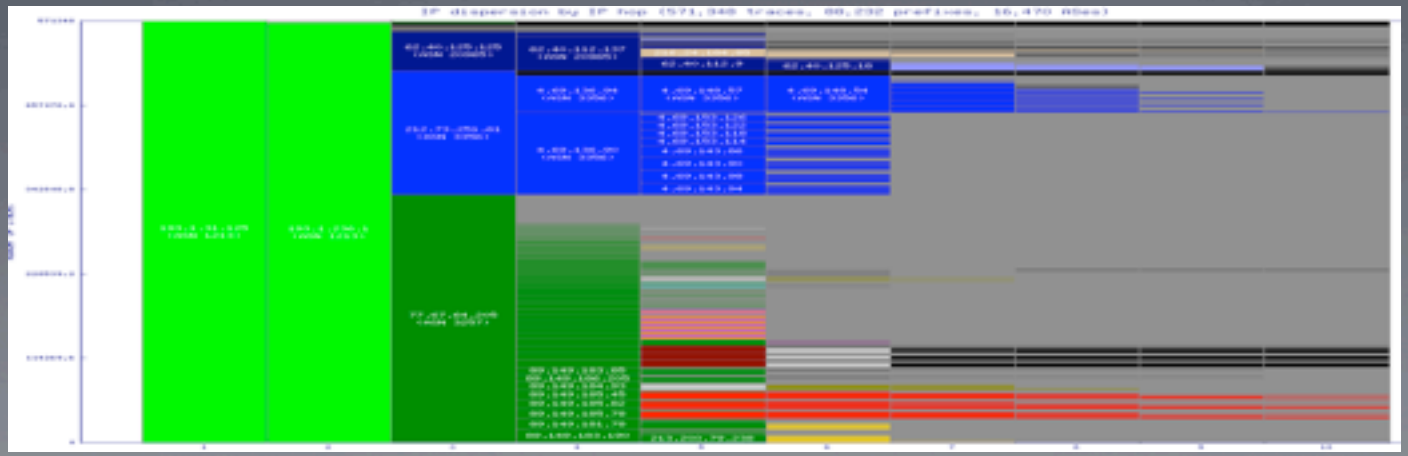

Irish monitor: shows fewer IP hops to other ASes.

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 June 2012, we have collected more than 17 Billion traces (6.7 TB uncompressed, 2.1 TB compressed).

Topology Measurement History

- 3.6 TB (skitter+ark compressed)
  - routed IPv4: 6.7TB (unc.) since Sep 2007
  - routed IPv6: 15.4GB since Dec 2008

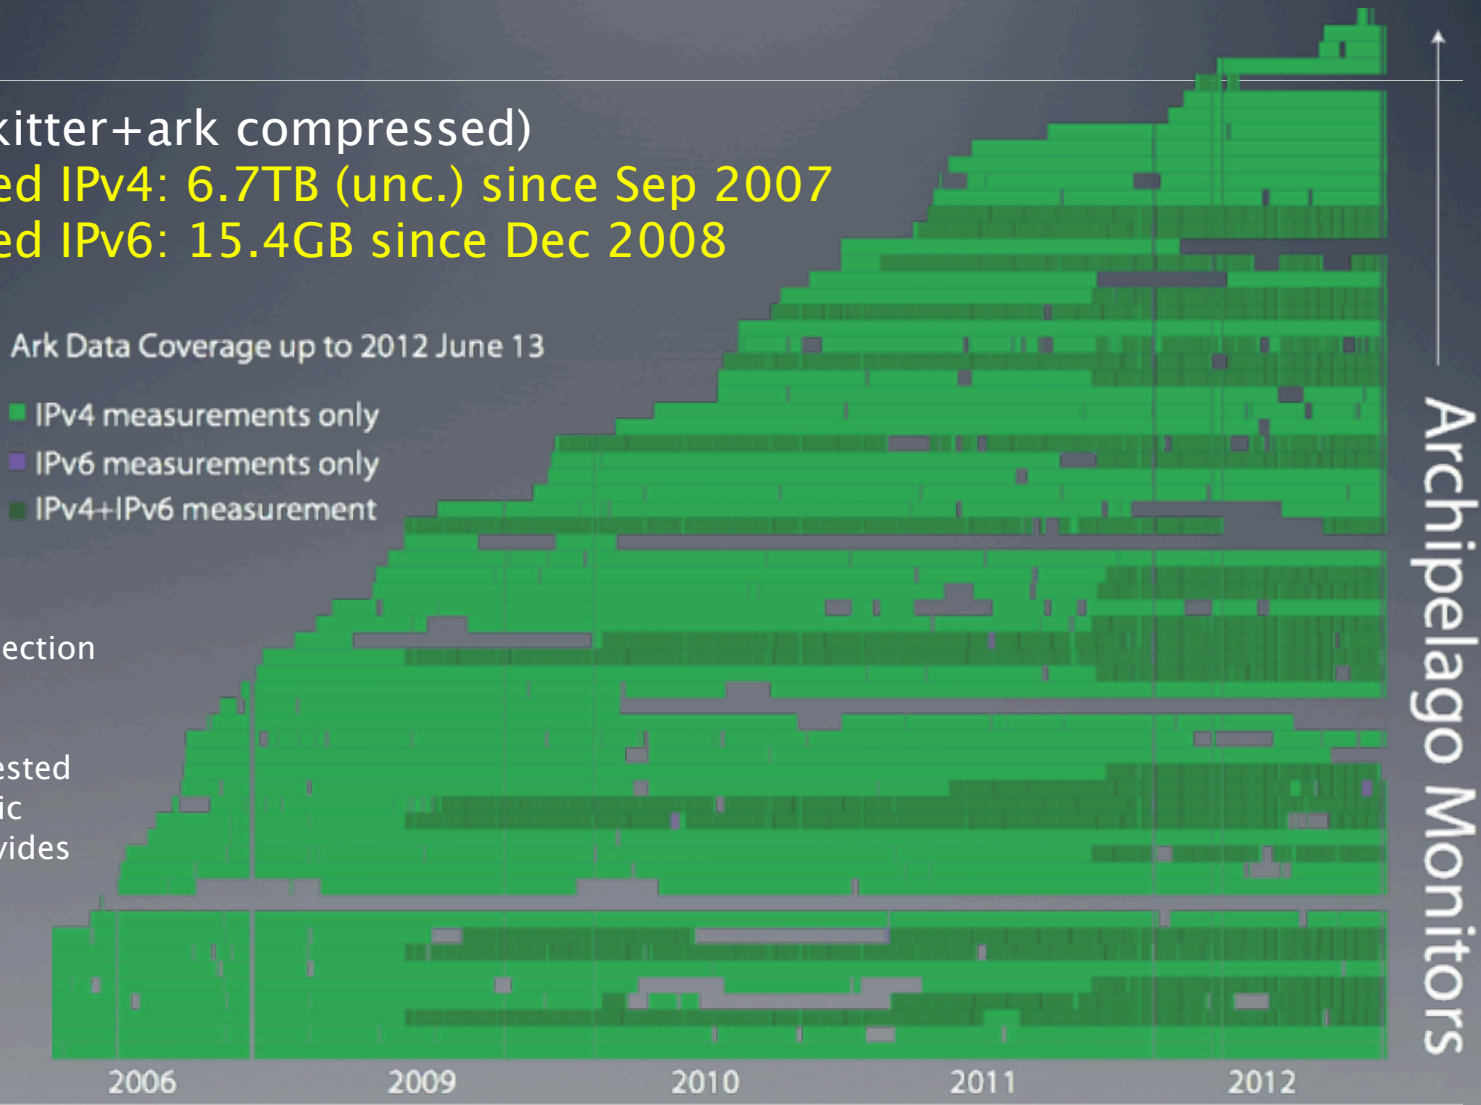

Raw traces area a collection of IP paths.

For researchers interested in a single microscopic snapshot, CAIDA provides the ITDK

http://www.caida.org/

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

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

IPv4 Routed /24 Topology

- ongoing large-scale topology measurements
- ICMP Paris traceroute to every routed /24 (10.1M)
  - ~60% of total IPv4 space (per Aug 2012 Route Views)
  - 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
  - 10,269 routed prefixes as of Sept 2012
- Each monitor probes a single random destination in each prefix using scamper

Internet Topology Data Kit Process

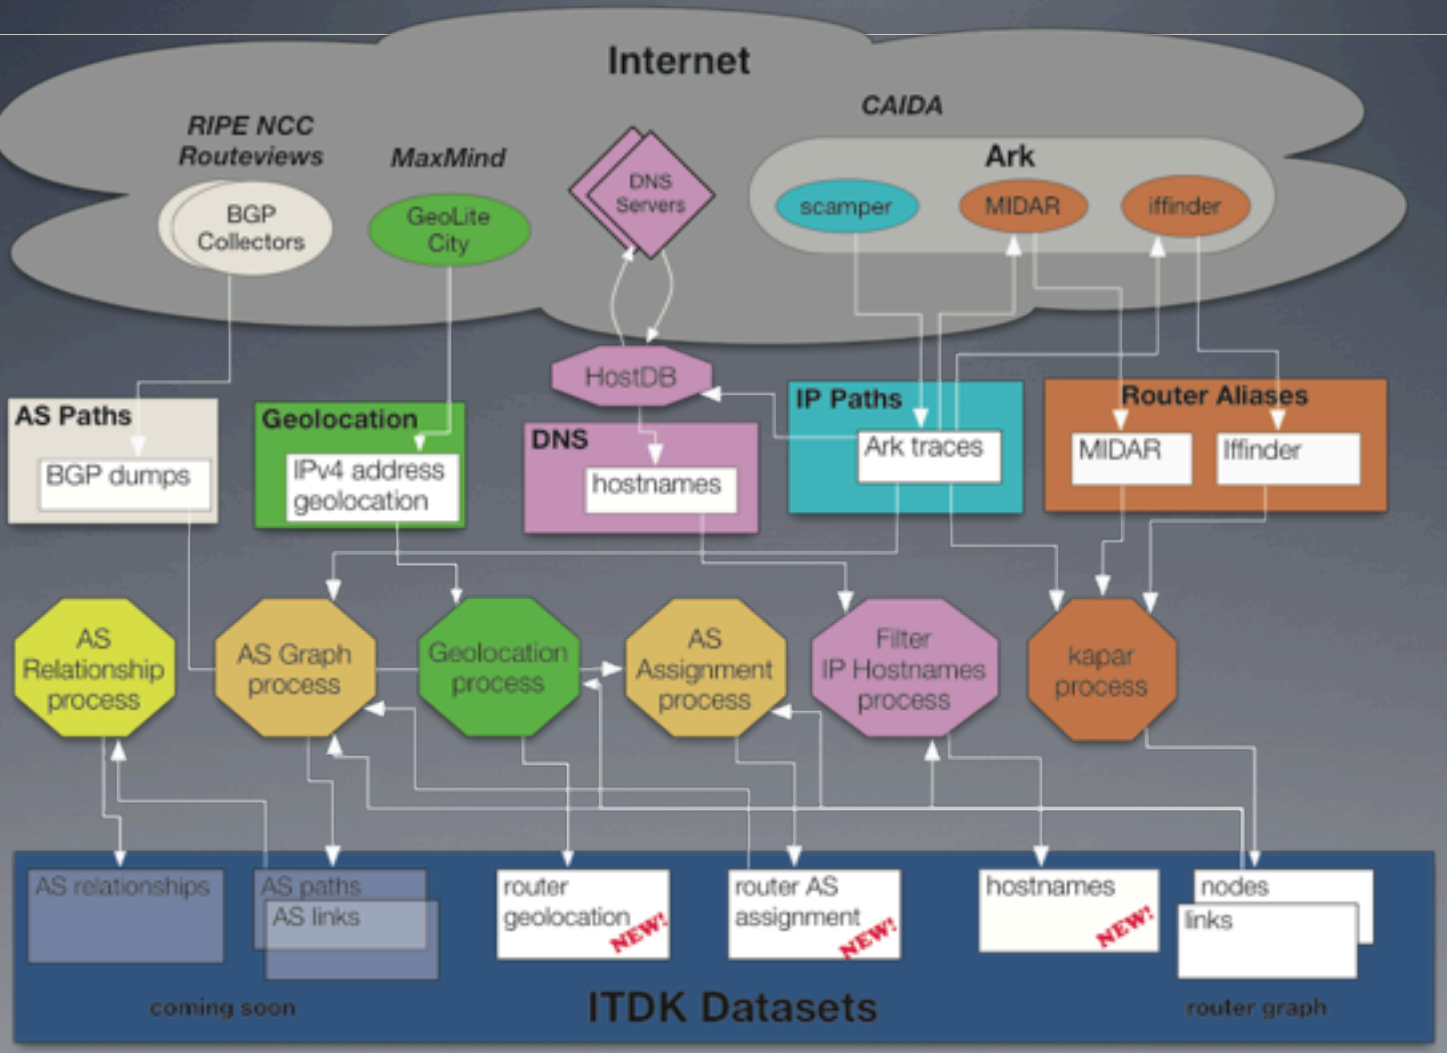

http://www.caida.org/

Internet Topology Data Kit

- Derived from two weeks of traceroute data probing IPv4 addresses.
- Last ITDK just posted on www.caida.org (July 2012)
- 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, false positives (inflating routers)
- Data files: routers, links, router-to-AS mappings, router geolocations, DNS lookups of IP addresses

Insights Enabled

- Probing technique performance comparison (w/ .NZ)
  M. Luckie, A. Dhamdhere, k. claffy, and D. Murrell, "Measured Impact of Crooked Traceroute", ACM SIGCOMM (CCR), 2011.
- Vulnerability assessment: ingress filtering (w/ NPS)
  R. Beverly, A. Berger, Y. Hyun, and k. claffy, "Understanding the Efficacy of Deployed Internet Source Address Validation Filtering", IMC 2009.
- Internet topology mapping: IP alias resolution
  - Compare accuracy of alias resolution techniques at Internet scale
  - Enhancements: (APAR++) [CCR 2010] , MIDAR [TON 2012]
  - Combine techniques (iffinder, kapar, ally, MIDAR) to improve overall accuracy
  - While others still saying it's impossible [AMS2009]
  - Daunting challenge as always: remains validation

Internet-scale IP Router 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: fingerprints to find aliases
    - IP ID: 16-bit header field supporting frag&reassembly
    - Two interfaces on same router probed closely in time will return similar IP ID values: over time, similar time-series velocity.
- Architecture paper to appear in TON2012

MIDAR Approach

http://www.caida.org/tools/measurement/midar/

- 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

IP ID over time

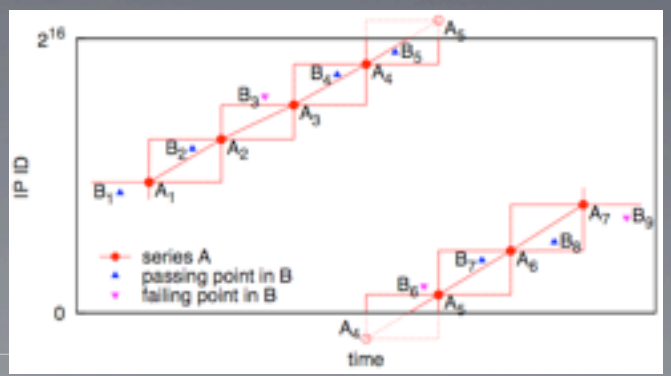

MIDAR Elimination Stage

- 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

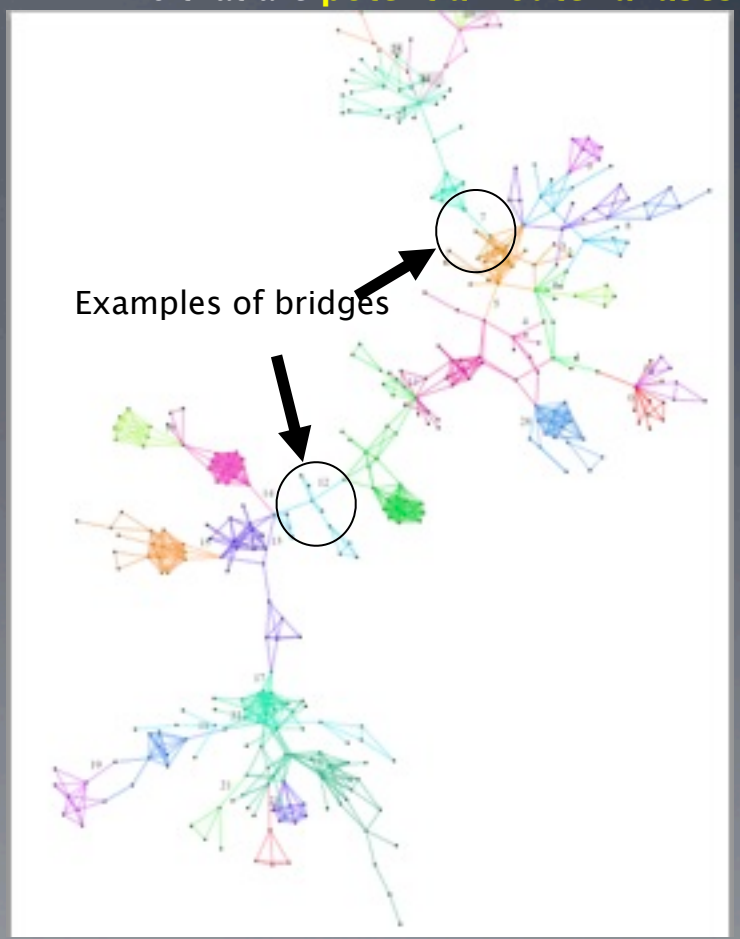

MIDAR Results

	2010-01	2010-04	2010-07	2011-04	2011-10	2012-07
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	2.34 M 1.86 M 1732 G
Shared pairs after Discovery stage	1.63 M	4.00 M	5.49 M	6.83 M	7.00 M	9.24 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	3.88 M 120 k 423 k

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

Internet Topology Data Comparison

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

AS Rank

- metric/indicator of influence over the global Internet interdomain routing system
- applications to: Internet science/modeling, infrastructure robustness/protection, public policy
- based on inferred economics of AS business relationships using data from public BGP 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/

Top Tier AS clique

- Largest set of full meshed, clique, ASes from the top 40 ASes by degree

Top Tier Ases **clique members** over time

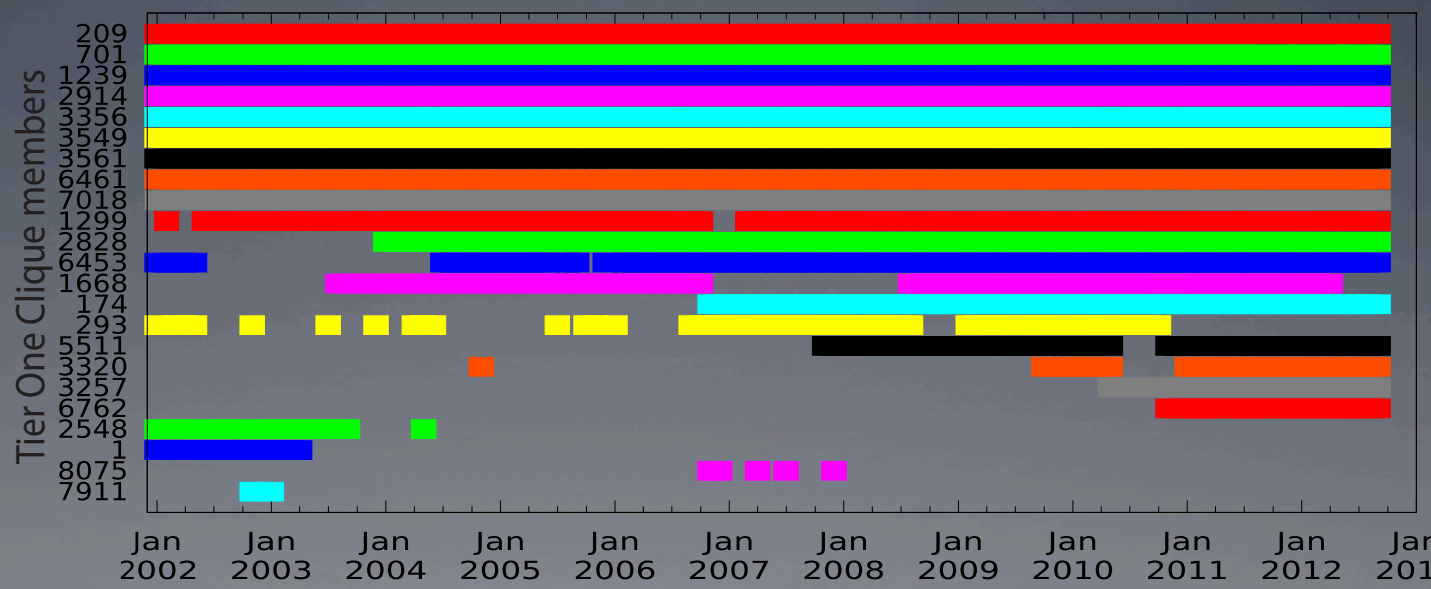

AS Rank: screen shot

Top 10 ranked
AS by customer cone

Level 3

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

The top ASes ranked by customer cone size are displayed below. **Dataset:** 2012.06.01
For information about a specific AS, enter its AS name, its AS number, or the name of the Org of which the AS is a member.
Look up an AS by number or name [Search]

Table **shows** 10 of 41377 ASes, **sorted by** number of ASes in customer cone [update view]

AS rank	AS number	AS name	Org name	customer cone						AS degree
				Number of			Percentages of all			
				ASes	IPv4 Prefixes	IPv4 Addresses	ASes	IPv4 Prefixes	IPv4 Addresses	
1	3356	LEVEL3	Level 3 Communications, Inc.	29,037	291,344	1,456,474,500	70%	69%	56%	3318
2	3549	LVLT-3549	Level 3 Communications, Inc.	25,105	260,300	928,688,621	60%	62%	36%	1499
3	1299	TELIANET	TeliaNet Global Network	21,887	217,829	933,429,605	52%	52%	36%	684
4	174	COGENT-174	Cogent Communications	21,206	215,777	849,203,950	51%	51%	33%	3539
5	3257	TINET-BACK...	Tinet Spa	18,211	206,207	796,599,260	44%	49%	31%	884
6	2914	NTT-COMMUN...	NTT America, Inc.	16,812	190,764	787,647,574	40%	45%	30%	791
7	701	UUNET	MCI WorldCom	14,781	188,837	879,888,016	35%	45%	34%	1812
8	1239	SPRINTLINK	U.S. Sprint	14,275	166,089	1,077,770,948	34%	39%	42%	969
9	6762	SEABONE-NET	Info-tel Communication S.r.l.	12,907	156,342	588,009,687	31%	37%	23%	264
10	6453	AS6453	TATA Communications formerly VSNL is Leading ISP	11,450	150,440	630,628,770	27%	36%	24%	549

data sources

geolocation	database	2012.06.25	netacuity	
organization	whois	2012.04.01	AFRINIC, APNIC, ARIN, LACNIC, RIPE	
topology	BGP	2012.06.01, 2012.06.02, 2012.06.03, 2012.06.04, 2012.06.05	ripe	rrc00, rrc01, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc11, rrc12, rrc13, rrc14, rrc15
			routeviews	eqix, isc, linx, routeviews2, saoppaulo, sydney

http://www.caida.org/

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
Org name:	TeliaNet Global Network
AS rank:	3
Country:	(no data)
Customer cone size:	21,887
AS degree:	684

Ranking

neighbor				type	projected peering: cone size (% of AS 1299's original cone size)		projected peering: cone size ratio	AS degree
AS rank	AS	AS name	Org name		of neighbor AS	of AS 1299		
352	3301	TELIANET-S...	TeliaNet Global Network	-- sibling	0.26%	99%	0.26	66
26818	31080	O2-AS	TeliaNet Global Network	-- sibling	0.00%	100%	0.00	5
1	3356	LEVEL3	Level 3 Communications, Inc.	-- peer	132%	100%	75.38	3,318
2	3549	LVLT-3549	Level 3 Communications, Inc.	-- peer	114%	100%	87.18	1,499
4	174	COGENT-174	Cogent Communications	-- peer	96%	100%	96.89	3,539
5	3257	TINET-BACK...	Tinet Spa	-- peer	83%	100%	83.20	884
6	2914	NTT-COMMUN...	NTT America, Inc.	-- peer	76%	100%	76.81	791
7	701	UUNET	MCI WorldCom	-- peer	67%	100%	67.53	1,812
8	1239	SPRINTLINK	U.S. Sprint	-- peer	65%	100%	65.22	969
9	6762	SEABONE-NET	Info-tel Communication S.r.l.	-- peer	58%	100%	58.97	264

Customers, providers, and peers

neighbor				type
AS rank	AS	AS name	Org name	
352	3301	TELIANET-S...	TeliaNet Global Network	-- sibling
26818	31080	O2-AS	TeliaNet Global Network	-- sibling
1	3356	LEVEL3	Level 3 Communications, Inc.	-- peer
2	3549	LVLT-3549	Level 3 Communications, Inc.	-- peer
4	174	COGENT-174	Cogent Communications	-- peer
5	3257	TINET-BACK...	Tinet Spa	-- peer
6	2914	NTT-COMMUN...	NTT America, Inc.	-- peer
7	701	UUNET	MCI WorldCom	-- peer
8	1239	SPRINTLINK	U.S. Sprint	-- peer
9	6762	SEABONE-NET	Info-tel Communication S.r.l.	-- peer

AS Rank Validation

- Interface to provide corrections to relationships

relationship correction page

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	

corrections

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

Geolocation Tools Comparison (to support viz and analysis)

- 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

AS Rank Visualization

Graphical view of providers, peers, and customers

provider

peer

customer

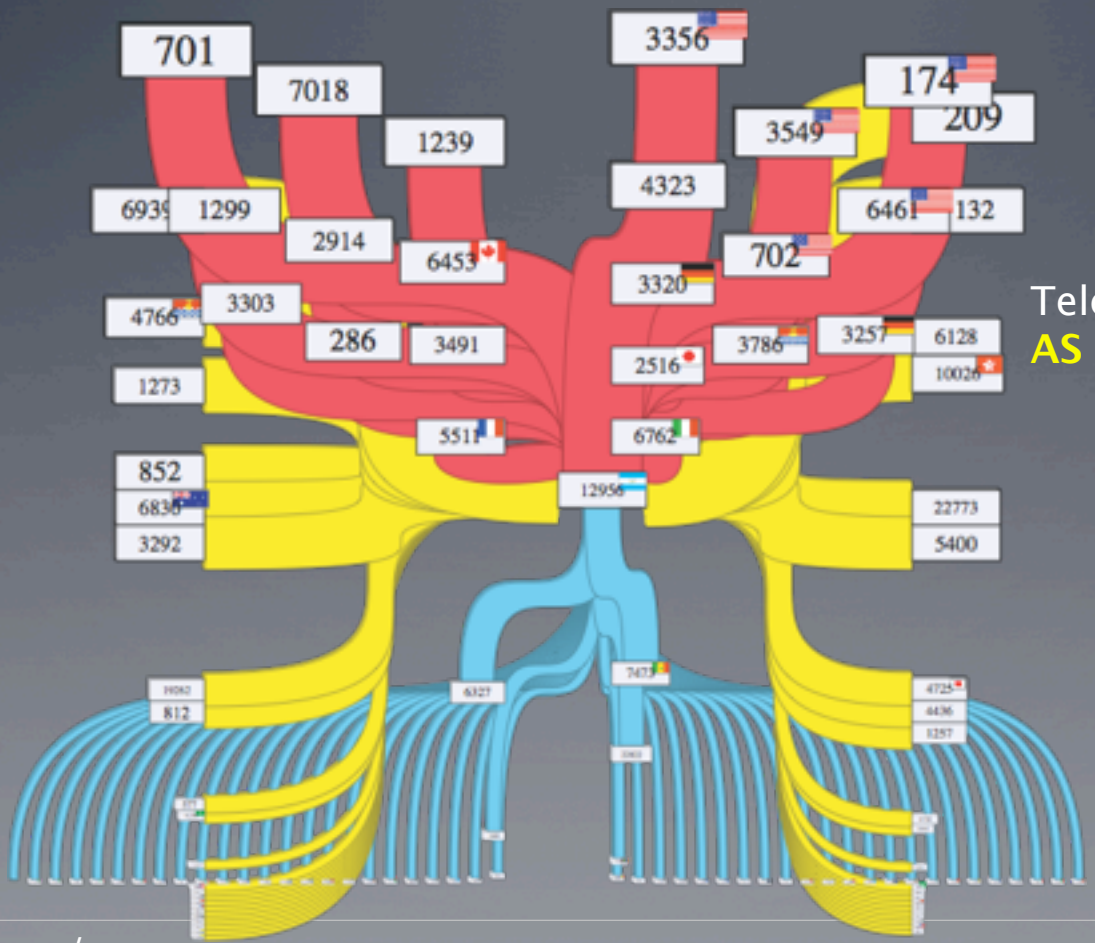

Telephonica (AS 12956)
AS neighbors

http://www.caida.org/

Location Graph

Semi-geographic view of all routers for a given AS

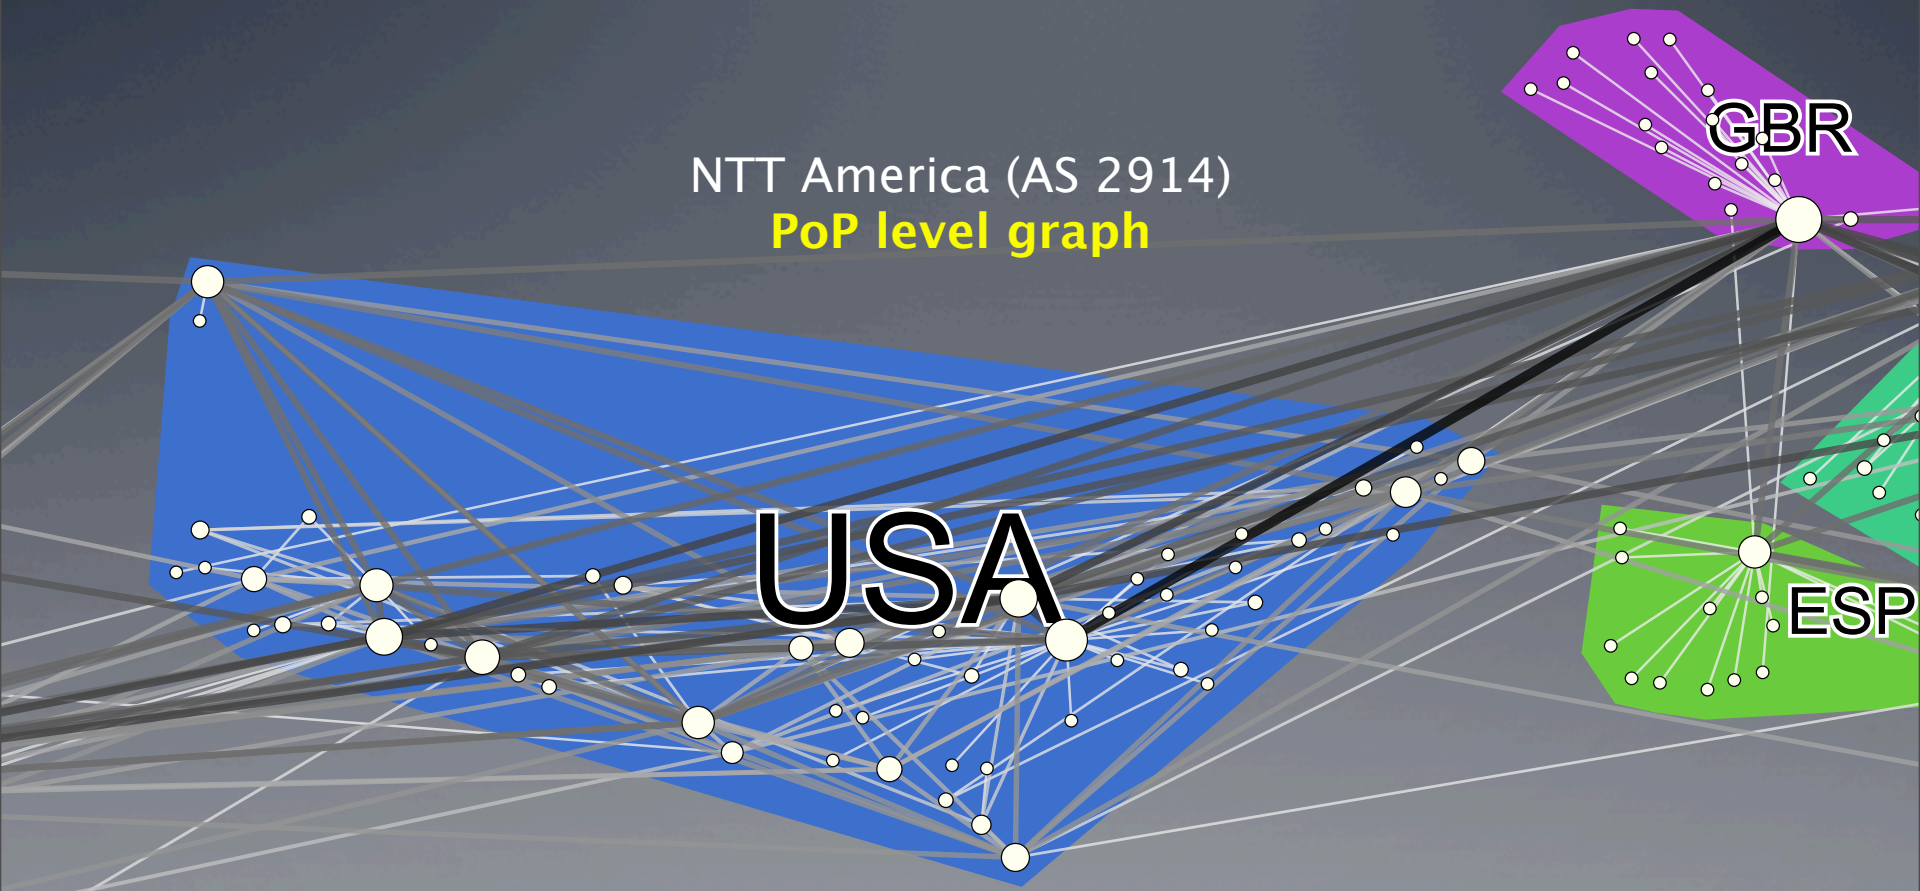

Integrated Visualization of Topological Connectivity

UCSD router (prototype)
single router graph

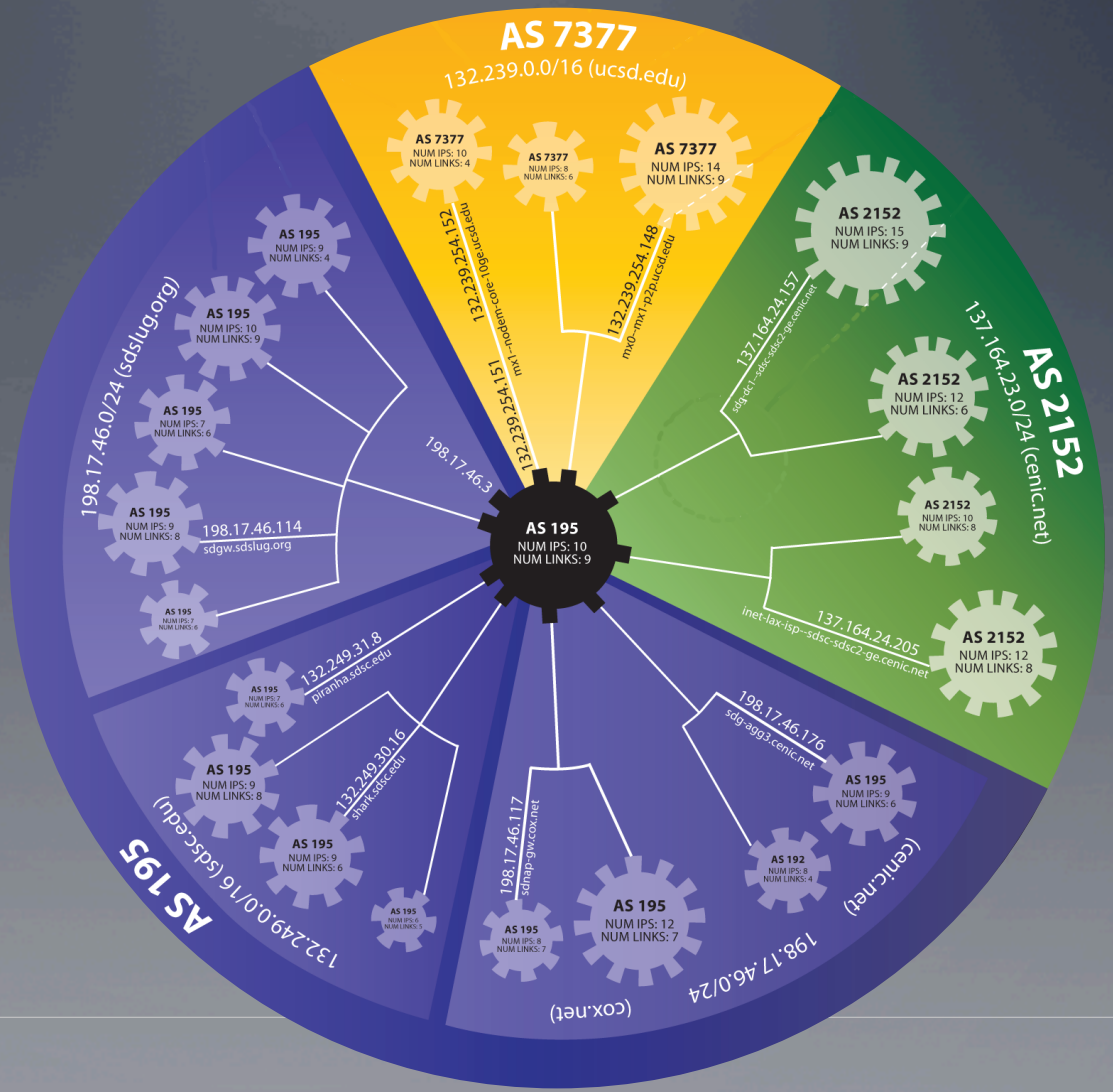

Published Experiments Using Ark

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.

Delivered 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
- AS Rank documentation, validation, new algorithms and interface
- Web-based interface to topo-on-demand service

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

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

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

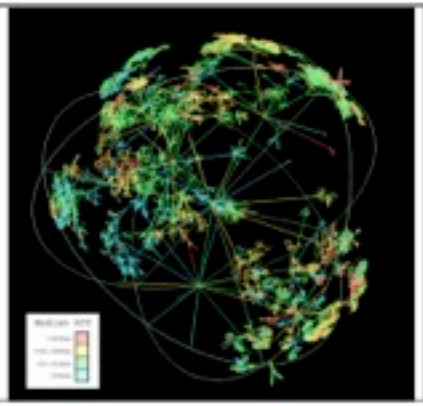

Walrus visualizations of round-trip time measurements made by CAIDA's macroscopic Internet topology monitor located in Herndon, VA, USA.

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

http://www.caida.org/

CAIDA 2011 Annual

http://www.caida.org/home/about/annualreports/2011/

- Research and Infrastructure Projects
- Tools
- Data
- Workshops
- Publications
- Presentations
- Web Site Usage
- Organizational Chart
- Funding Sources
- Operating Expenses

UC San Diego SDSC