Leveraging the Science and Technology of Internet Mapping for Homeland Security

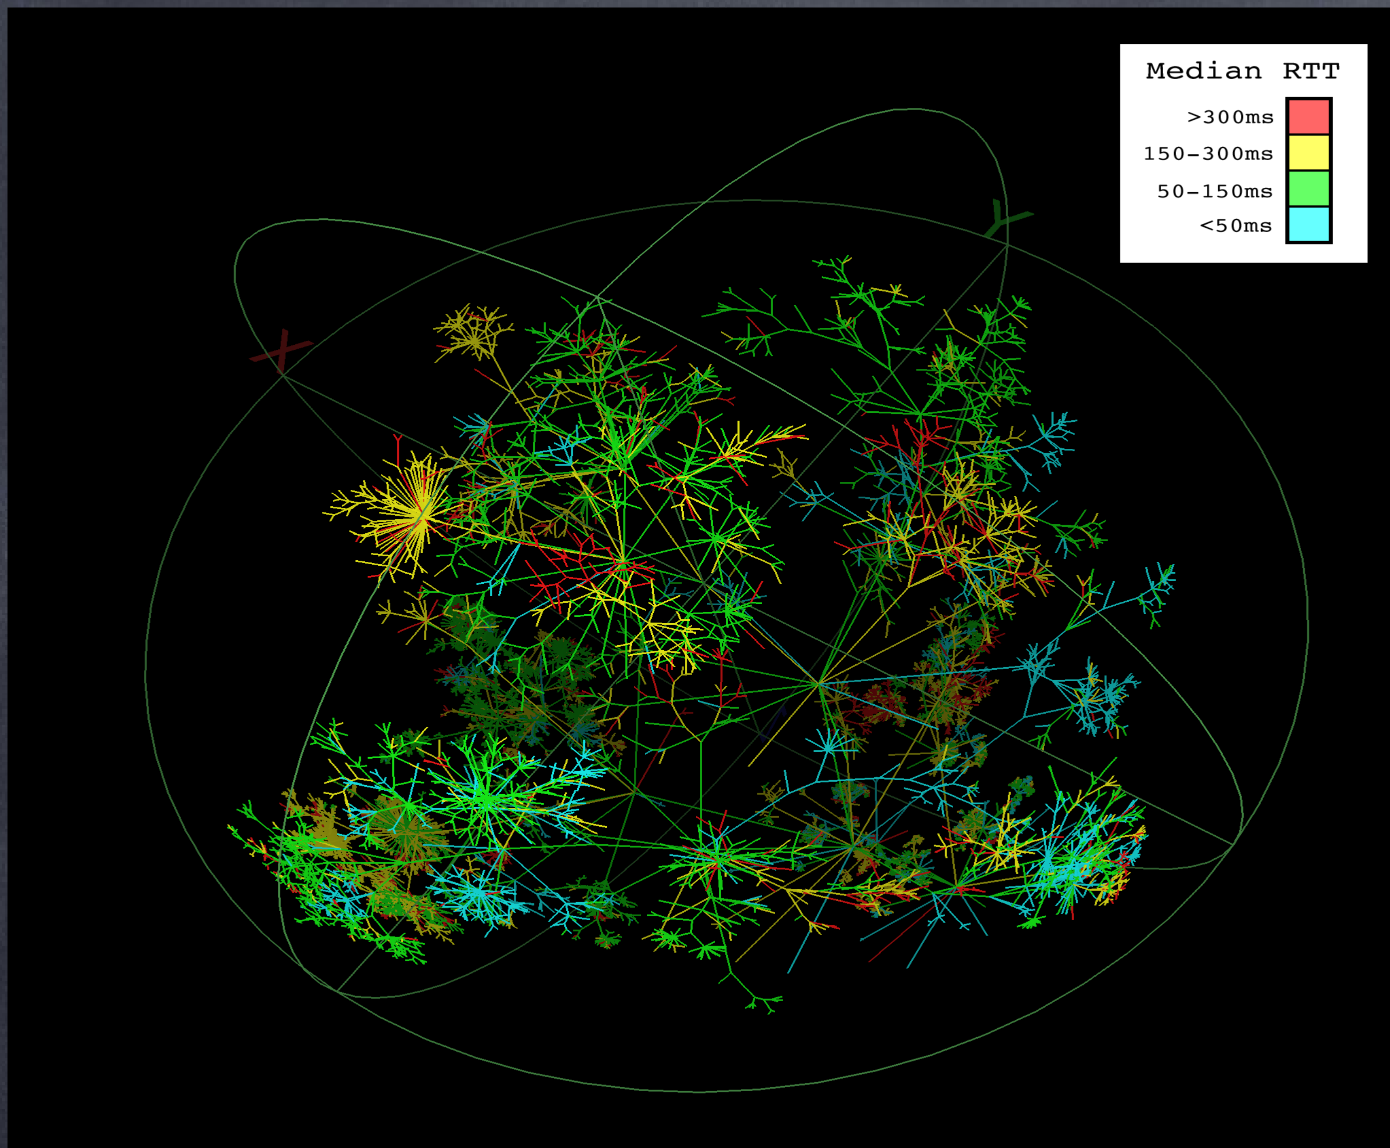

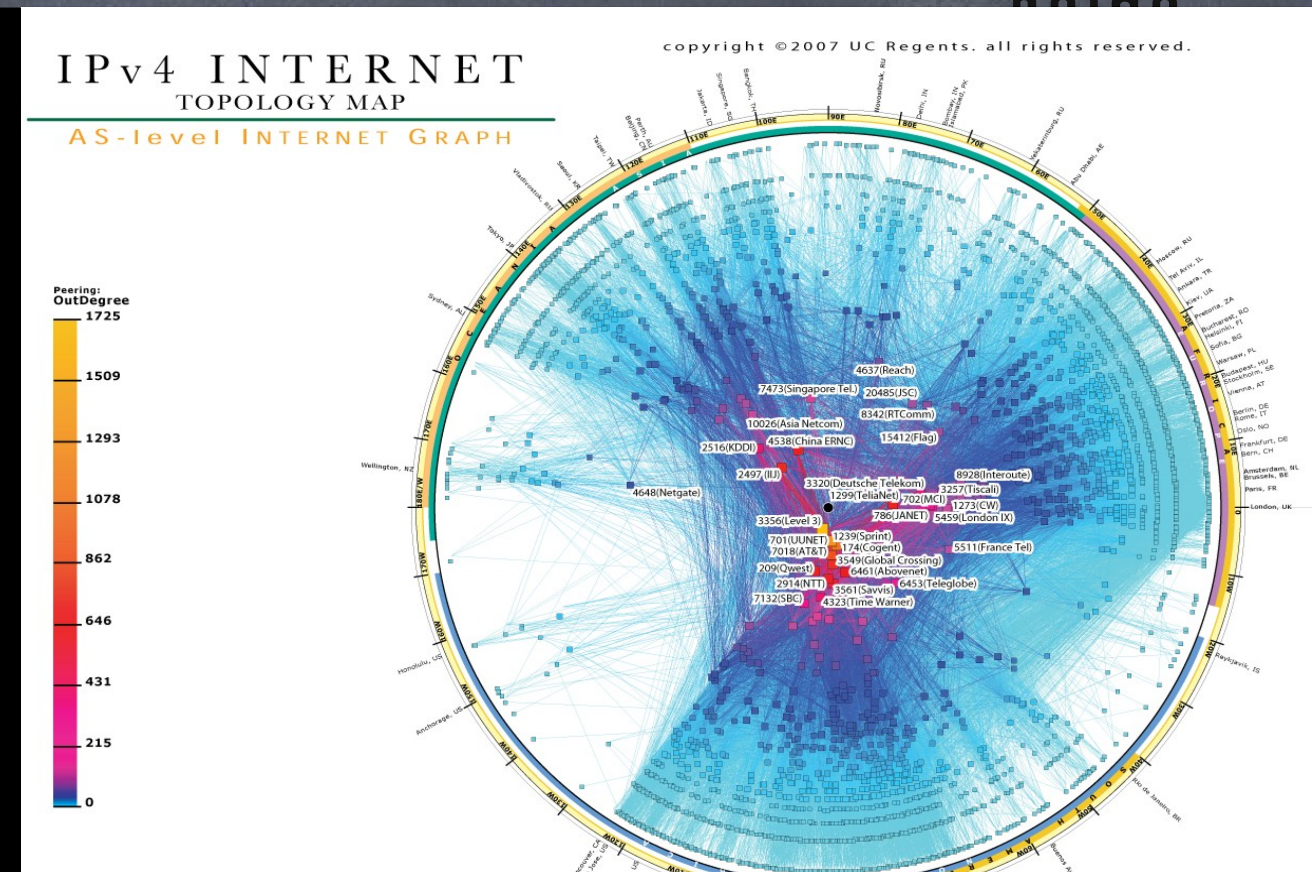

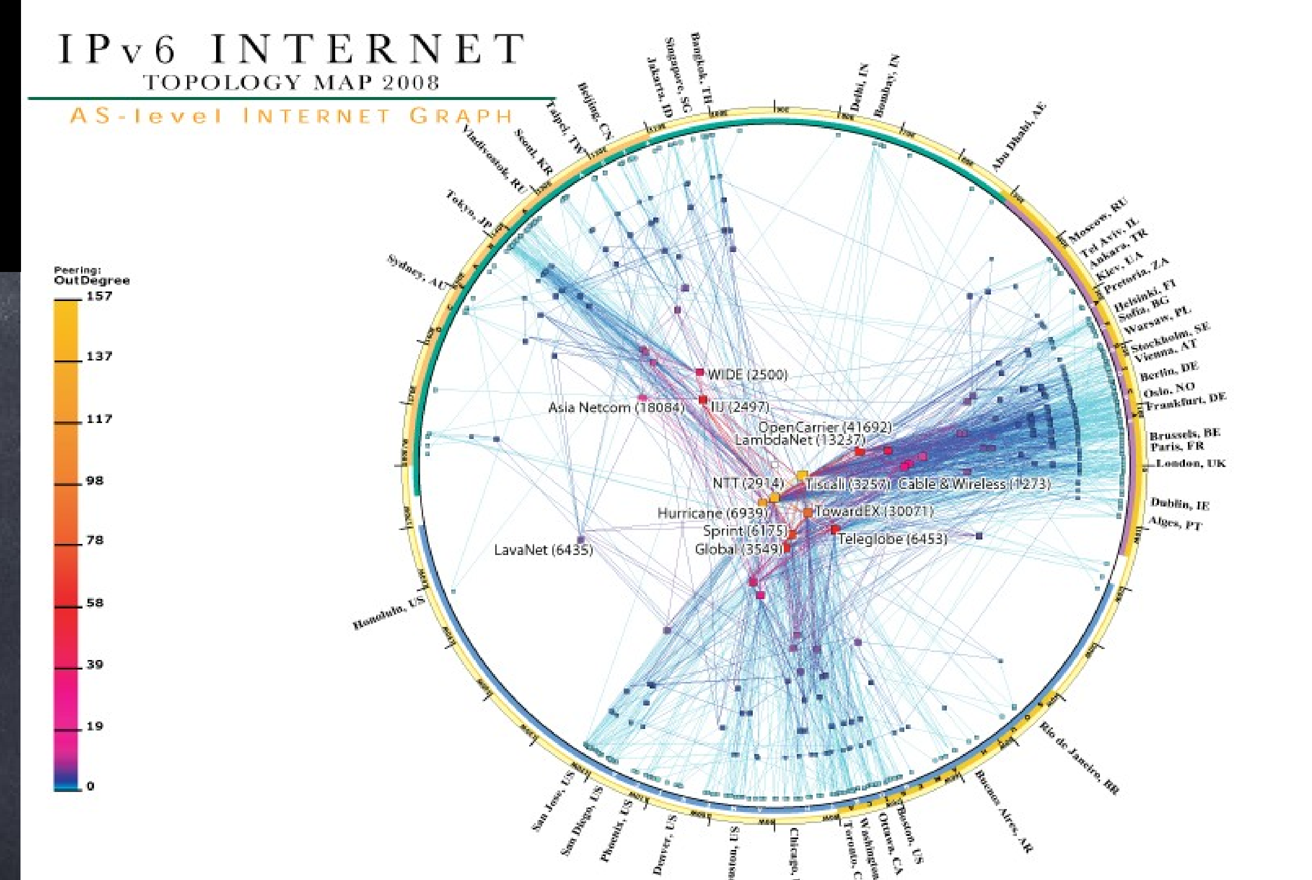

Young Hyun, Ken Keys, Amogh Dhamdhere, Bradley Huffaker, Joshua Polterock, Marina Fomekov, Dima Krioukov, kc claffy

CAIDA
DHS – PI meeting
SRI Menlo Park, CA
1-2 September 2010

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...

Technical Approach

- Integrate 6 strategic measurement and analysis capabilities:
  - new architecture for continuous topology measurements (Archipelago, or “Ark”),
  - Topology analysis techniques, e.g., IP alias resolution
  - 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/

approach

Archipelago (Ark)

- CAIDA's measurement infrastructure
- Built on decade of achievements, from SIGCOMM to MOMA
- Launch 12 Sept 2007
- 46 active IPv4 probers
    - 16 in US
- 12 active IPv6 probers
- 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 supports researchers with an environment for easy development and rapid prototyping 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
  - geolocation technology assessment
  - 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

Insights enabled

- Probing technique performance comparison (w/.nz)
- Vulnerability assessment: ingress filtering (w/NPS)
- Internet topology mapping: IP alias resolution
    - compare performance and accuracy of known alias resolution techniques used at Internet scale
    - enhancements: (APAR++), MIDAR (radargun++)
    - combine techniques (iffinder, kapar, ally, MIDAR) →
    - MAARS: most accurate complete IP-to-router mapping
    - while others still saying it's impossible, AMS2009
    - daunting challenge as always: remains validation

2010 technical (infra.) accomplishments

- 46 monitors active, 12 probing Ipv6

- IPv4 topology data
  - 3.123TB data served by PREDICT, data.caida.org
  - Sep 2007 to July 2010 (35 months):
    - 7.8B traceroutes; 1K cycles
  - per month: 350M traceroutes; ~140 GB data
  - key input to, e.g., AS links and alias resolution
  - Each team collects traces from 9.1 million /24s

- IPv6 topology data

Ark monitors/data over time

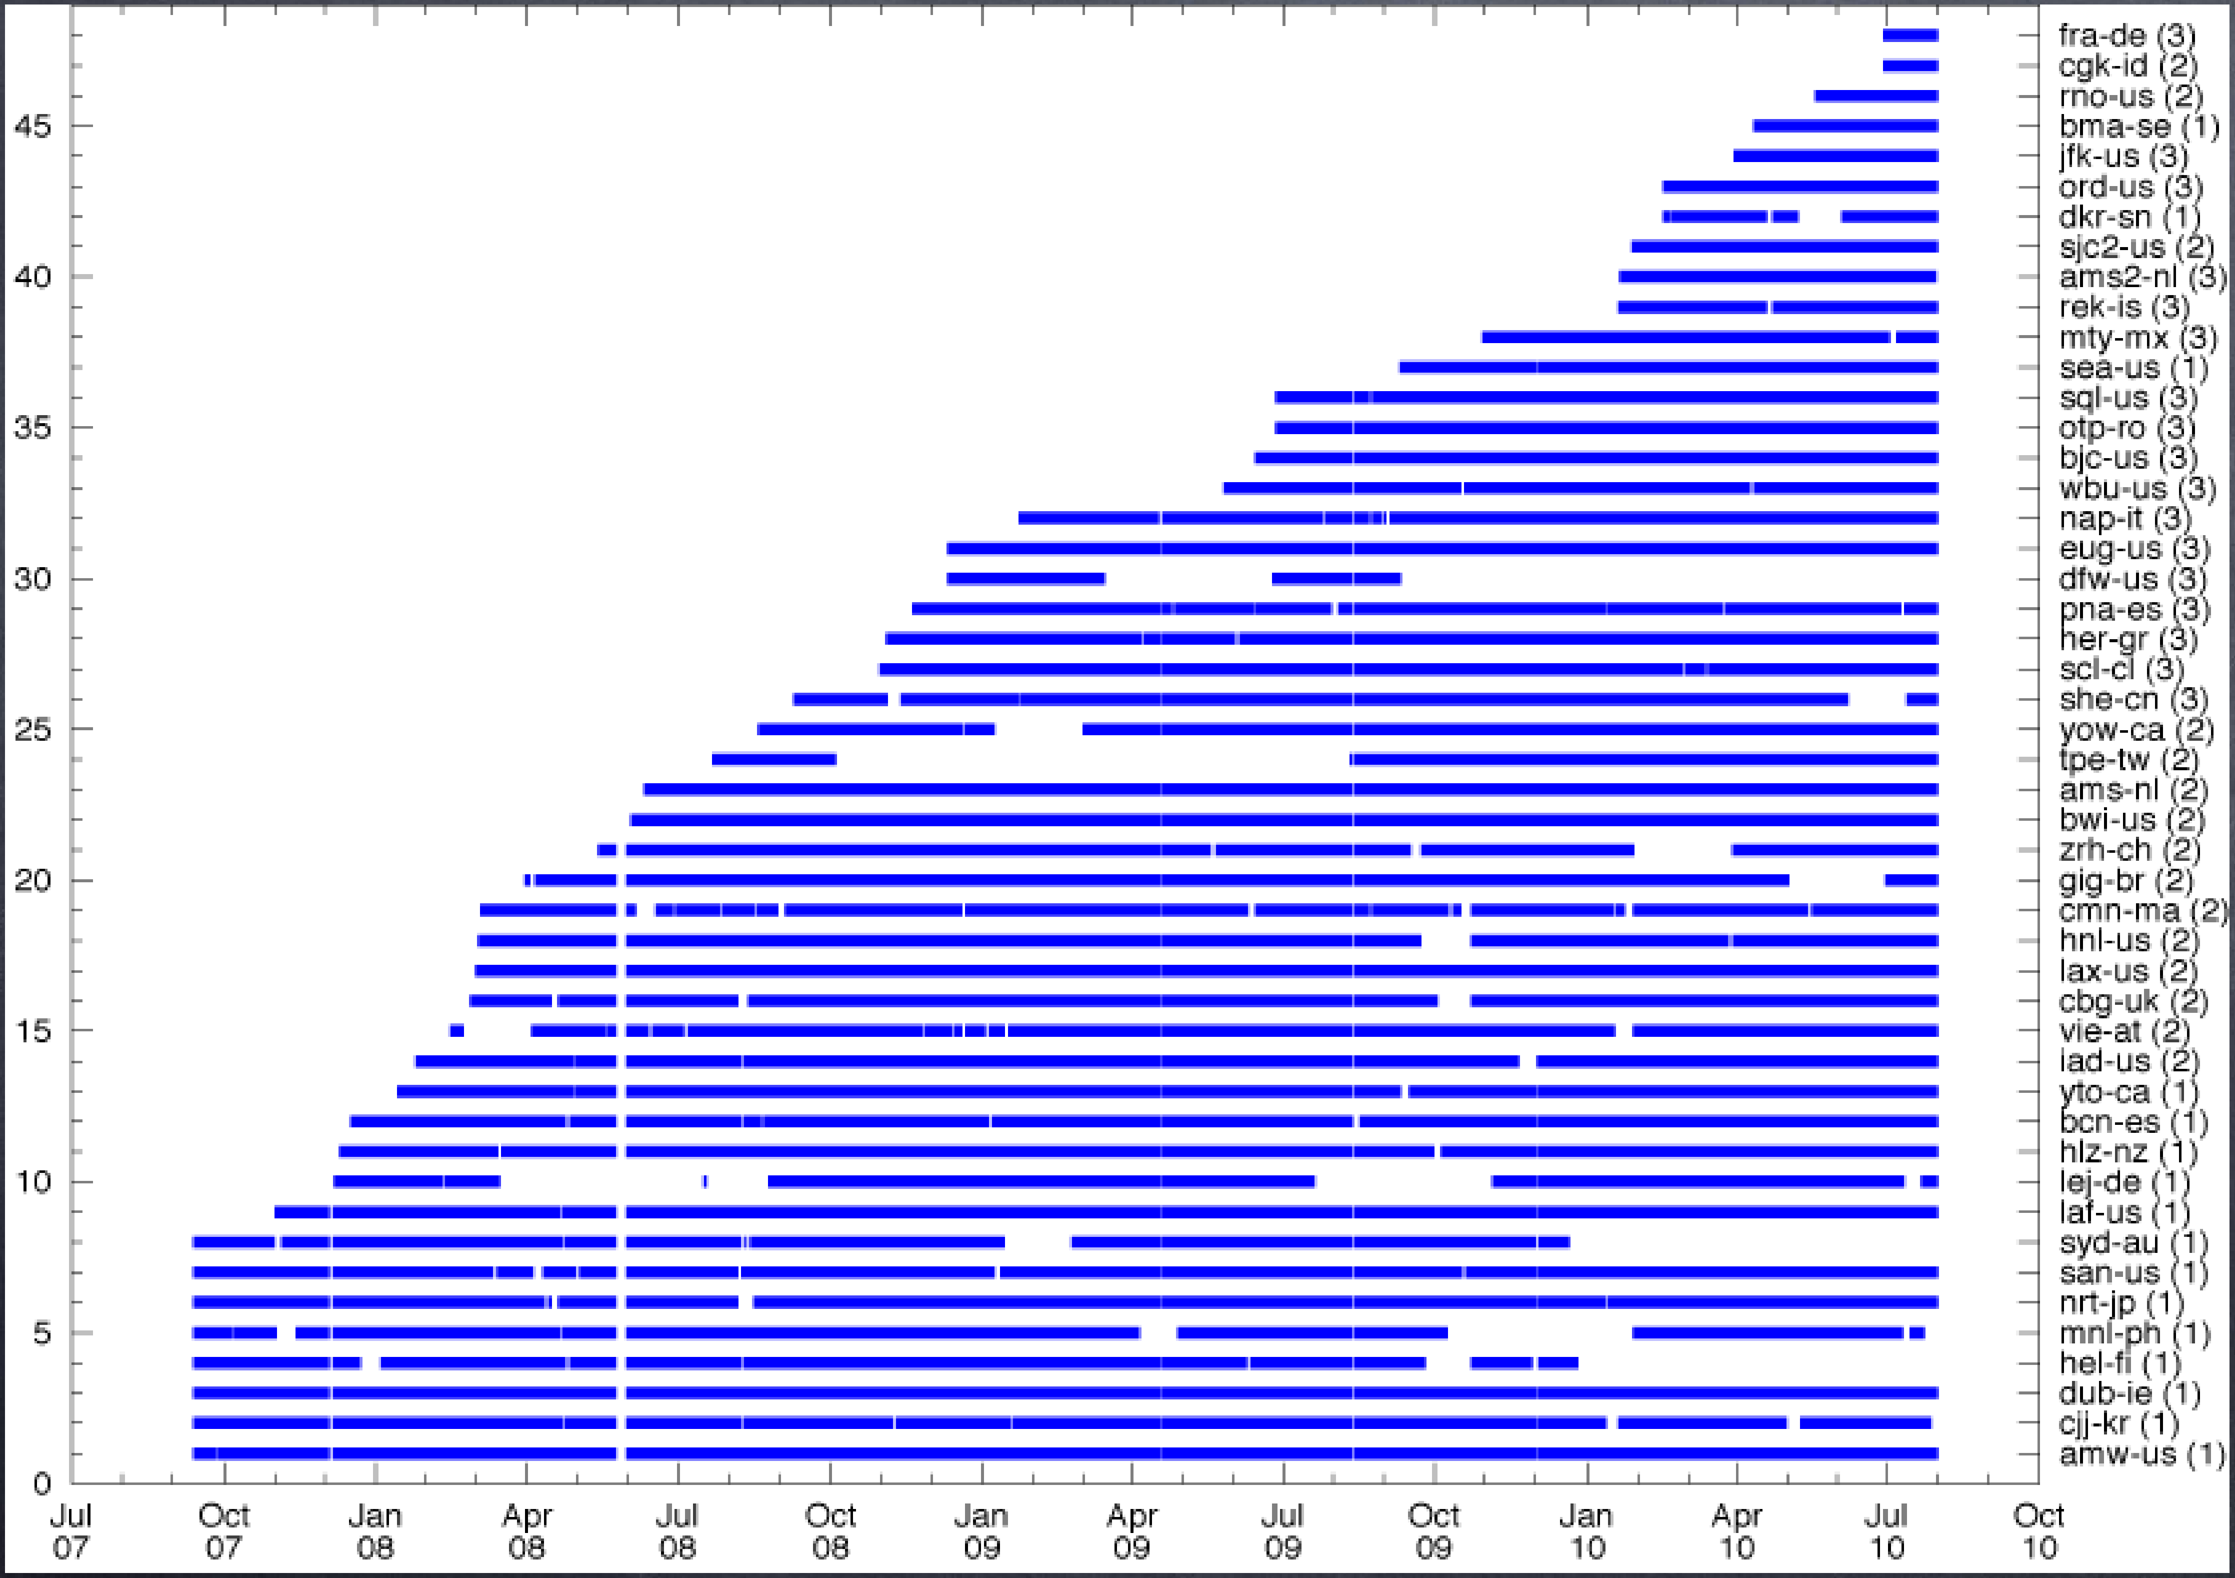

2010 technical accomplishments

- AIMS2 workshop report → CCR
- AS-level & router-level graph “ITDK” (Jan,Apr,Jul)
  http://www.caida.org/data/active/internet-topology-data-kit/
- Dual AS-router graph (June)
  - Preliminary dual graph PAM 2010 paper http://www.caida.org/publications/papers/2010/as_assignment/
- Tool to calculate topology statistics – topostats (Feb)
  http://www.caida.org/tools/utilities/topostats/
- Supporting software: mper, Marinda, MIDAR, kapar
- AS Rank revival

Internet Topology Data Kit (ITDK)

caida
www.caida.org

- Two router-level topologies
  1)optimized for **accuracy:**
  MIDAR+iffinder,
  highest confidence aliases with low false positives.
  2)optimized for **completeness:**
  adds kapar results,
  more alias coverage, more false positives (inflating routers)

- Data files: routers, links, router-to-AS mappings, DNS

Topology Data Architecture

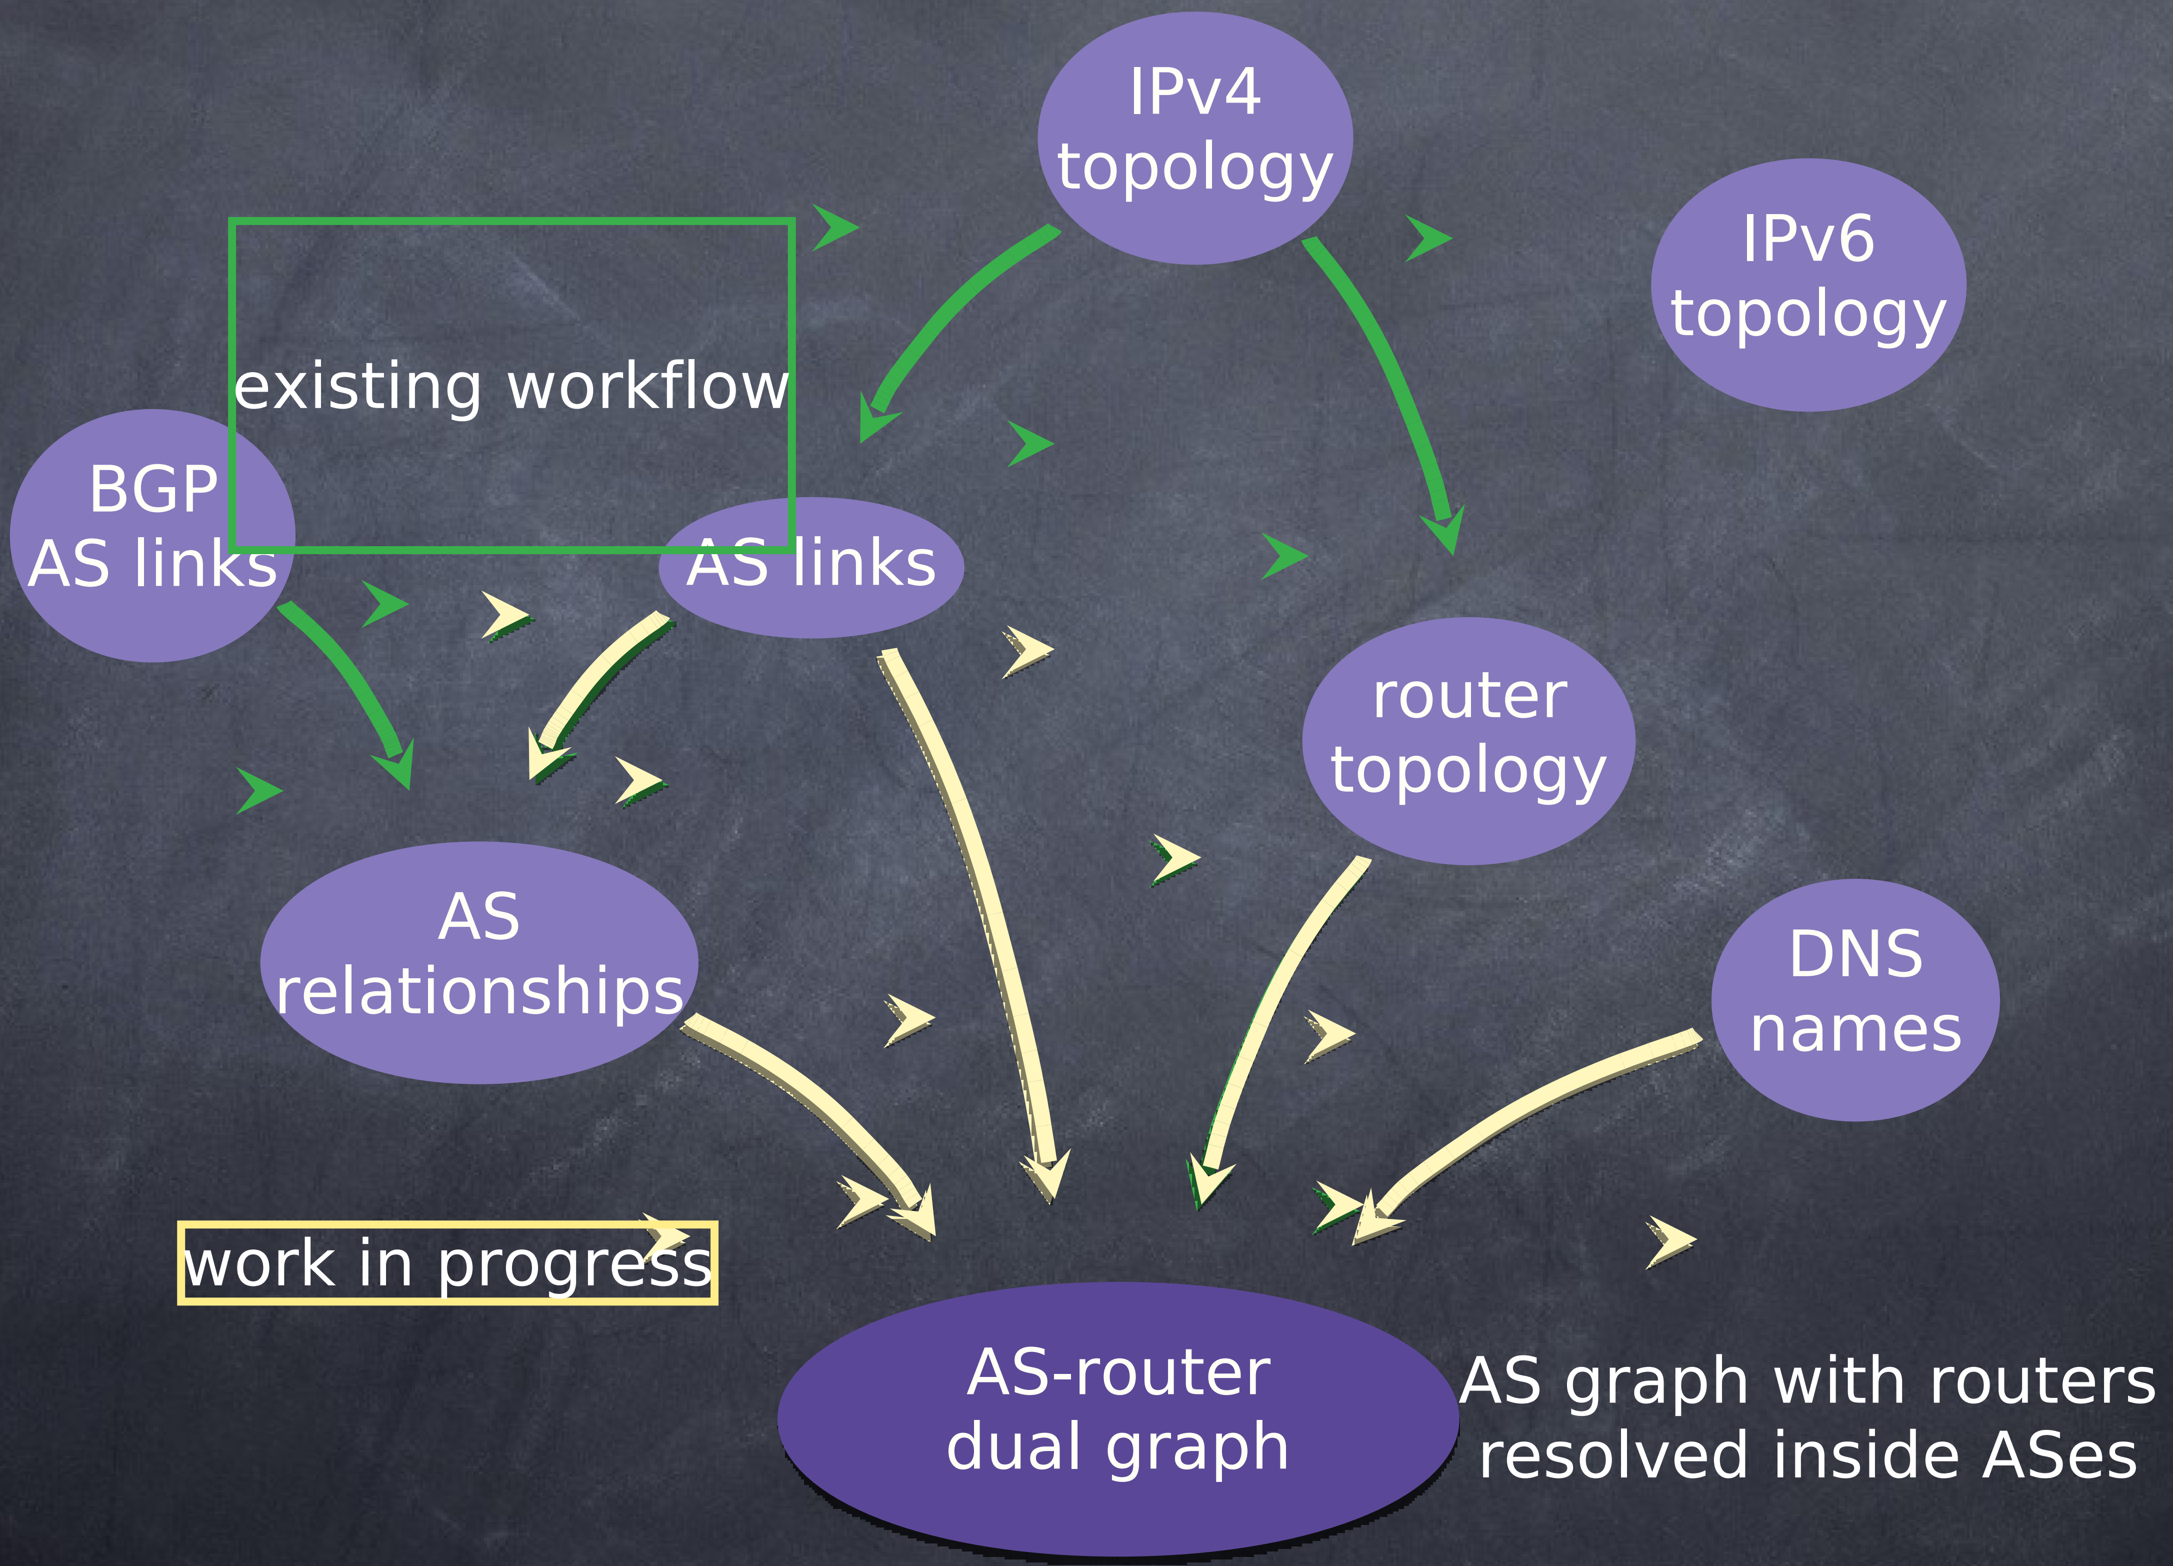

Data: Internet Topology Data Kit (ITDK)

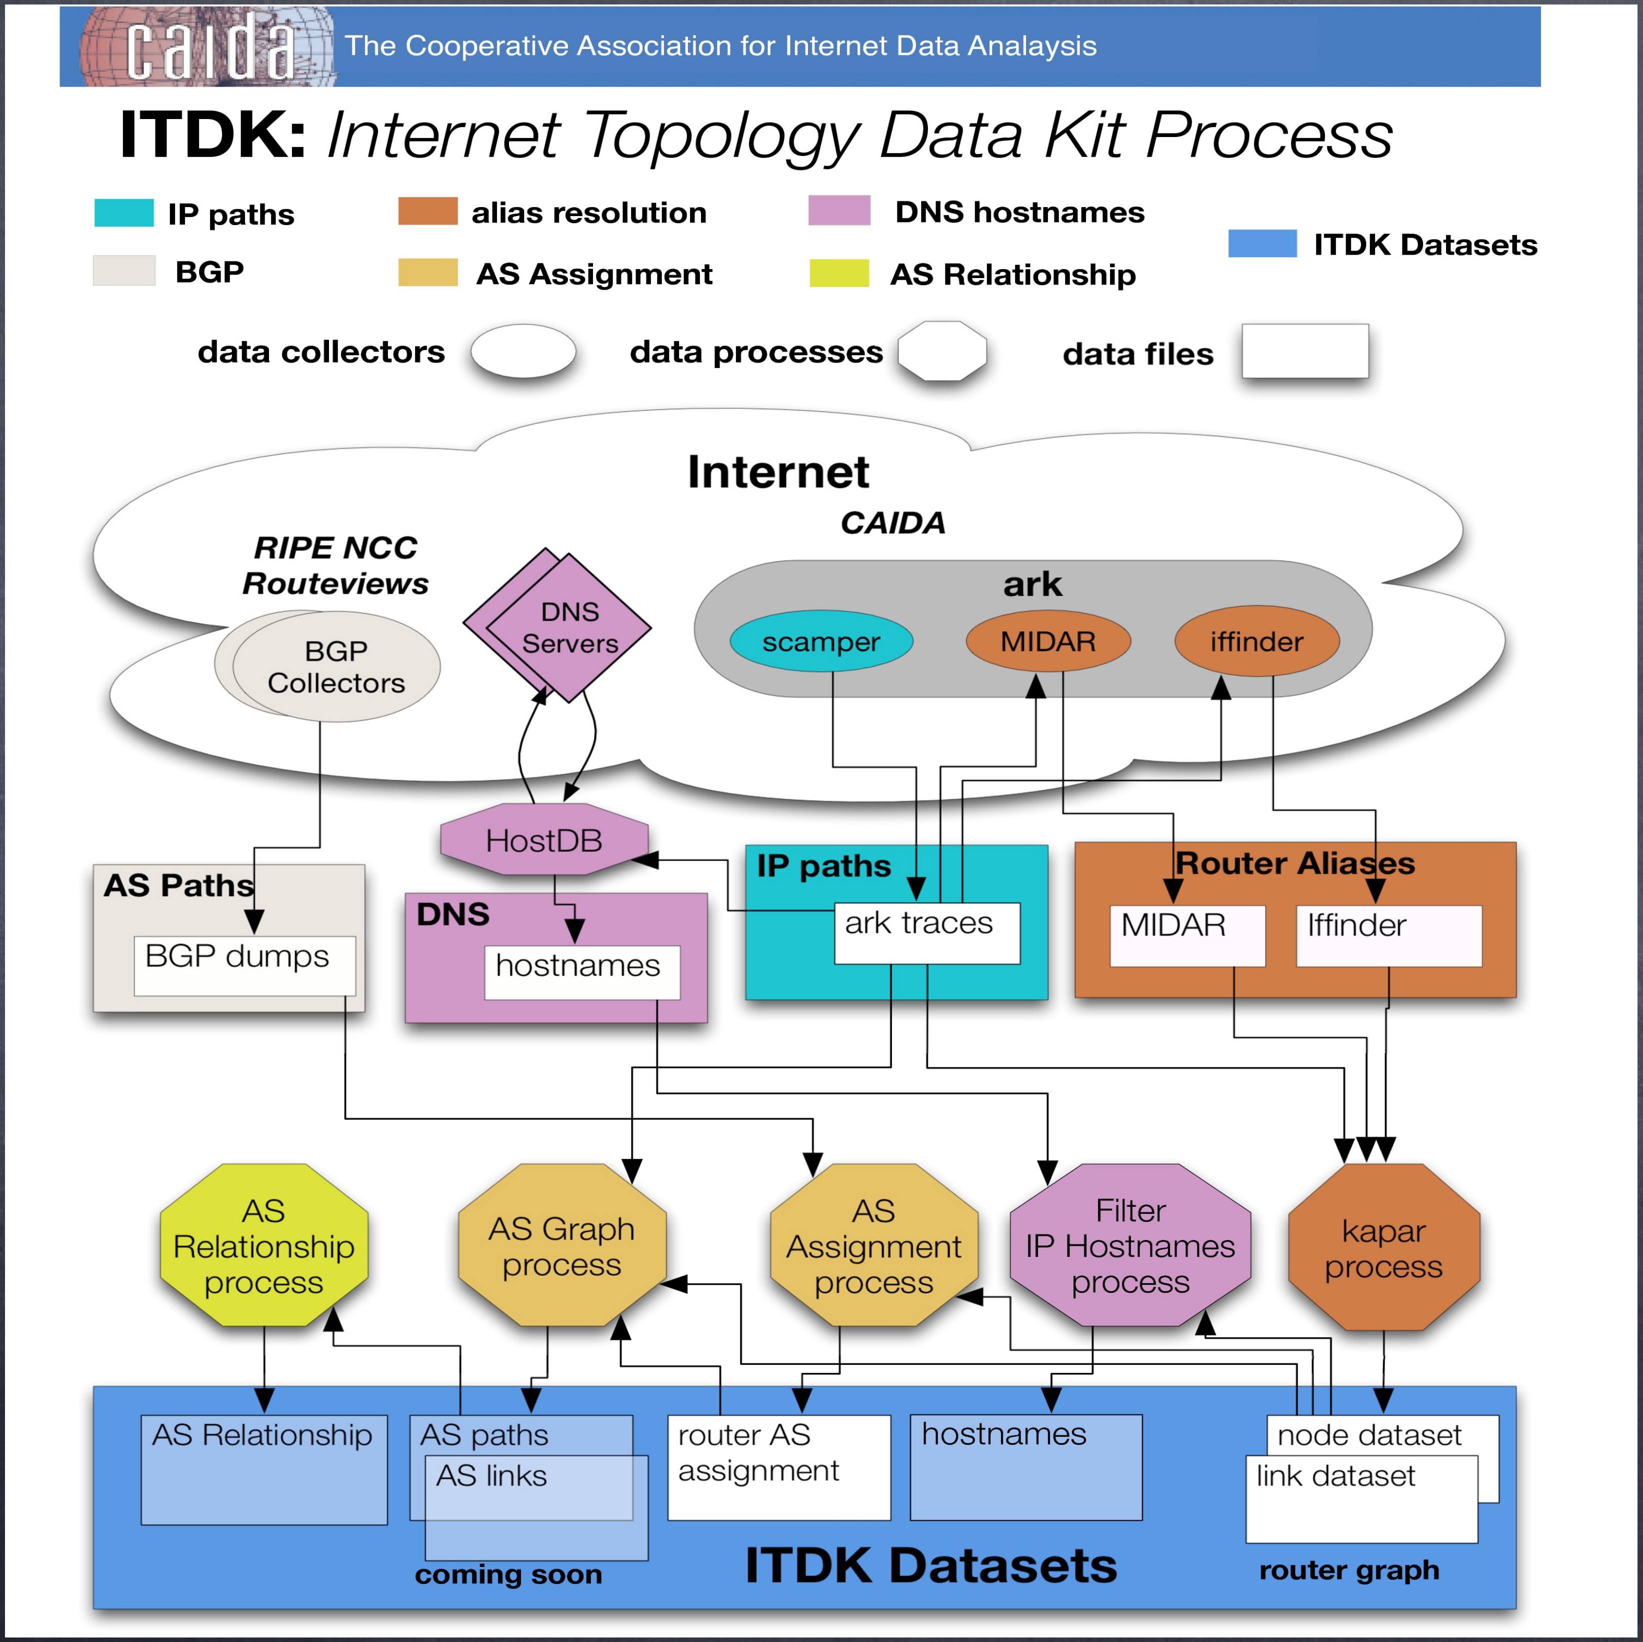

Data: IPv4 Routed /24 Topology

- ongoing large-scale topology measurements
  - ICMP Paris traceroute to every routed /24 (9.1 million)
    - ~ 138.7 /8-equivalents of routed space (Aug 2010)
    - 10.1% increase since Oct 2009
    - more routed space than unrouted space in IPv4
  - running *scamper* (Matthew Luckie, U. Waikato)
- dynamically assign measurements to teams
  - 3 teams active
  - 15/16-member team probes every /24 in 2-3 days at 100pps
    - only one monitor probes each /24 per cycle (=one pass through all /24's)

Alias Resolution

- goal: collapse observed interfaces into routers
- earlier at CAIDA: iffinder, kapar (APAR++)
- past year: MIDAR (Radargun++)
  - two interfaces on same router respond in similar way
  - IP ID values in responses can be 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.

Alias Resolution: myths?

// *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 timeseries must be monotonic
  - sliding window for scalable probing
  - 4 probing methods: TCP, UDP, ICMP, "indirect" (TTL expired)
  - multiple monitors
  - stages: estimation, discovery, elimination, corroboration

MIDAR Elimination Stage

- potential alias set found in Discovery stage: 378 IPs, 977 suspected pairs
- Testing pairwise 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
- Covers all pairs with, e.g., 411 timeseries instead of 1954.
- More efficient, reduces chance of rate limiting

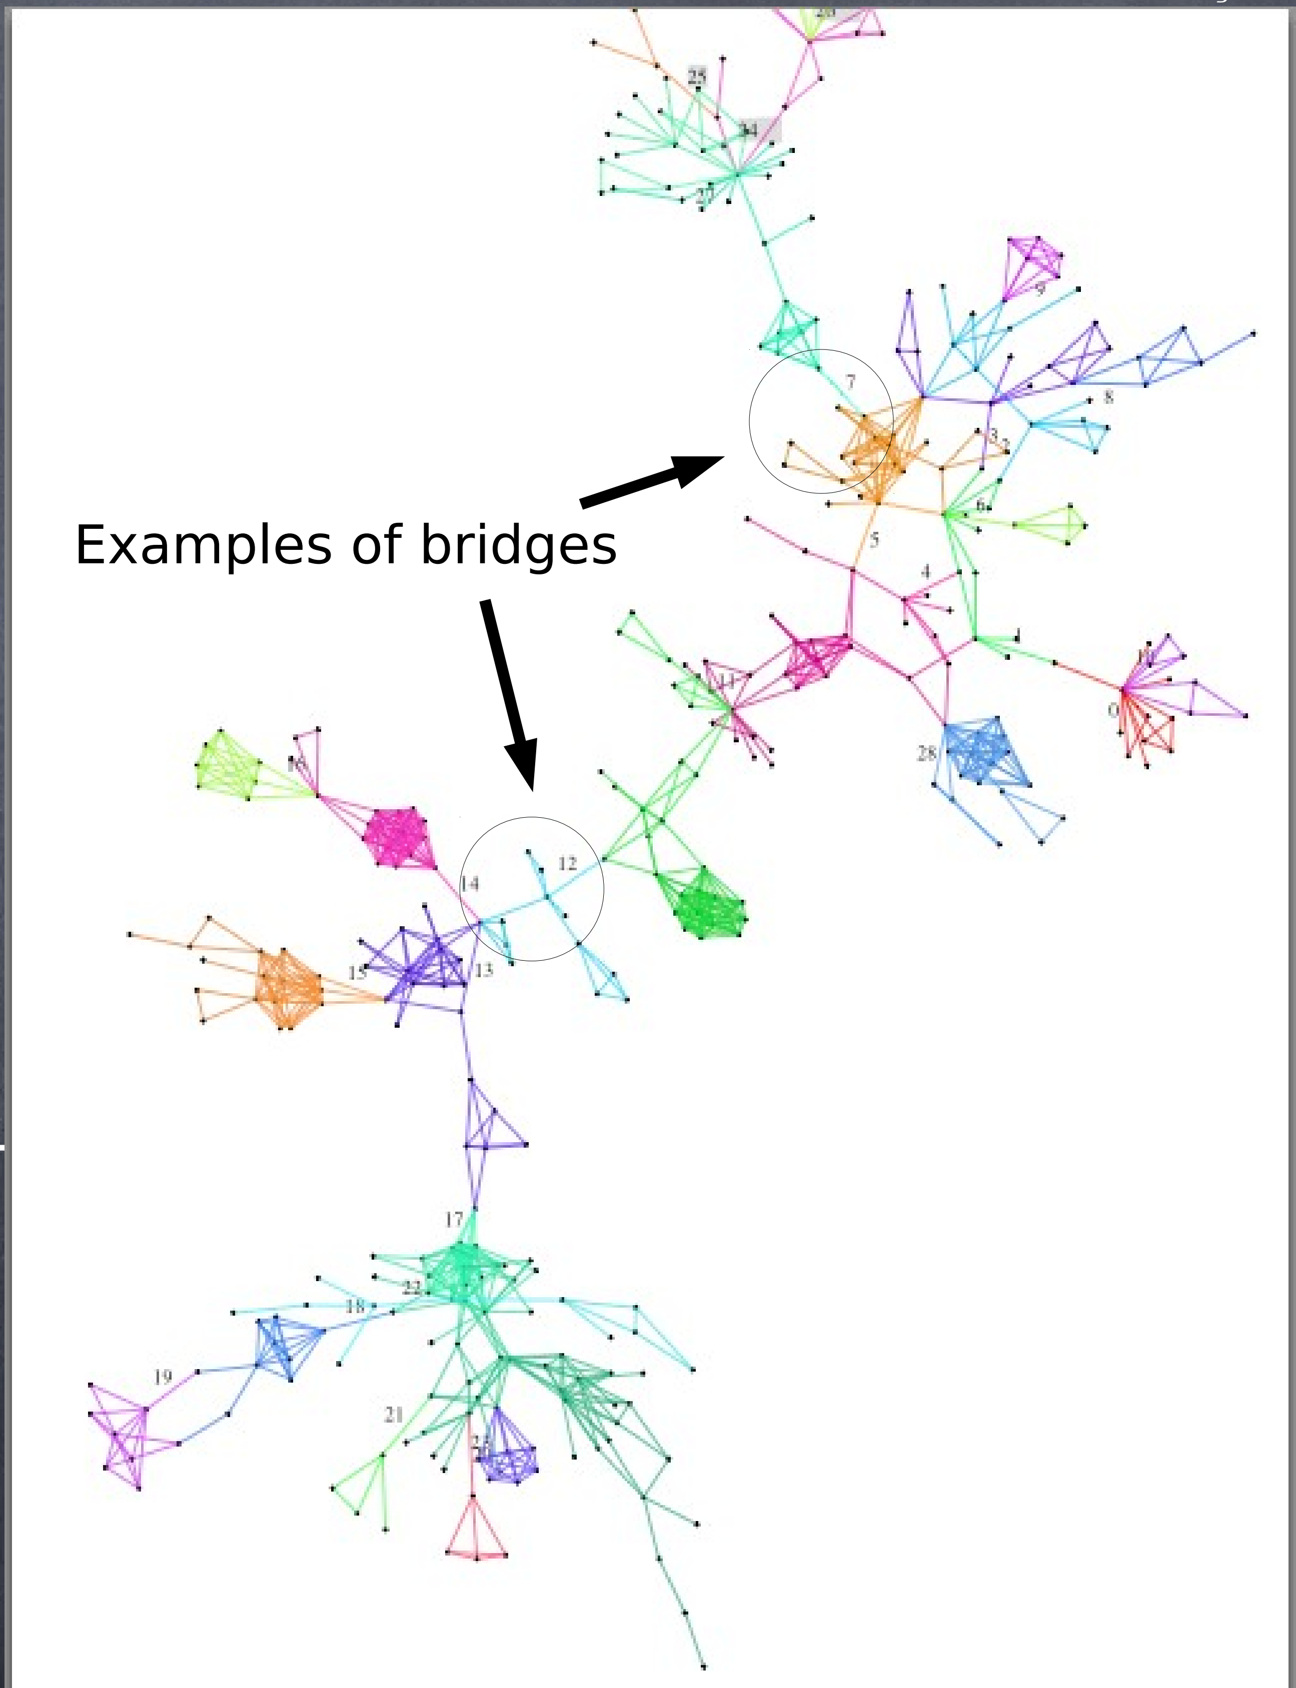

MIDAR Results

	2010-01	2010-04	2010-07
input address	1.12 M	1.50 M	1.90 M
monotonic address	0.99 M	1.20 M	1.44 M
Possible pairs	486 G	724 G	1038 G
Shared pairs after Discovery stage	1.63 M	4.00 M	5.49 M
Final results			
·Shared pairs	0.433 M	1.36 M	1.67 M
·Routers	69 k	108 k	121 k
·Addresses on routers	189 k	383 k	426 k

caida
www.caida.org

Eliminating false positives in iffinder

- Iffinder: UDP probe to address A, response from B → A,B aliases.
- More vantage points reveal more pairs.
- A -> B: probe was sent to A, and the response came from B.
- graph suggests likely clusters (routers) and bridges (false)
- MIDAR confirms this intuition:
  - red: not aliases per MIDAR
  - green: are aliases per MIDAR
  - grey: MIDAR can't test
- other meta-data (e.g. DNS) helps validate

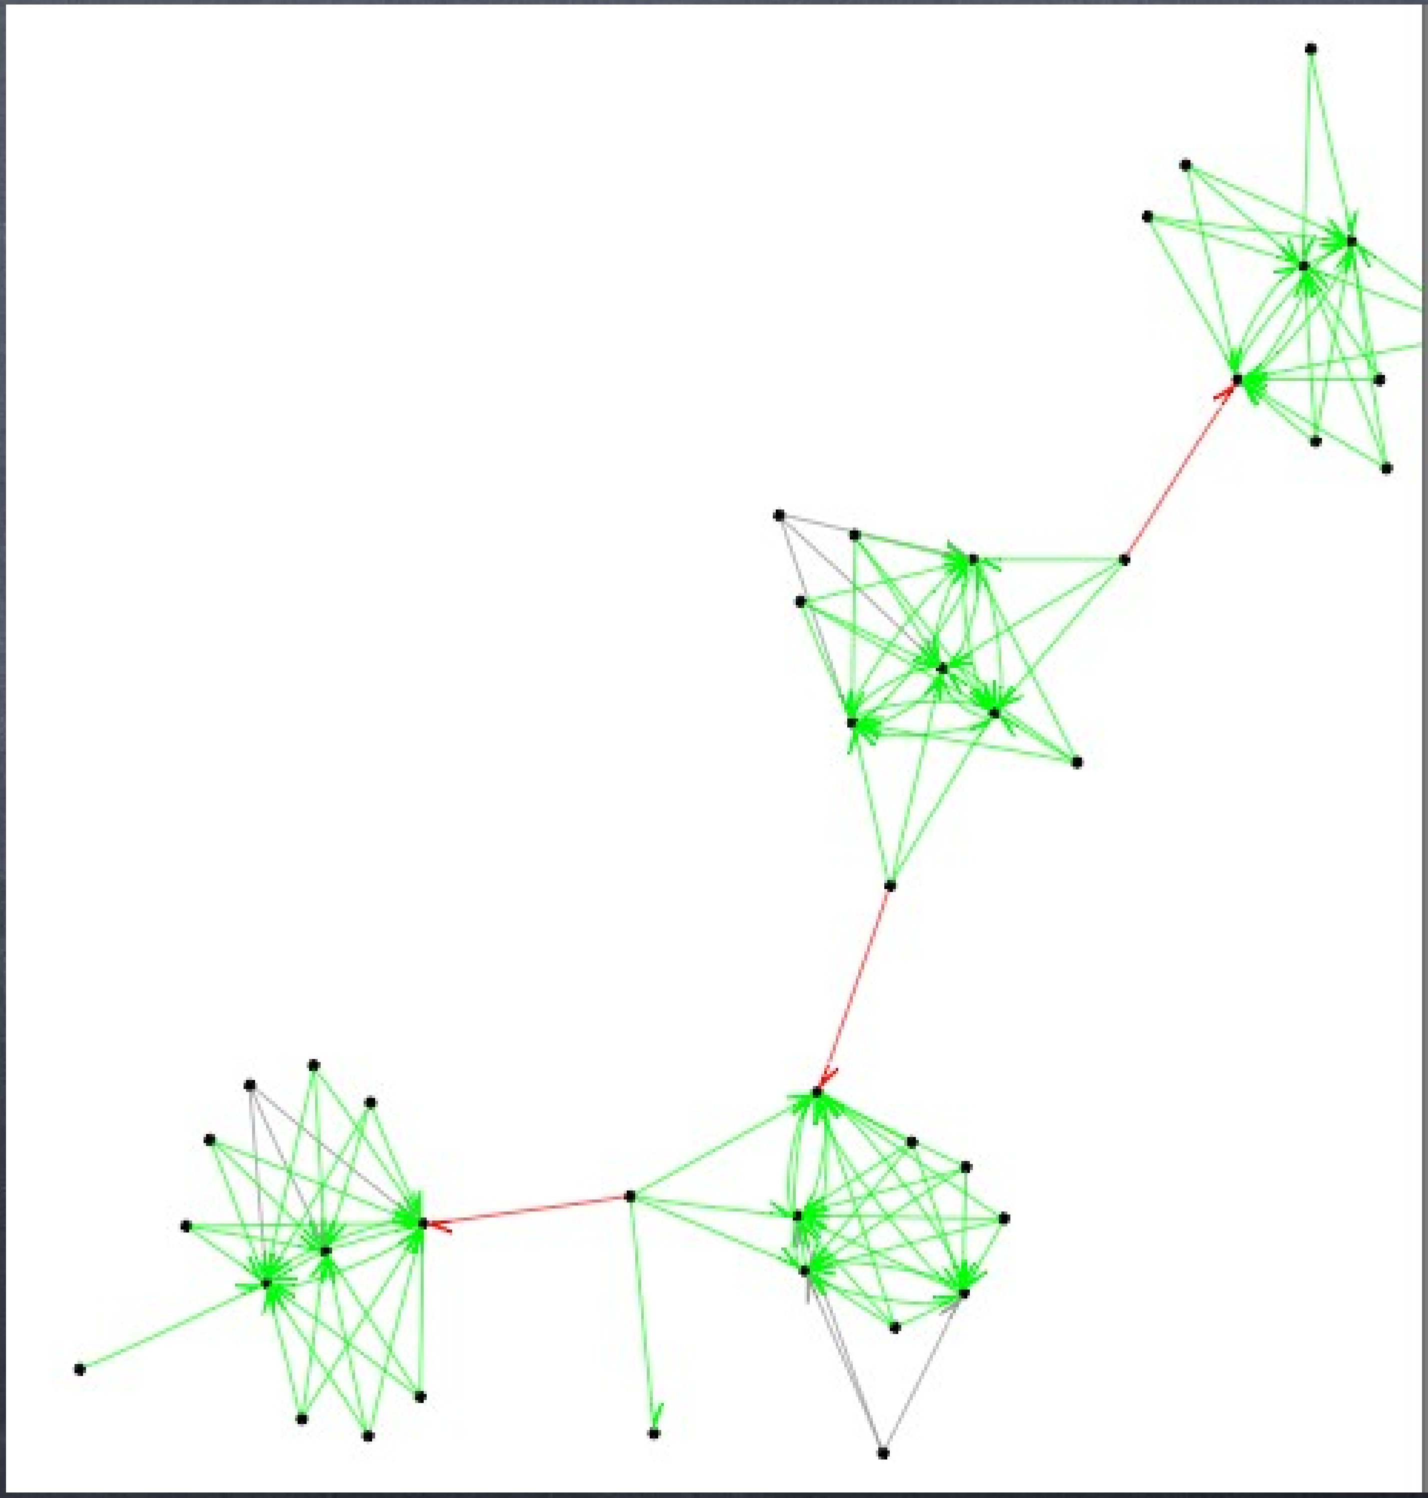

Eliminating false positives in iffinder

- To identify false positives [without MIDAR], we apply the "responder cluster" algorithm.
- Clusters of addresses that respond from each other (blue) → assume routers
- For other addresses, if its responders share a cluster (green) → cluster.
- If responders in different clusters (red), separate
- All "bridges" that are false according to both intuition and MIDAR are also classified as false by this algorithm.
- Conservative (discards aliases), which yields combined MIDAR+iffinder "high-accuracy" router-level graph

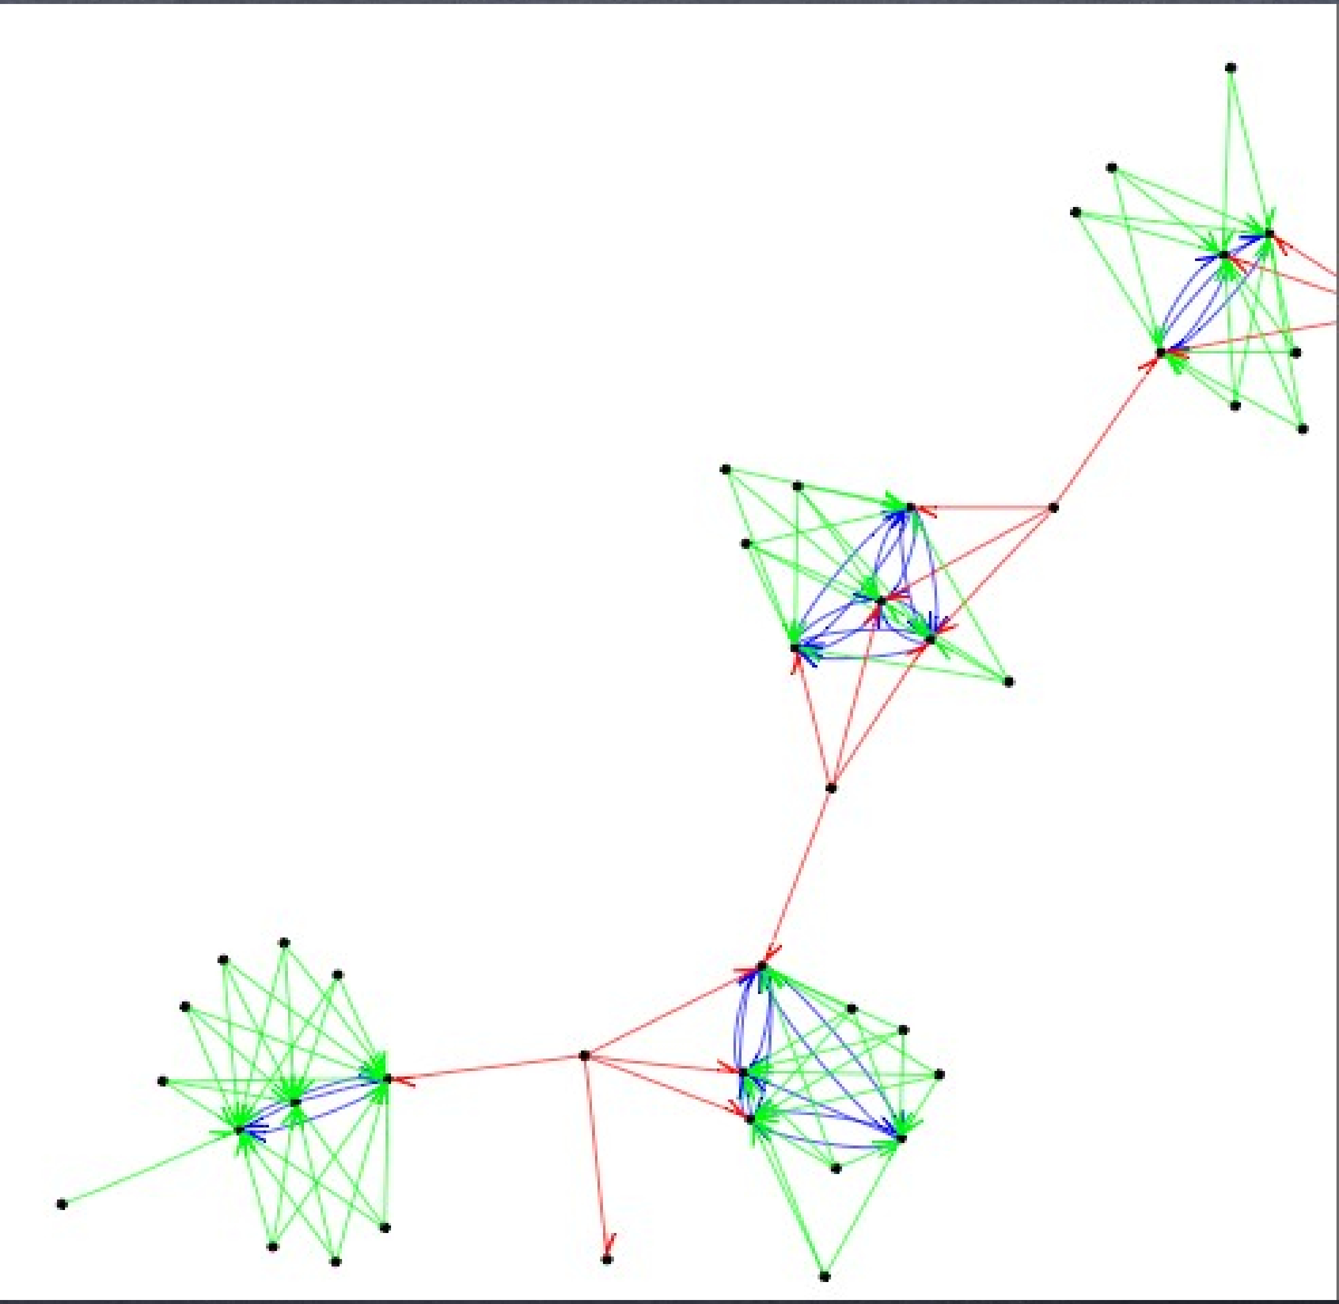

Statistics Pages

- per-monitor analysis of IPv4 topology data

www.caida.org/projects/ark/statistics/

AS dispersion by AS hop

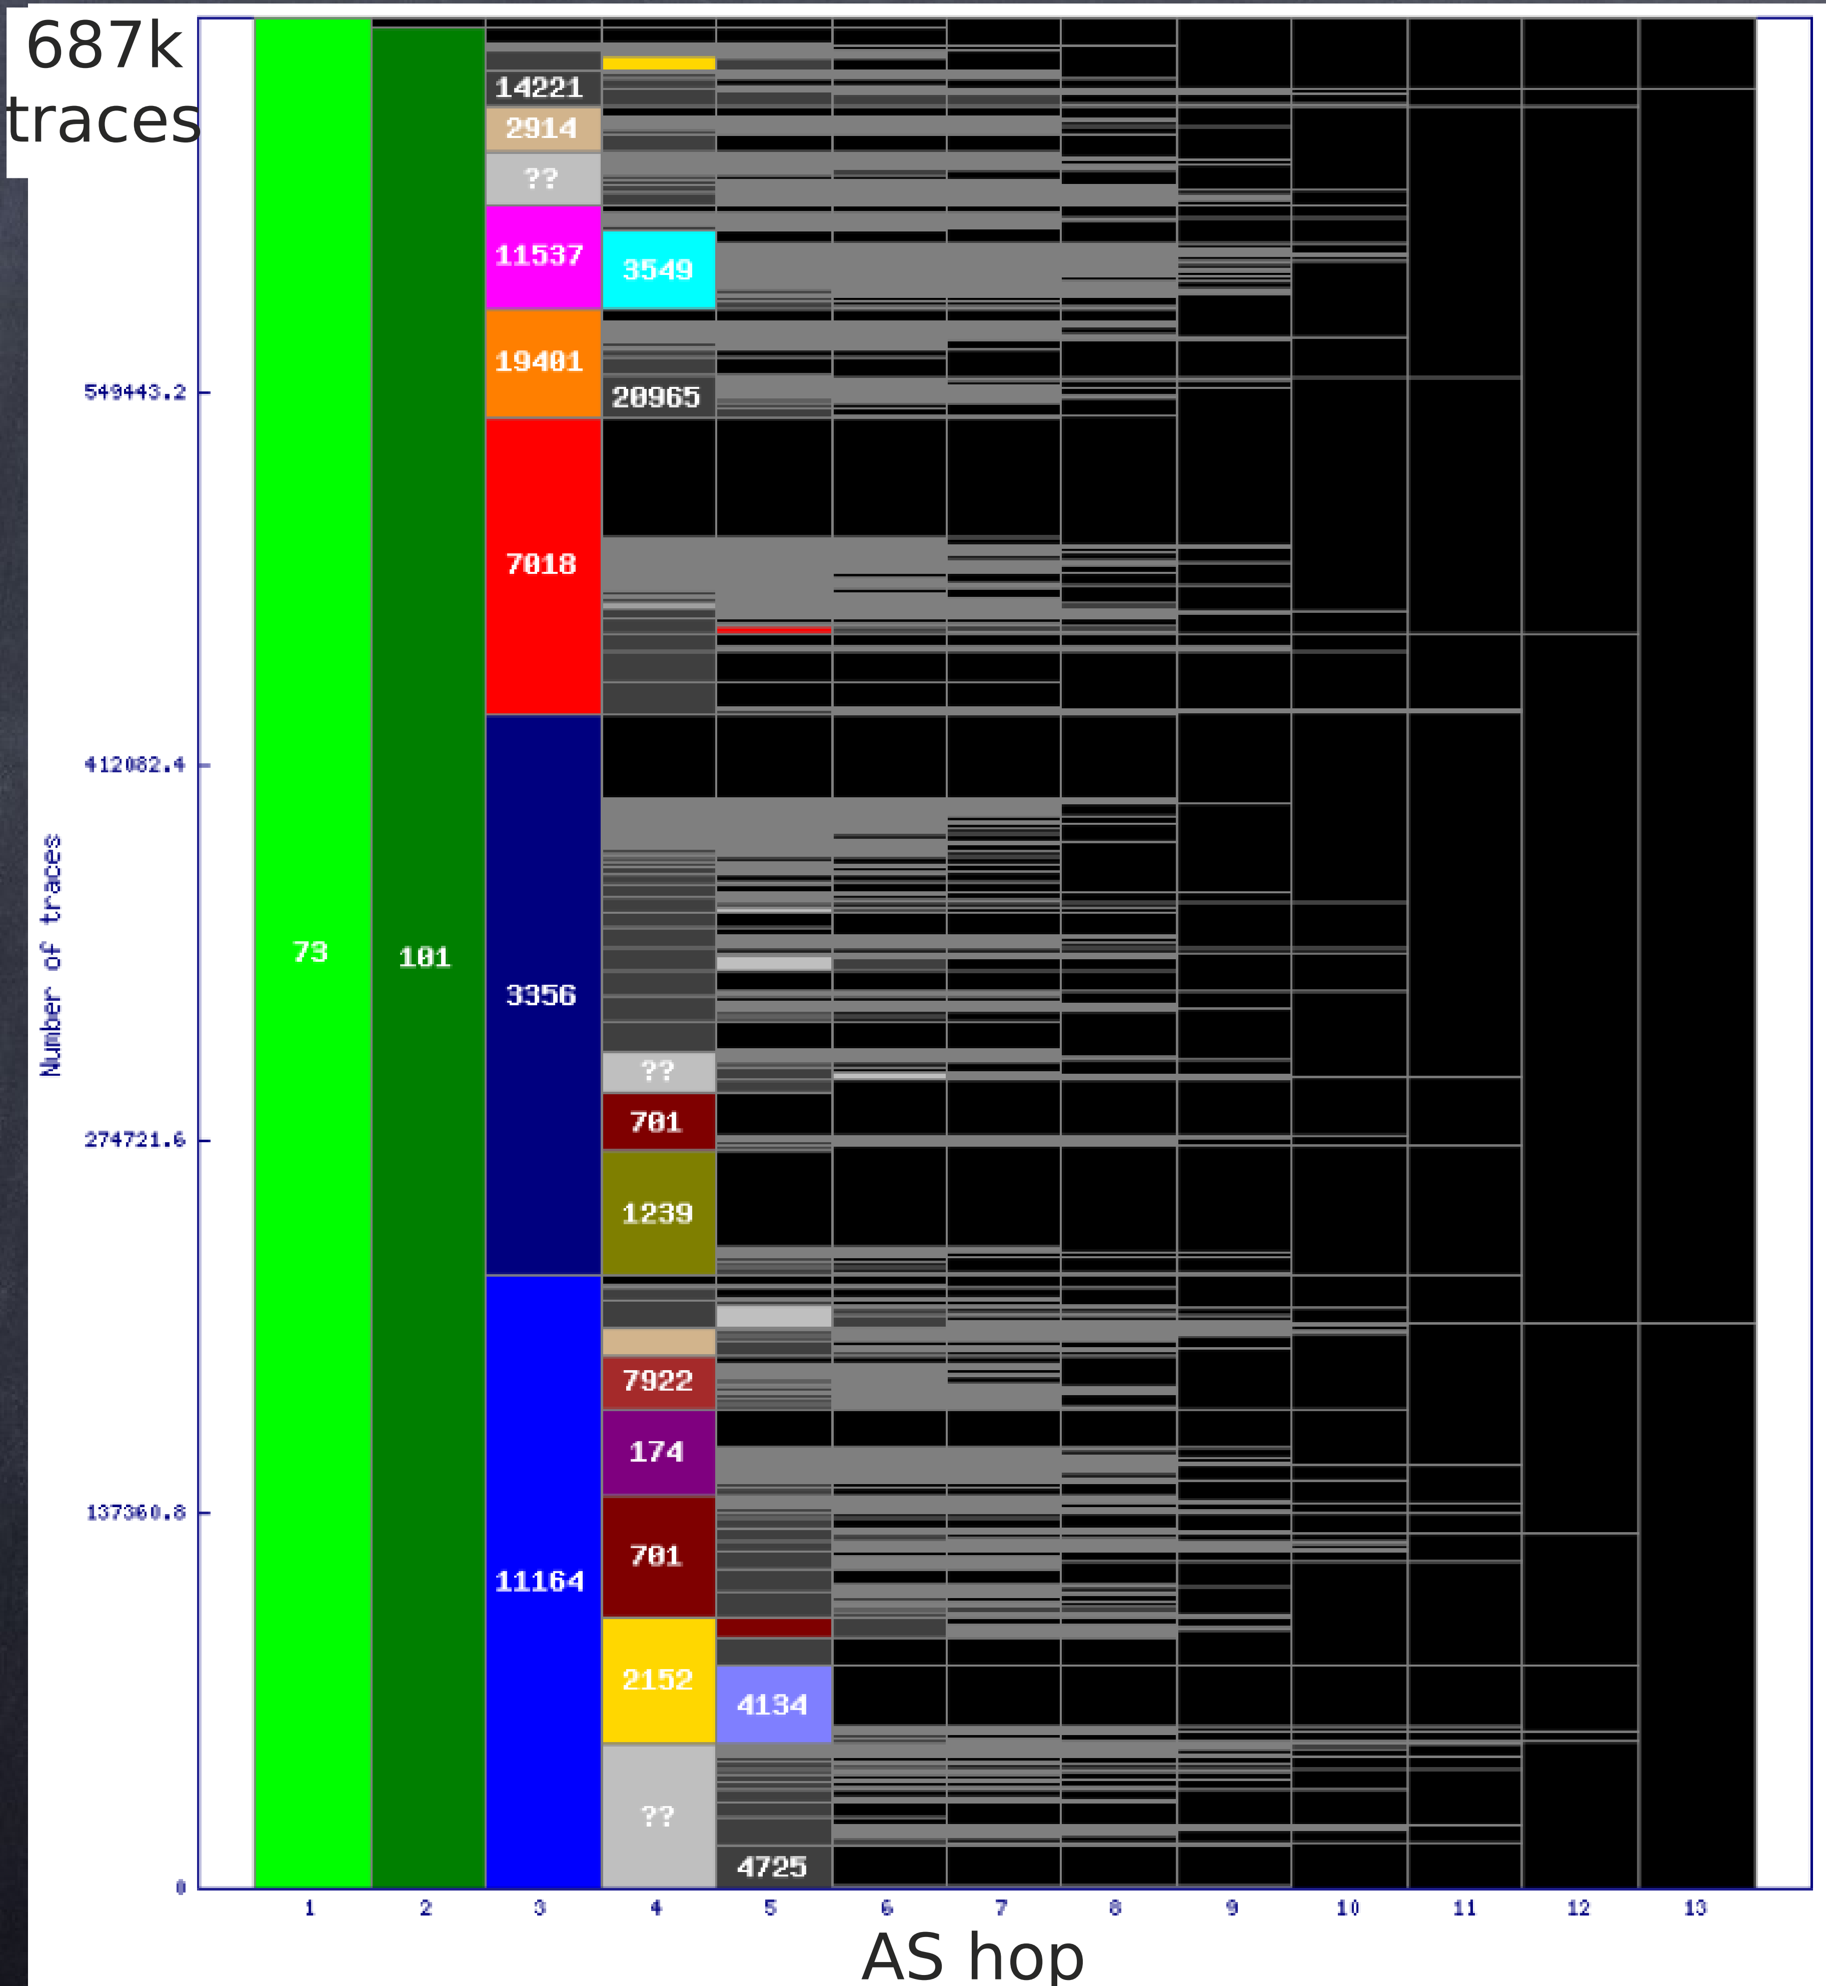

73	WASHINGTON-AS - University of Wash
101	WASH-NSF-AS - University of Washingt
11164	TRANSITRAIL - National LambdaRail, L
3356	LEVEL3 Level 3 Communications
7018	ATT-INTERNET4 - AT&T WorldNet Serv
701	UUNET - MCI Communications Service
2152	CSUNET-NW - California State Universi
1239	SPRINTLINK - Sprint
19401	NLR - National LambdaRail
11537	ABILENE - Internet2
174	COGENT Cogent/PSI
4134	CHINANET-BACKBONE No.31,Jin-rong
3549	GBLX Global Crossing Ltd.
2914	NTT-COMMUNICATIONS-2914 - NTT A
7922	COMCAST-7922 - Comcast Cable Com
20965	GEANT The GEANT IP Service
4725	ODN SOFTBANK TELECOM Corp.
14221	WASHINGHTON-RD-AS - University of

sea-us monitor

AS dispersion by IP hop

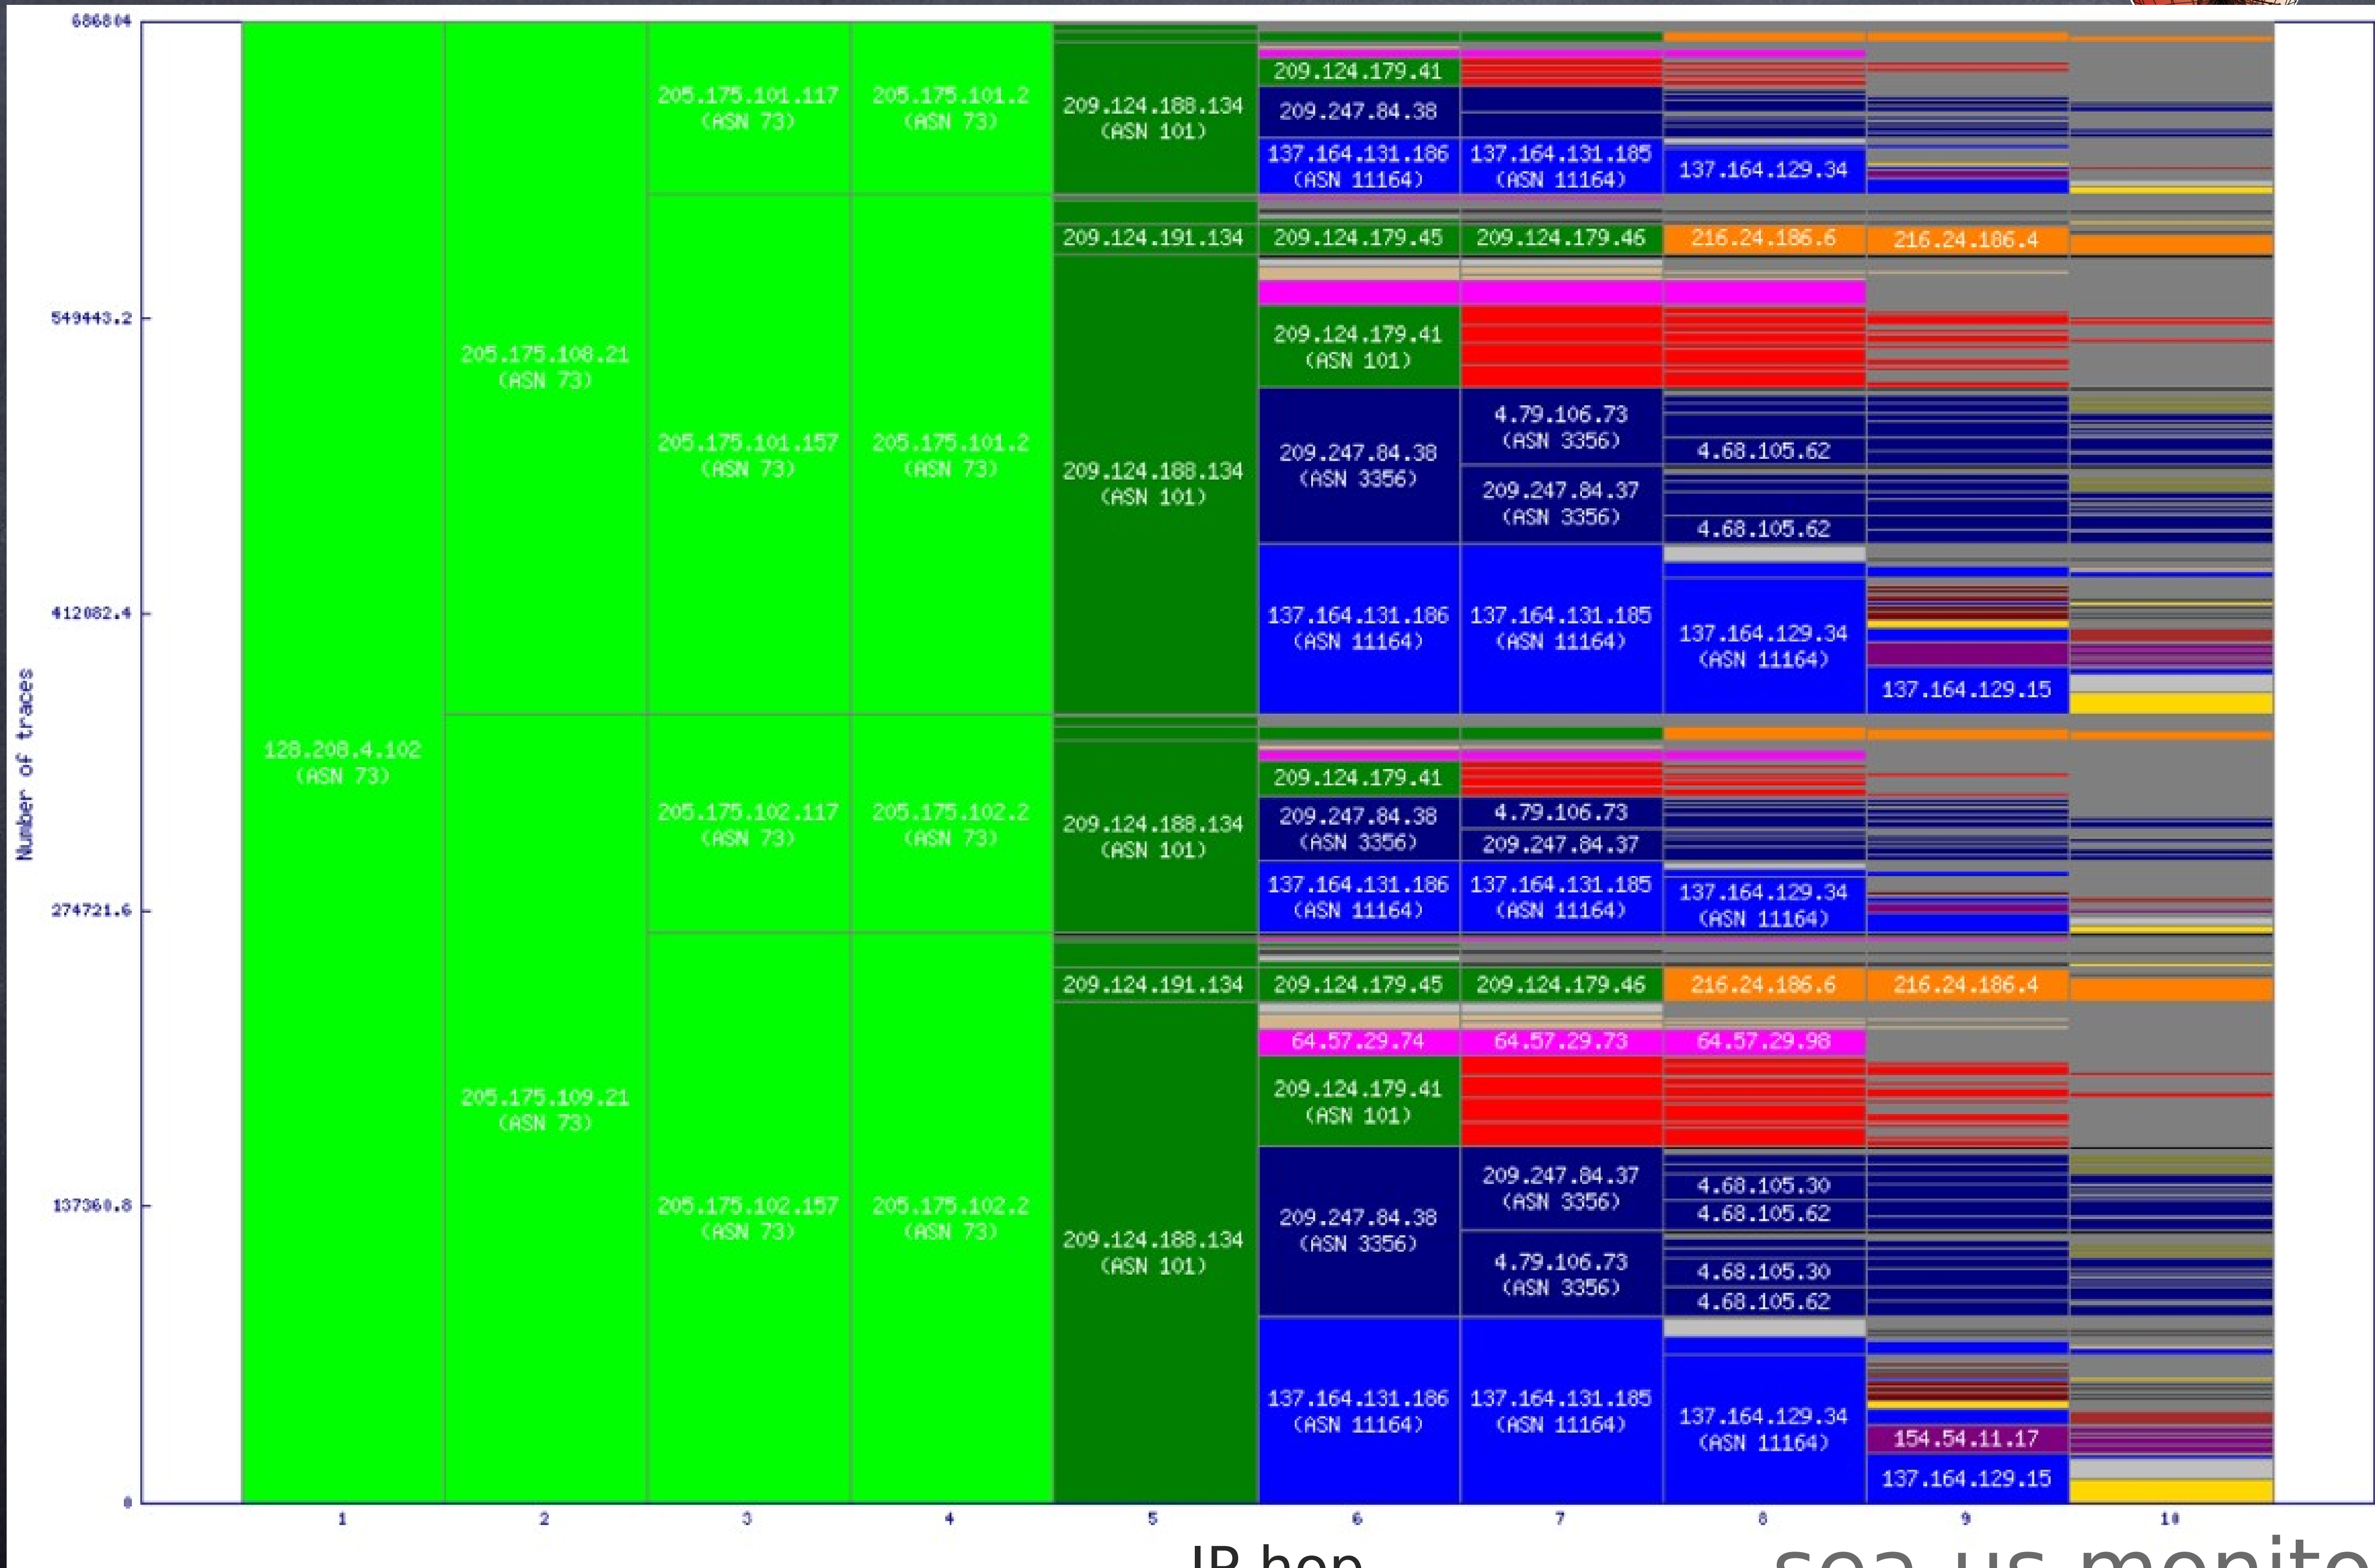

IP hop

sea-us monitor

AS dispersion by IP hop: see load balancing

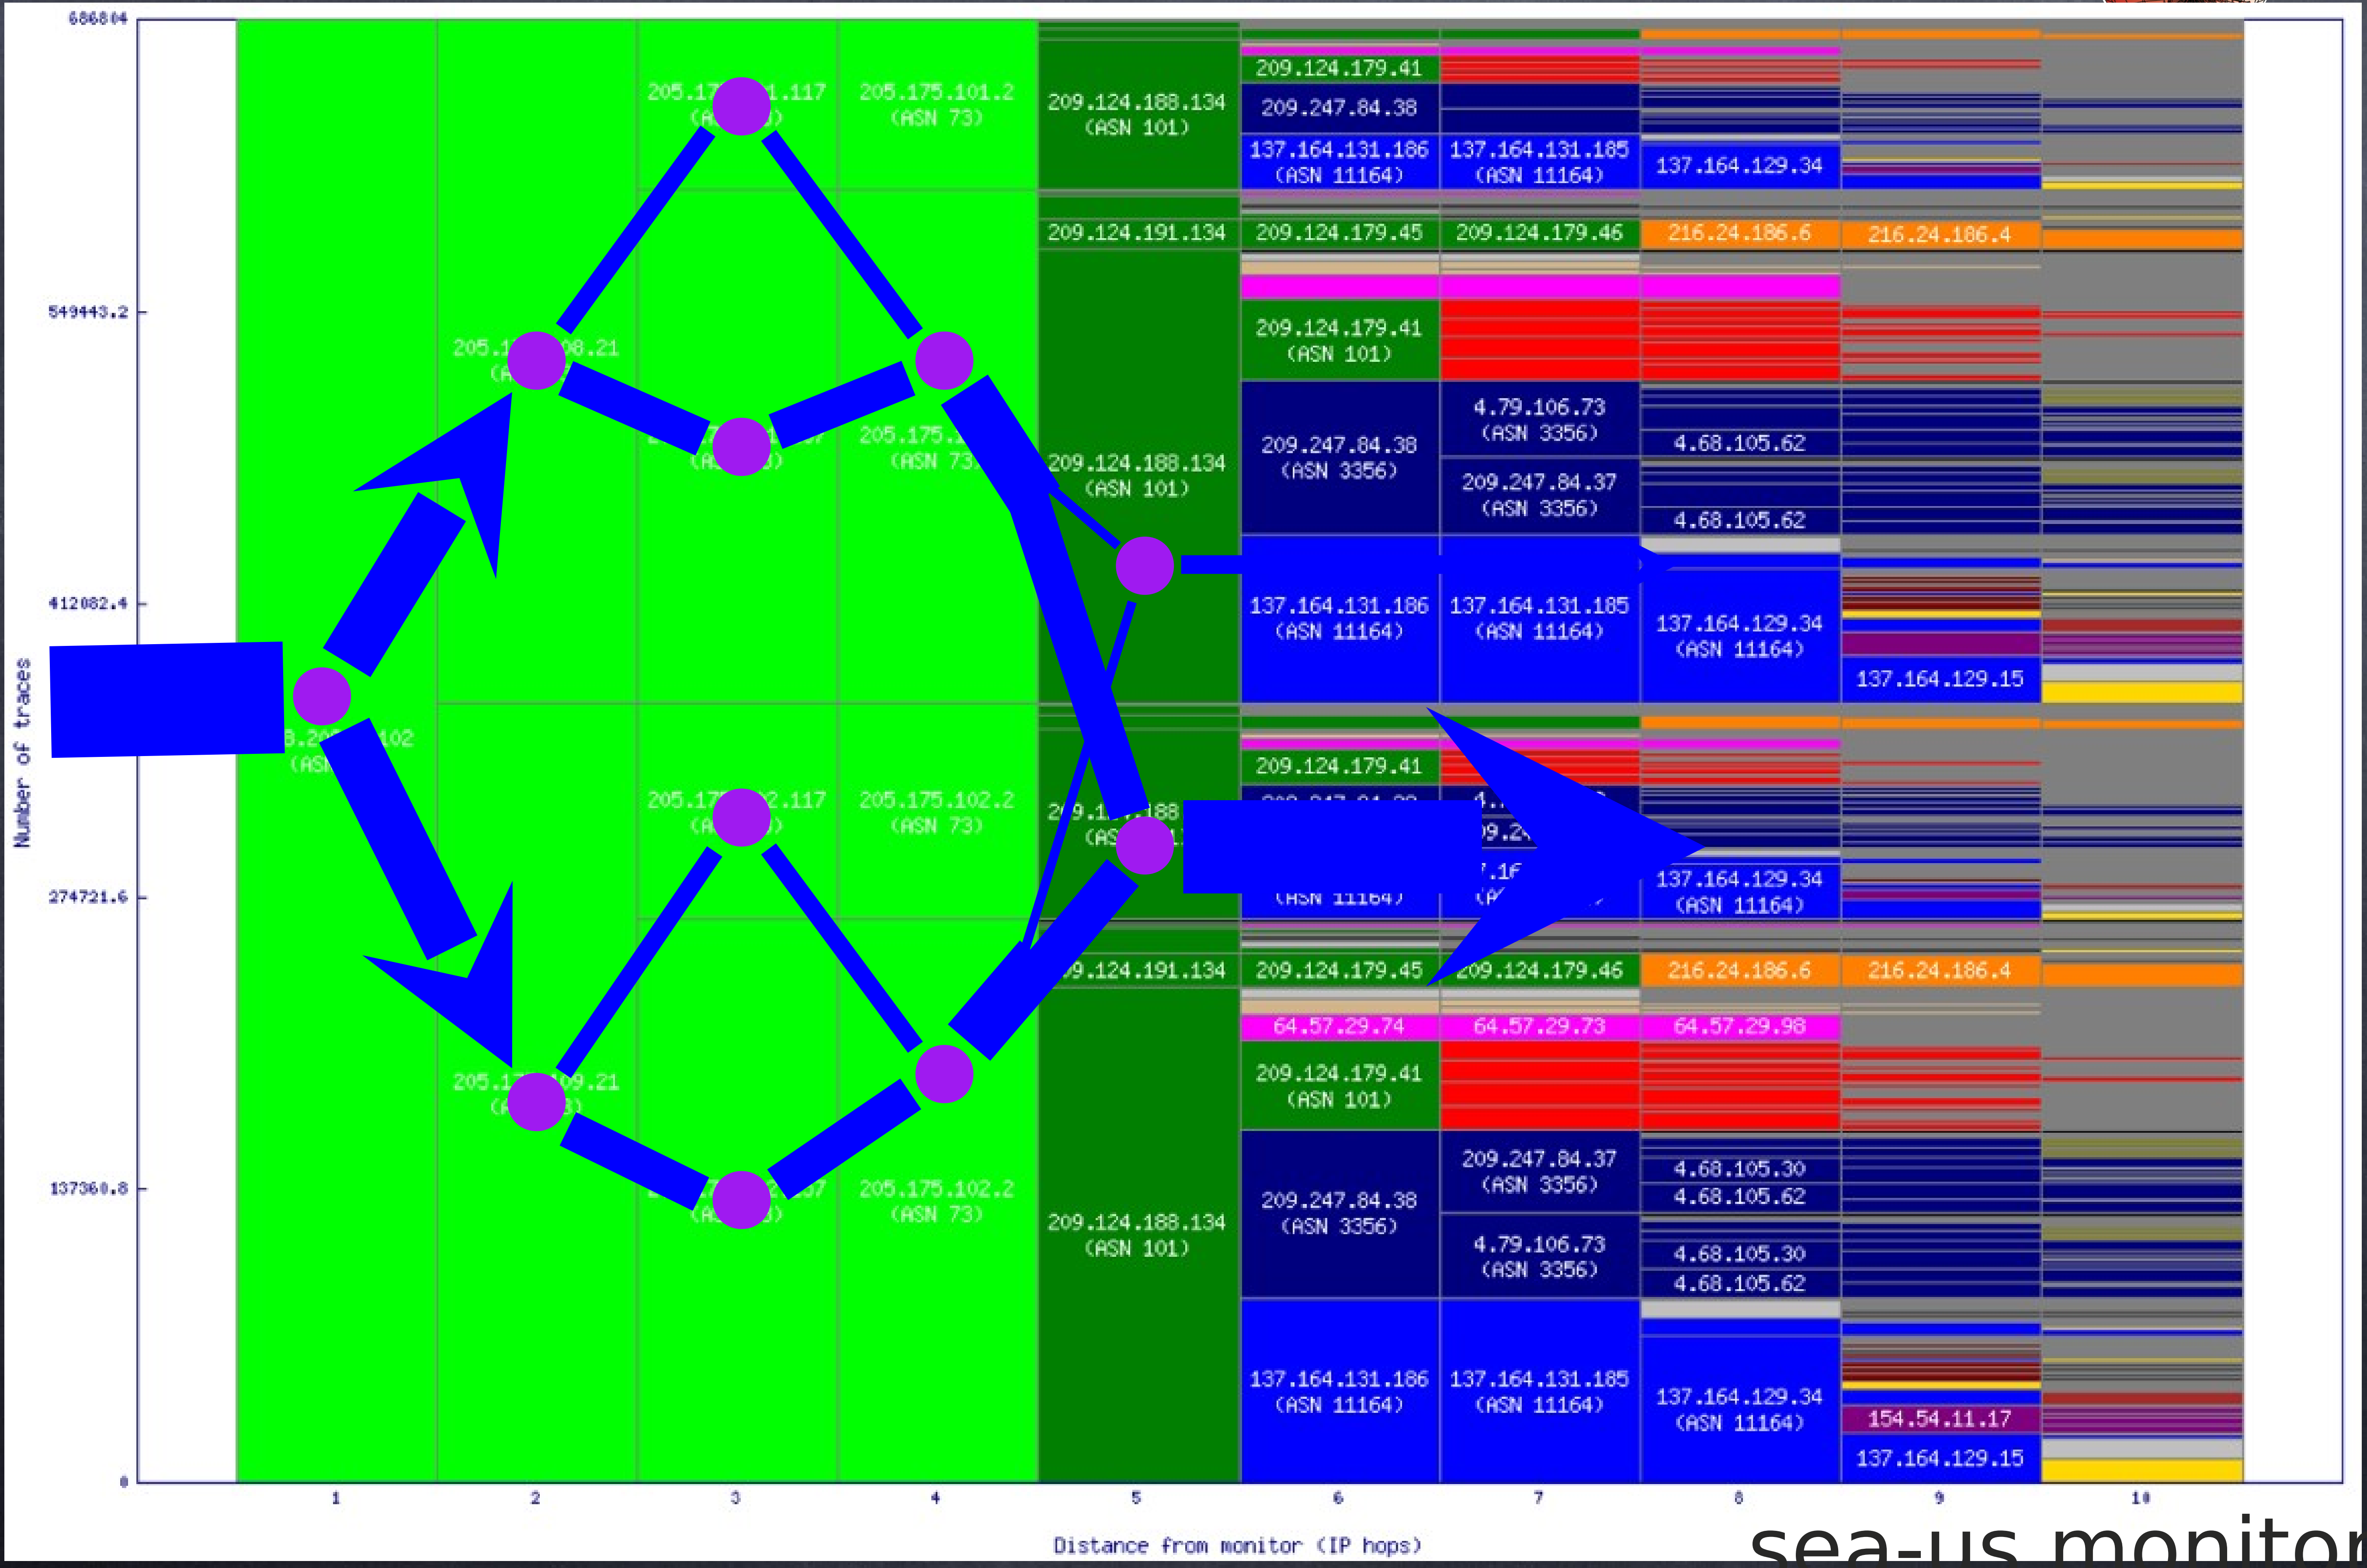

sea-us monitor

Statistics Pages

- work in progress: RTT plotted by country
  - geolocate destinations with NetAcuity
  - color each country by median RTT of destinations

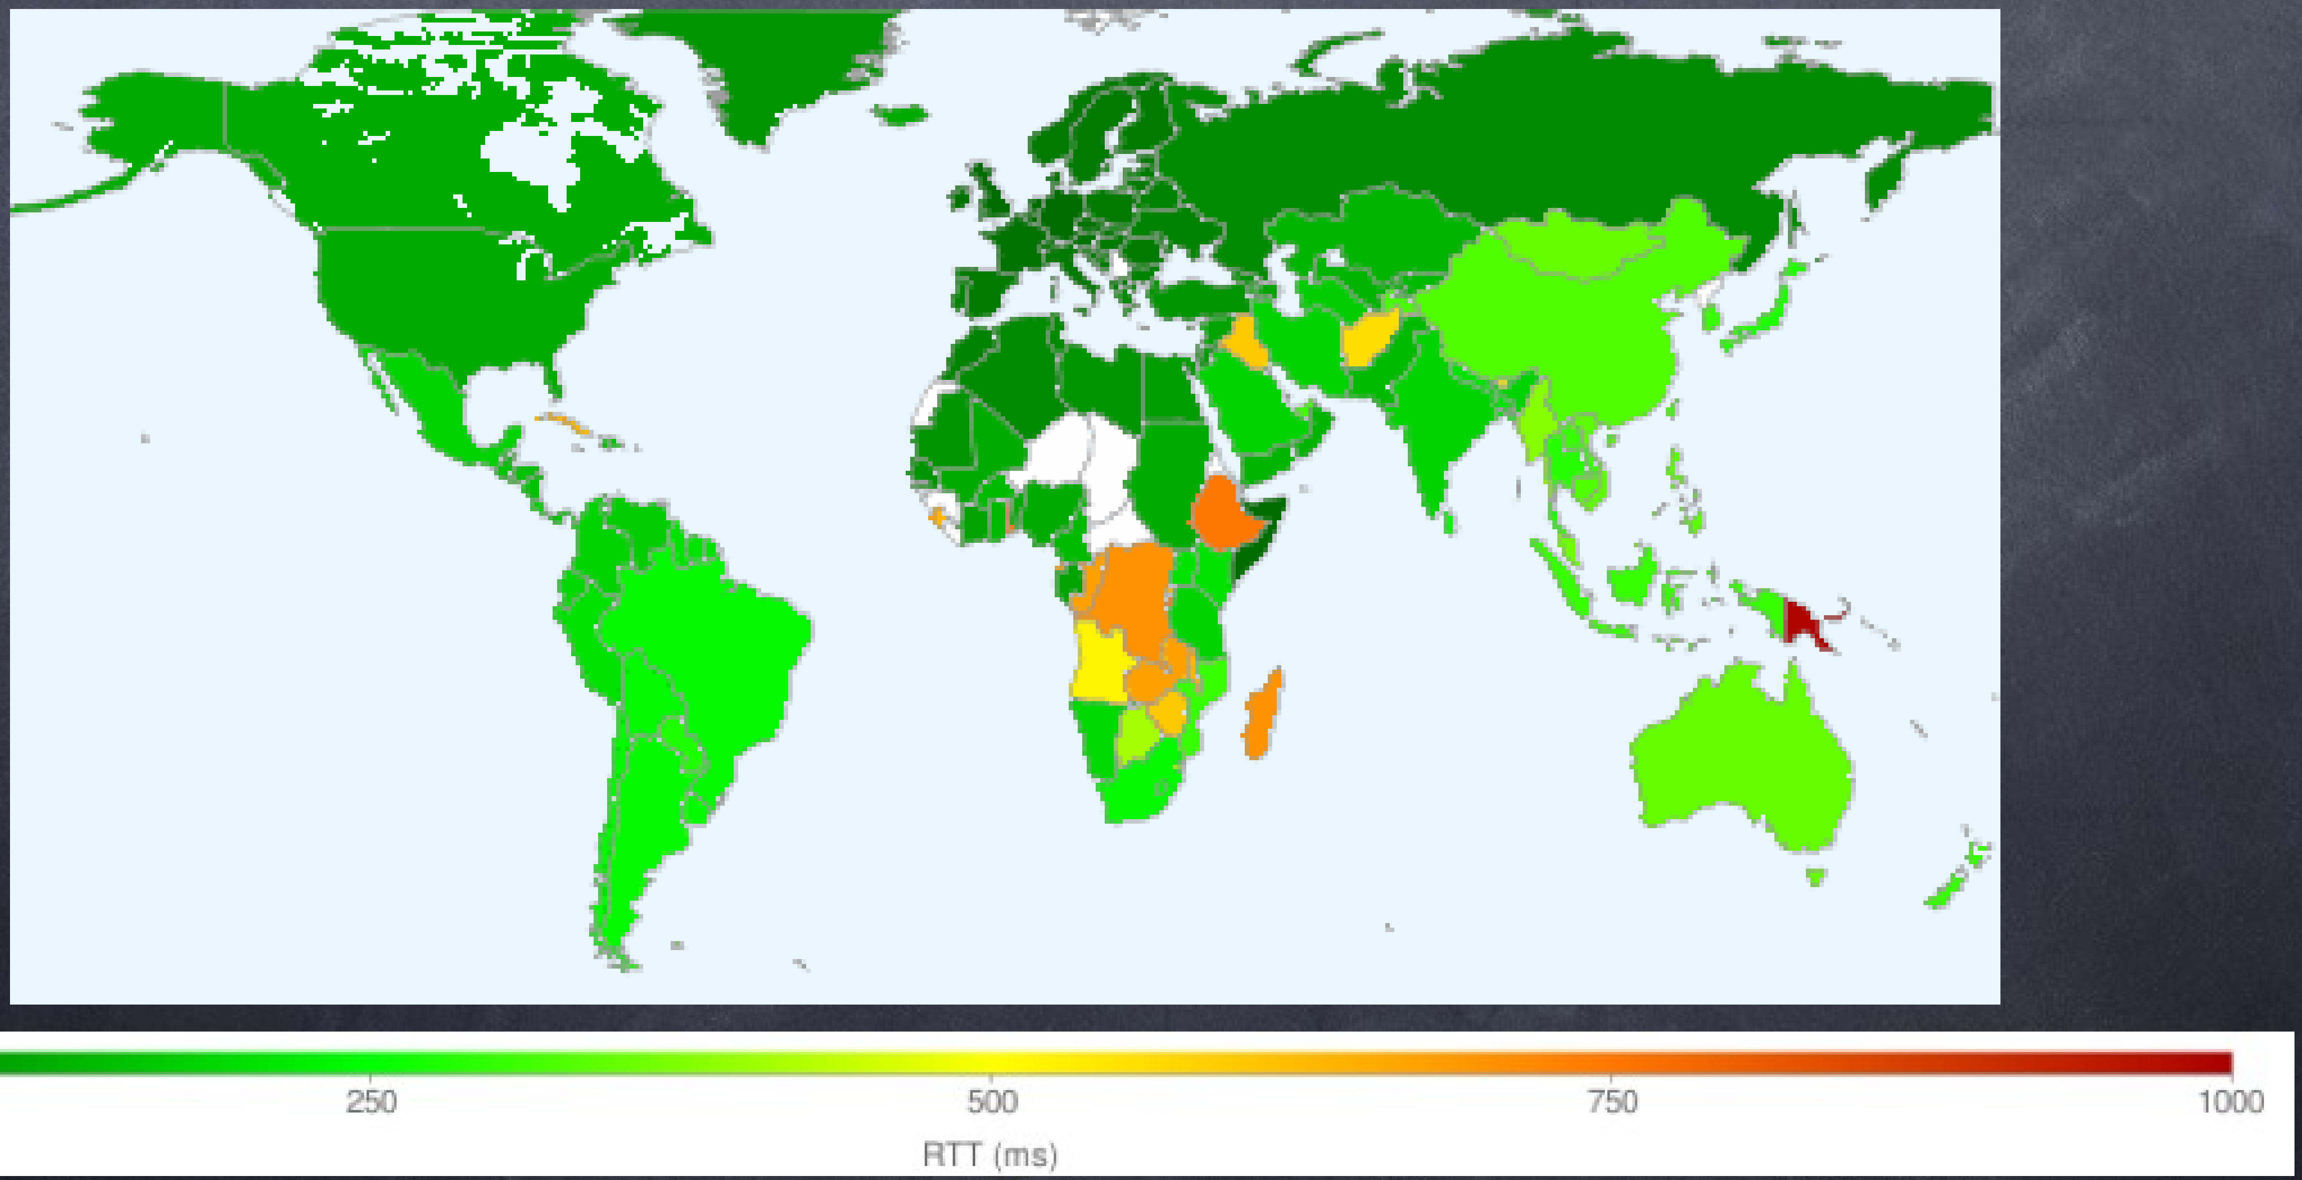

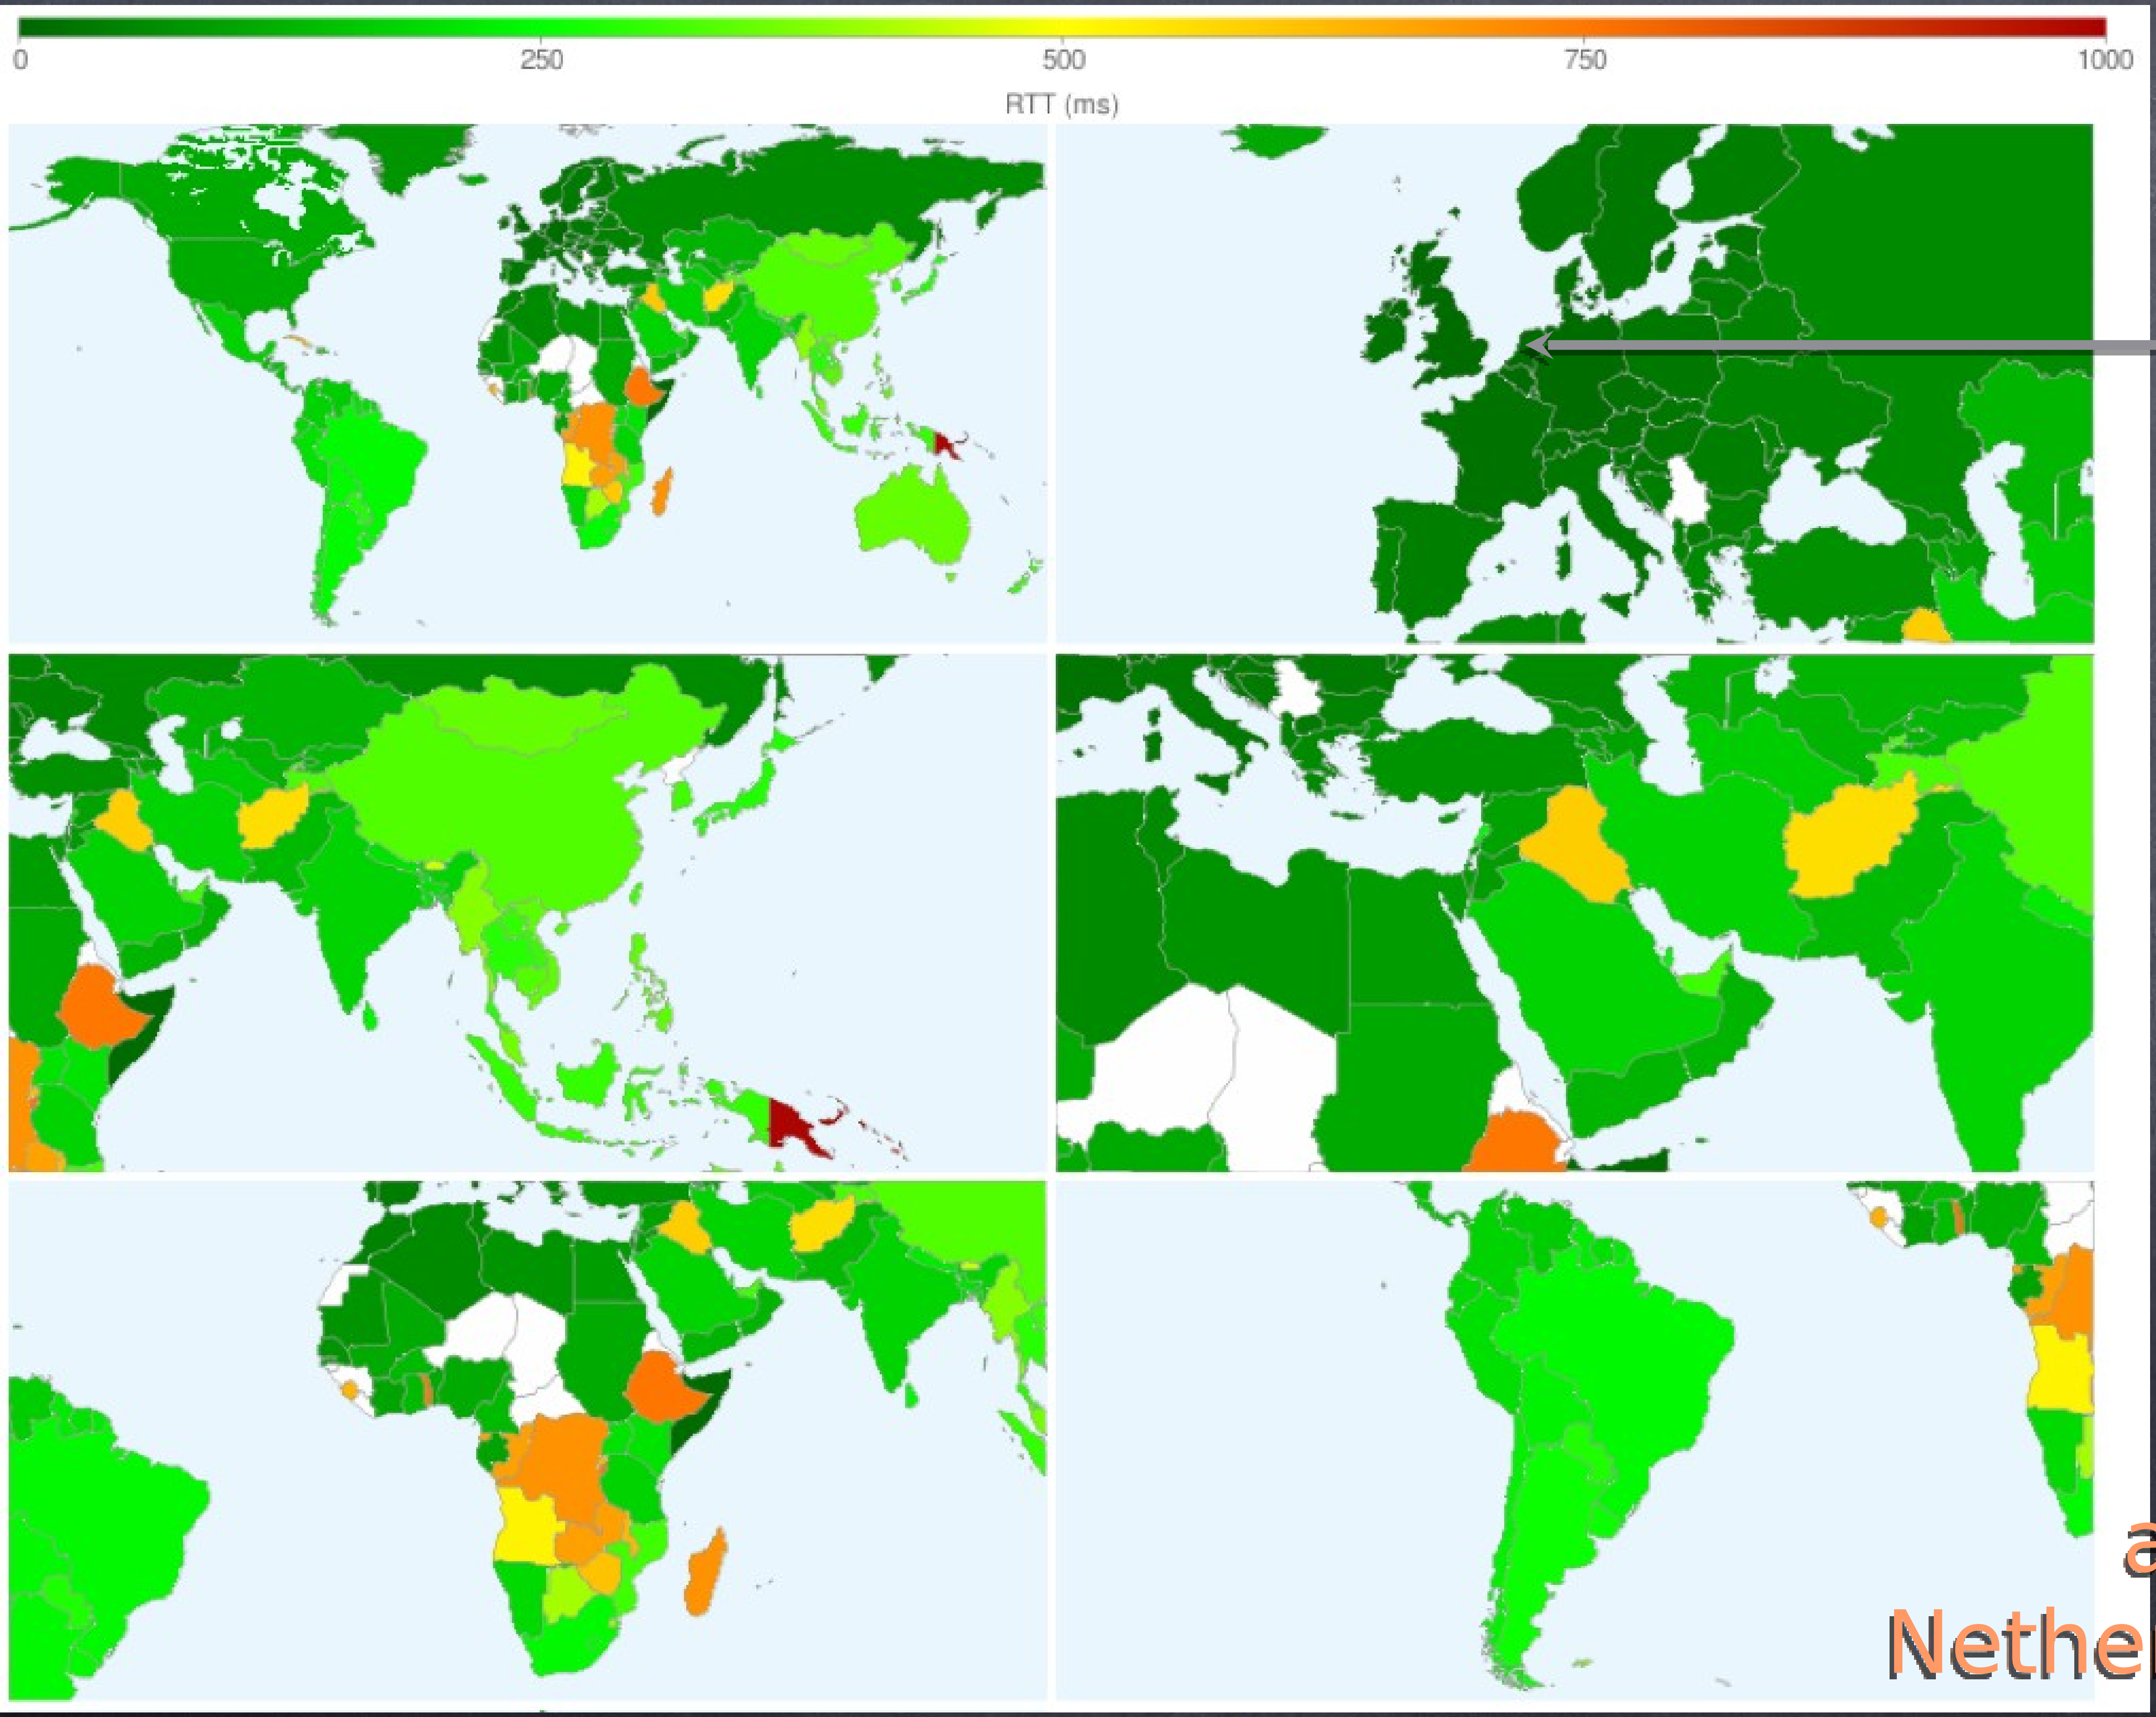

view
from
ams-nl
Netherlands

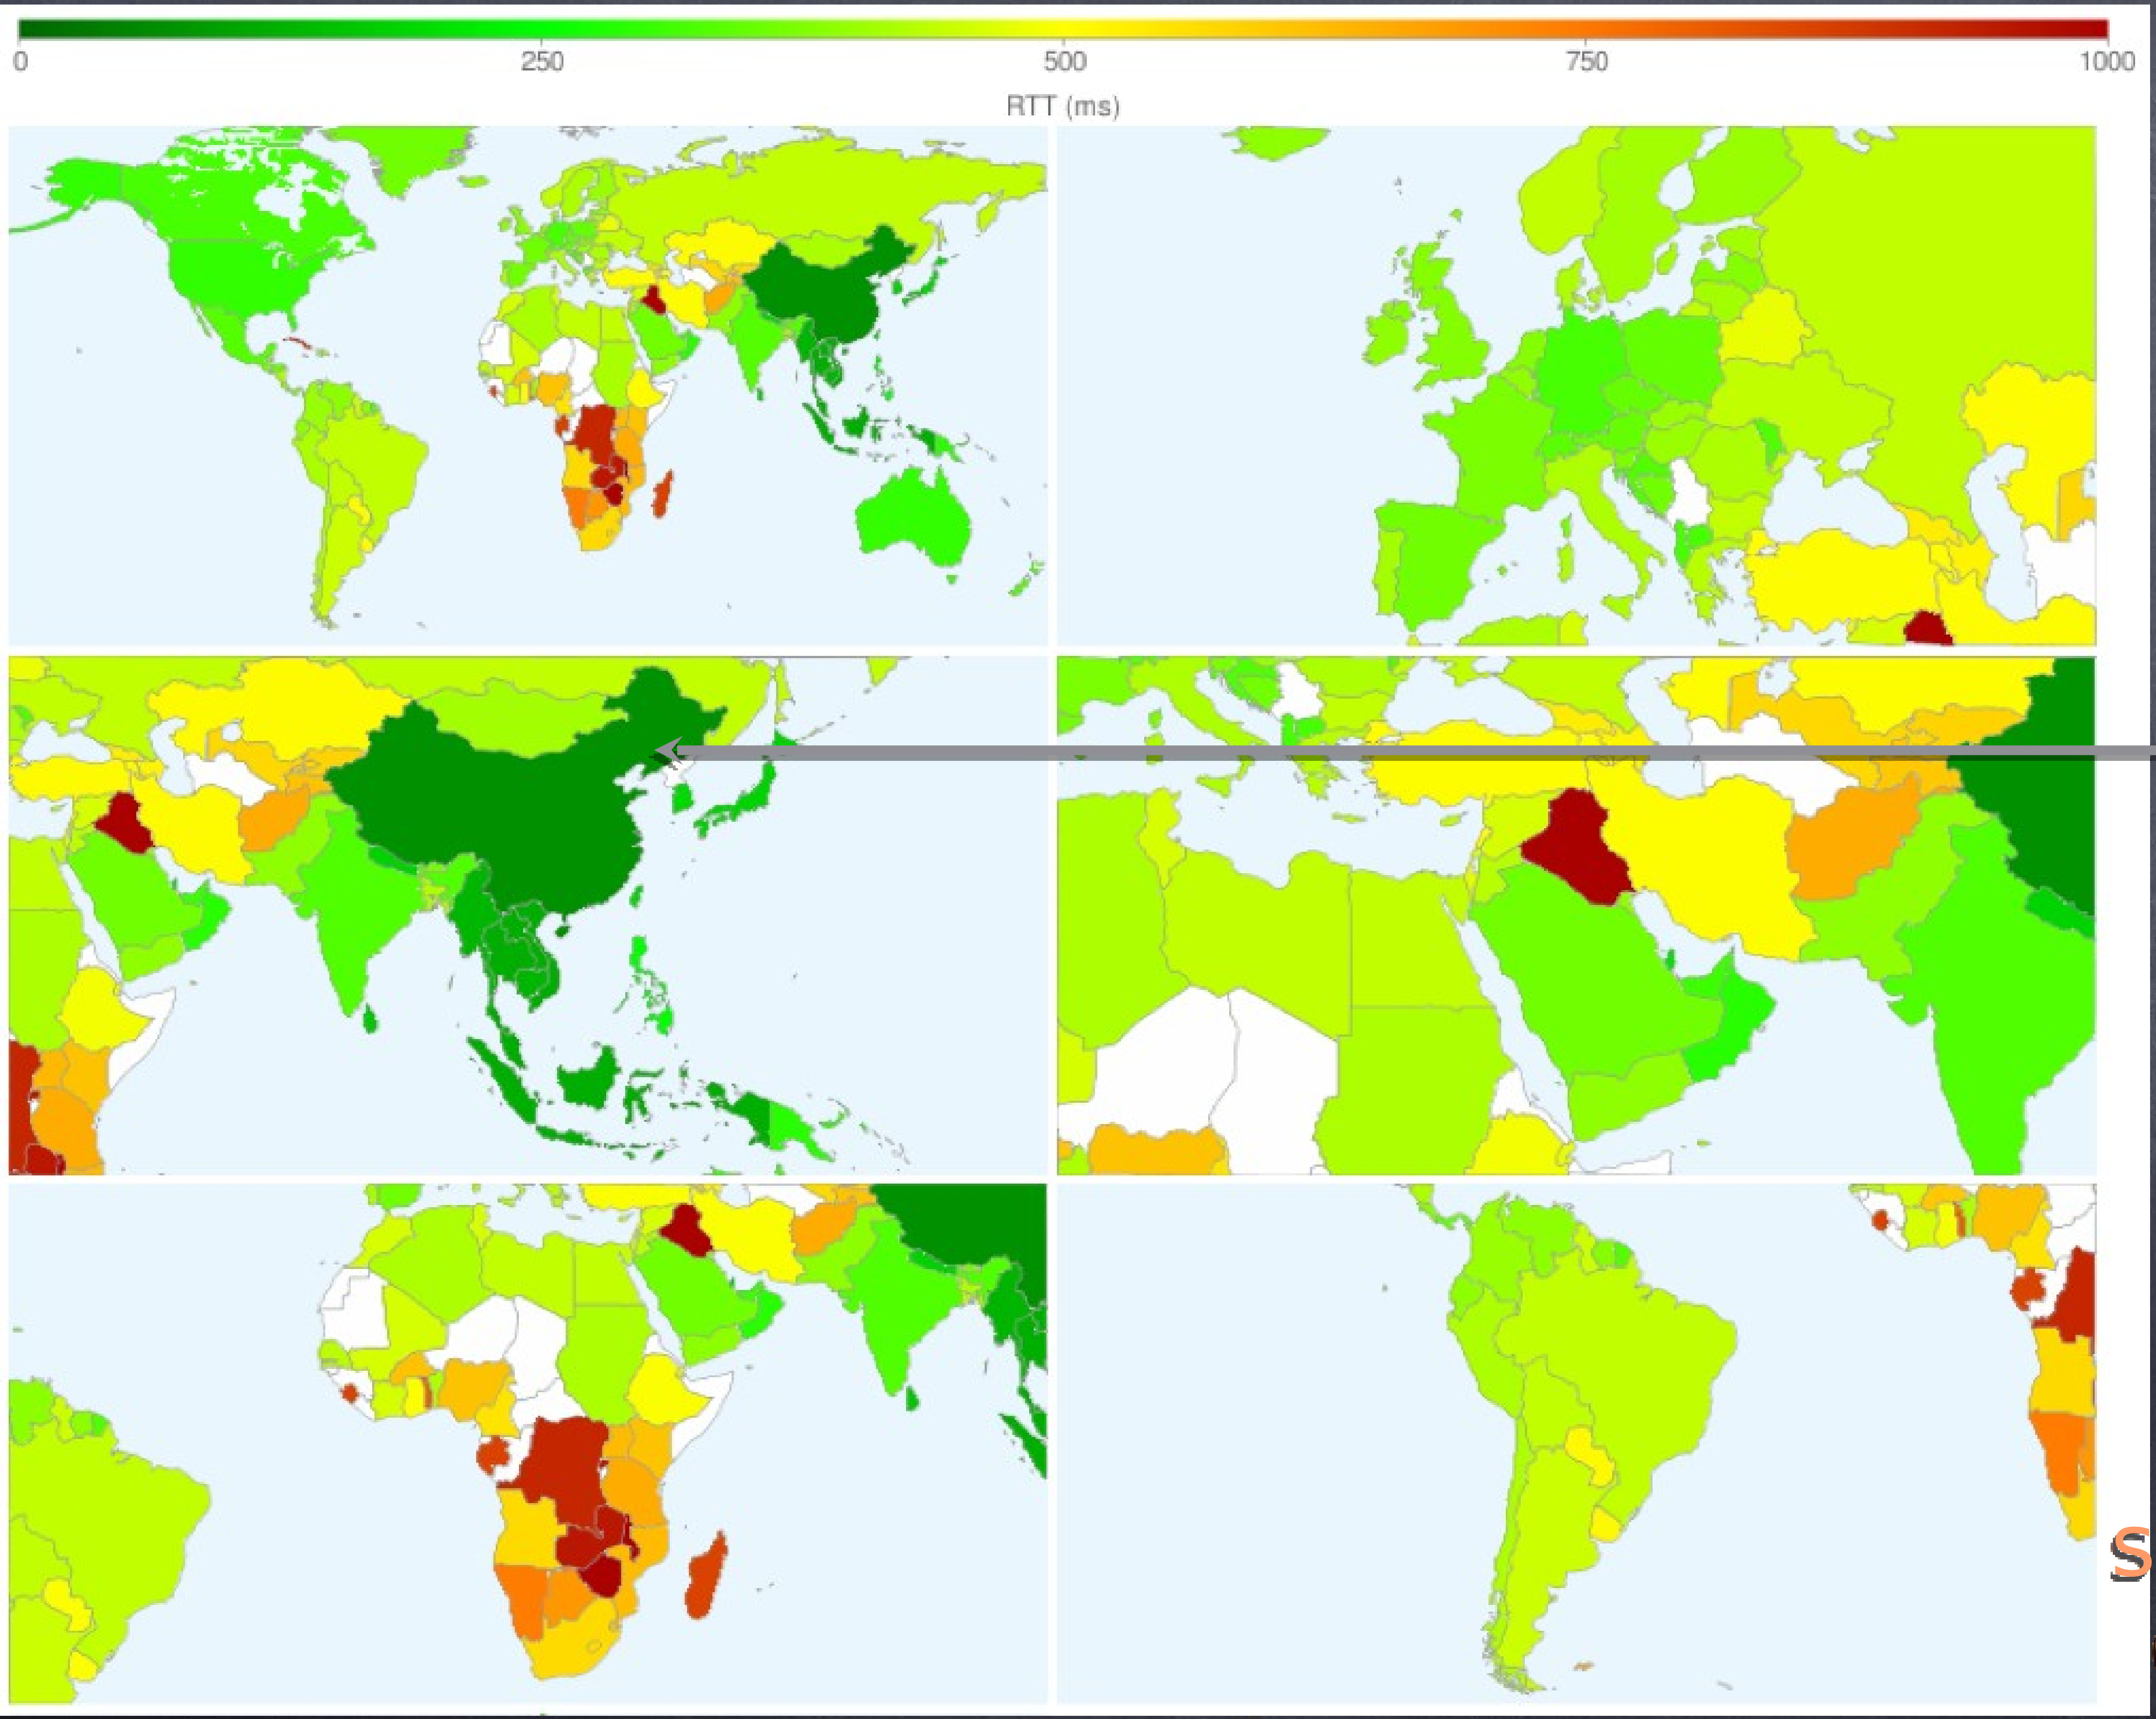

view
from
she-cn
China

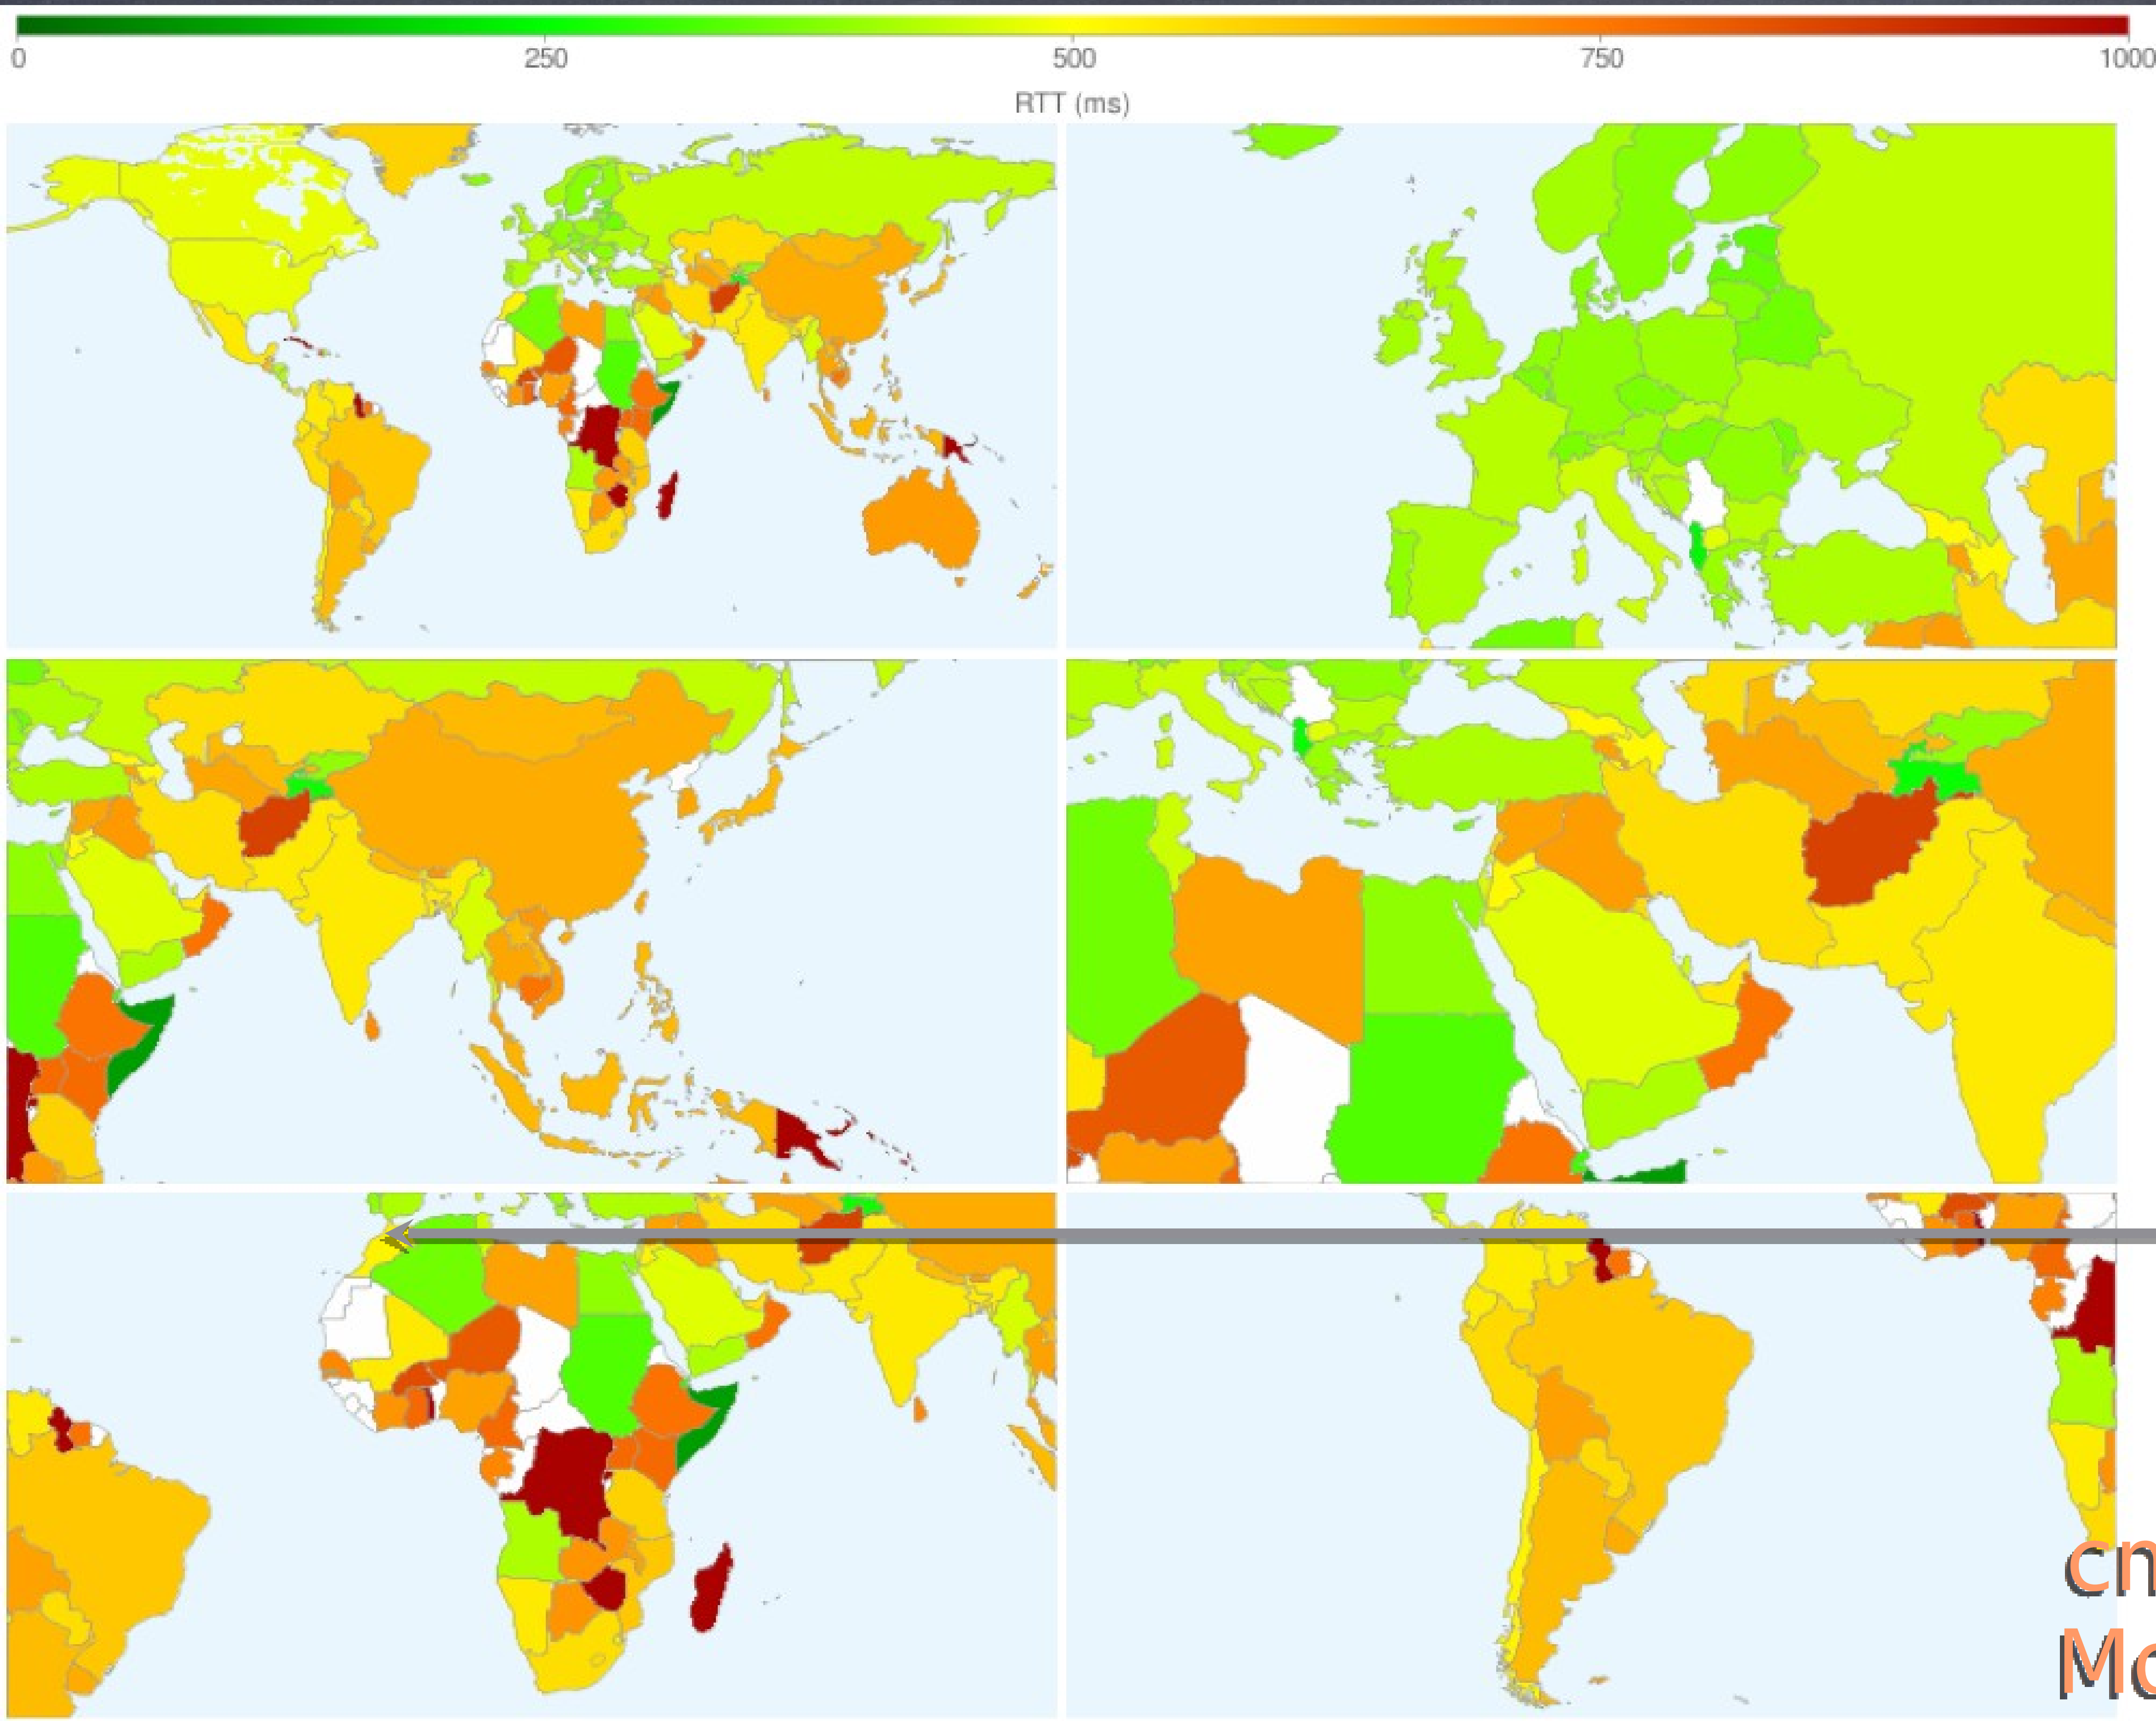

view from cmn-ma Morocco

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

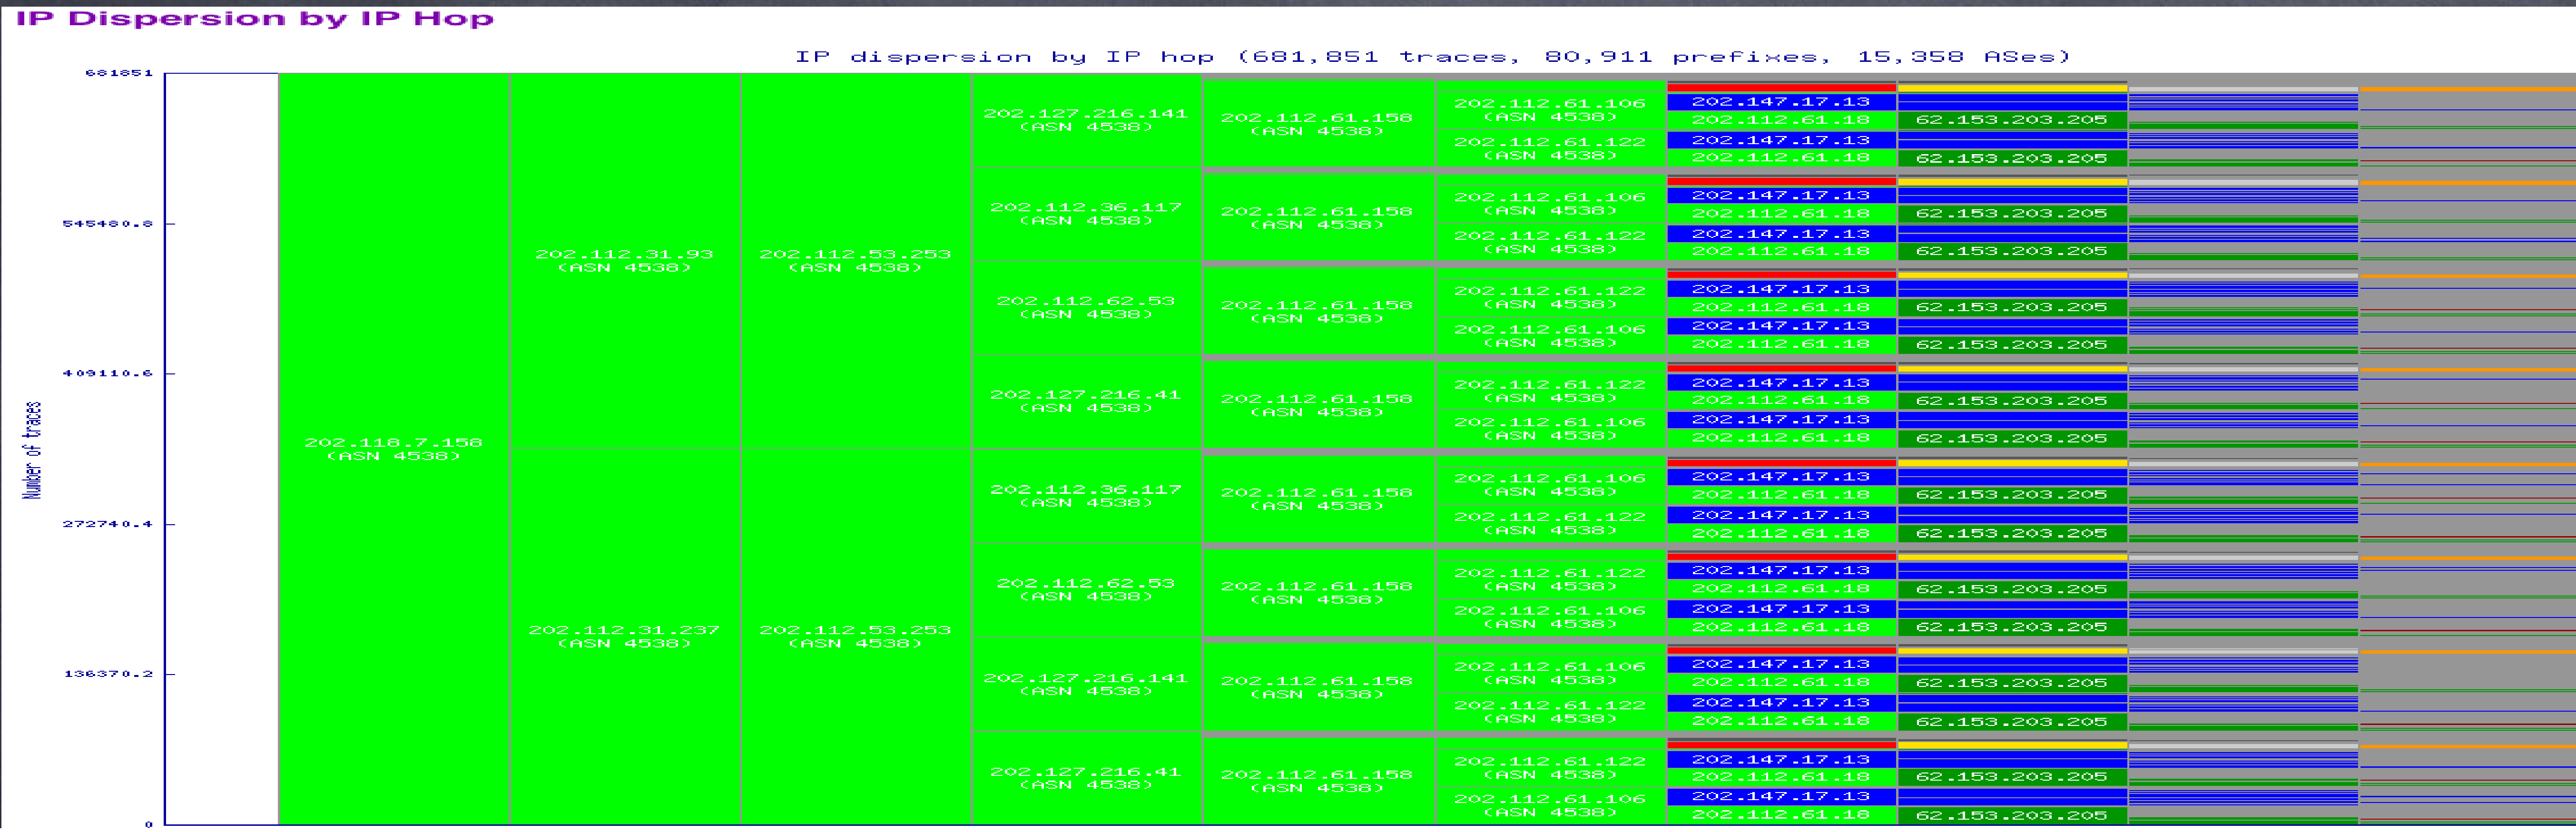

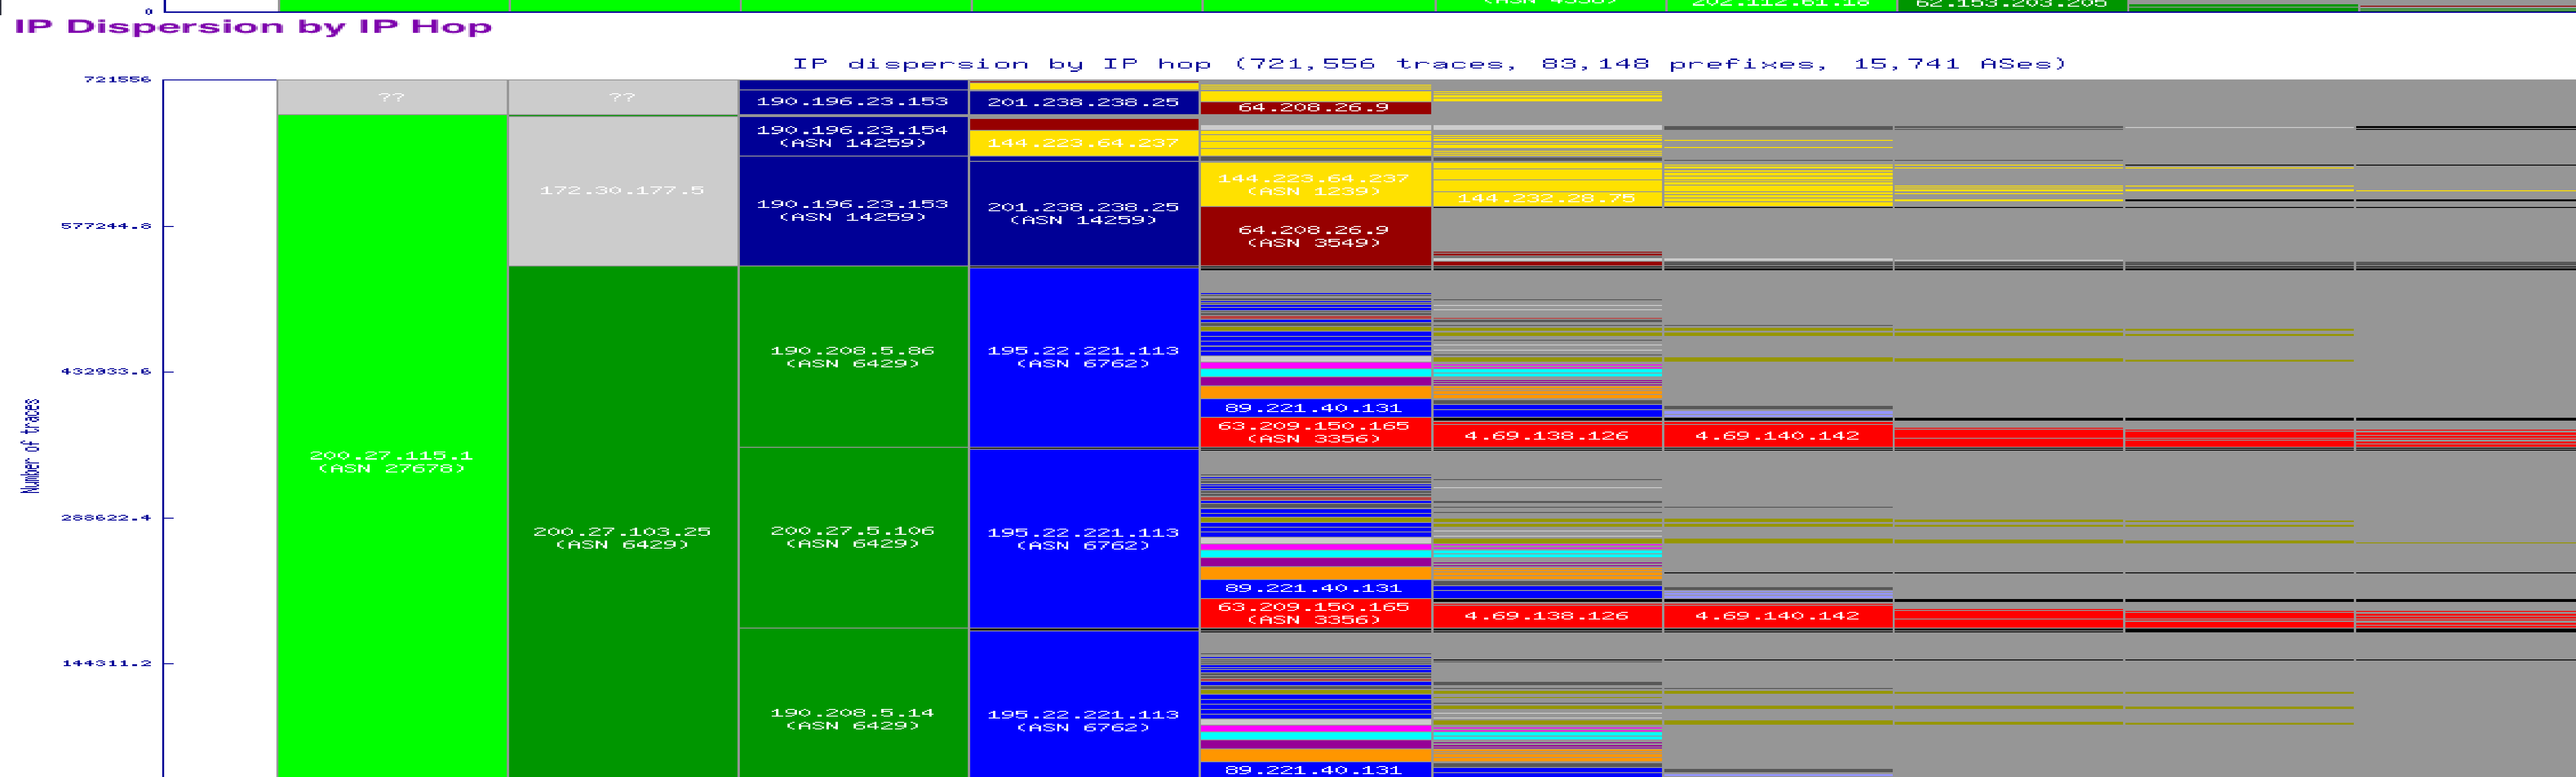

AS Rank

- Autonomous Systems rank by "customer cone"

rank	AS number	AS name	customer cone: customer cone size	customer cone: percentage of all ASes	AS degree
1	3356	LEVEL3 Level 3 Commu	31112	92%	2632
2	7018	AT&T WorldNet Servic	29978	89%	2283
3	701	MCI Communications S	29820	88%	2066
4	174	Cogent/PSI	29328	87%	2533
5	3549	Global Crossing Ltd.	29035	86%	1365
6	1239	Sprint	29012	86%	1381
7	209	Qwest Communications	28983	86%	1387
8	6939	Hurricane Electric,	27227	81%	1552
9	4323	tw telecom holdings,	27198	81%	1291
10	1299	TeliaNet Global Netw	27117	80%	561

data sources

country	ASN allocation	2010.04.22	IANA
	delegated	2010.08.19	AFRINIC,APNIC,ARIN,IANA,LACNIC,RIPENCC
	whois	2010.04.01	AFRINIC,APNIC,ARIN,LACNIC,RIPE
name	ASN allocation	2010.04.22	IANA
	autnum.txt	2010.08.19	potaroo.net
	whois	2010.04.01	AFRINIC,APNIC,ARIN,LACNIC,RIPE
topology	BGP	2010.01.29	Ripe NCC RCC12,routeviews2

AS Rank

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

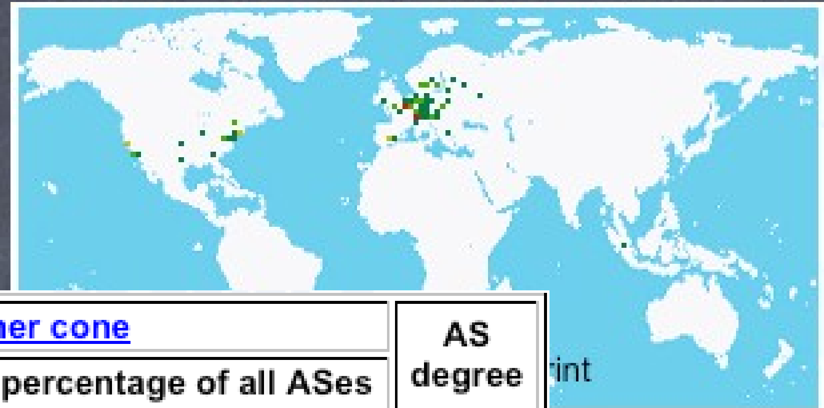

AS number:	1299
AS name:	TeliaNet Global Network
rank:	10
customer cone size:	27117
degree:	561

rank	AS number	AS name	customer cone		AS degree
			customer cone size	percentage of all ASes	
5	3549	Global Crossing Ltd.	29035	86%	1365
6	1239	Sprint	29012	86%	1381
7	209	Qwest Communications	28983	86%	1387
8	6939	Hurricane Electric,	27227	81%	1552
9	4323	tw telecom holdings,	27198	81%	1291
10	1299	TeliaNet Global Netw	27117	80%	561
11	2914	NTT America, Inc.	26832	79%	650
12	6453	TATA Communications	26236	78%	530
13	3561	Savvis	25690	76%	425
14	9002	ReTN.net Autonomous	25146	74%	

Ranking

Customers, providers, and peers

rank	neighbor AS	neighbor name	type
4	174	Cogent/PSI	↑ provider
6	1239	Sprint	↔ peer
5	3549	Global Crossing Ltd.	↑ provider
7	209	Qwest Communications	↑ provider
8	6939	Hurricane Electric,	↔ peer
9	4323	tw telecom holdings,	↔ peer
11	2914	NTT America, Inc.	↓ customer
12	6453	TATA Communications	↓ customer
13	3561	Savvis	↓ customer
15	1273	Cable and Wireless p	↓ customer

AS Rank visualization

- Graphical view of customers, providers and peers.

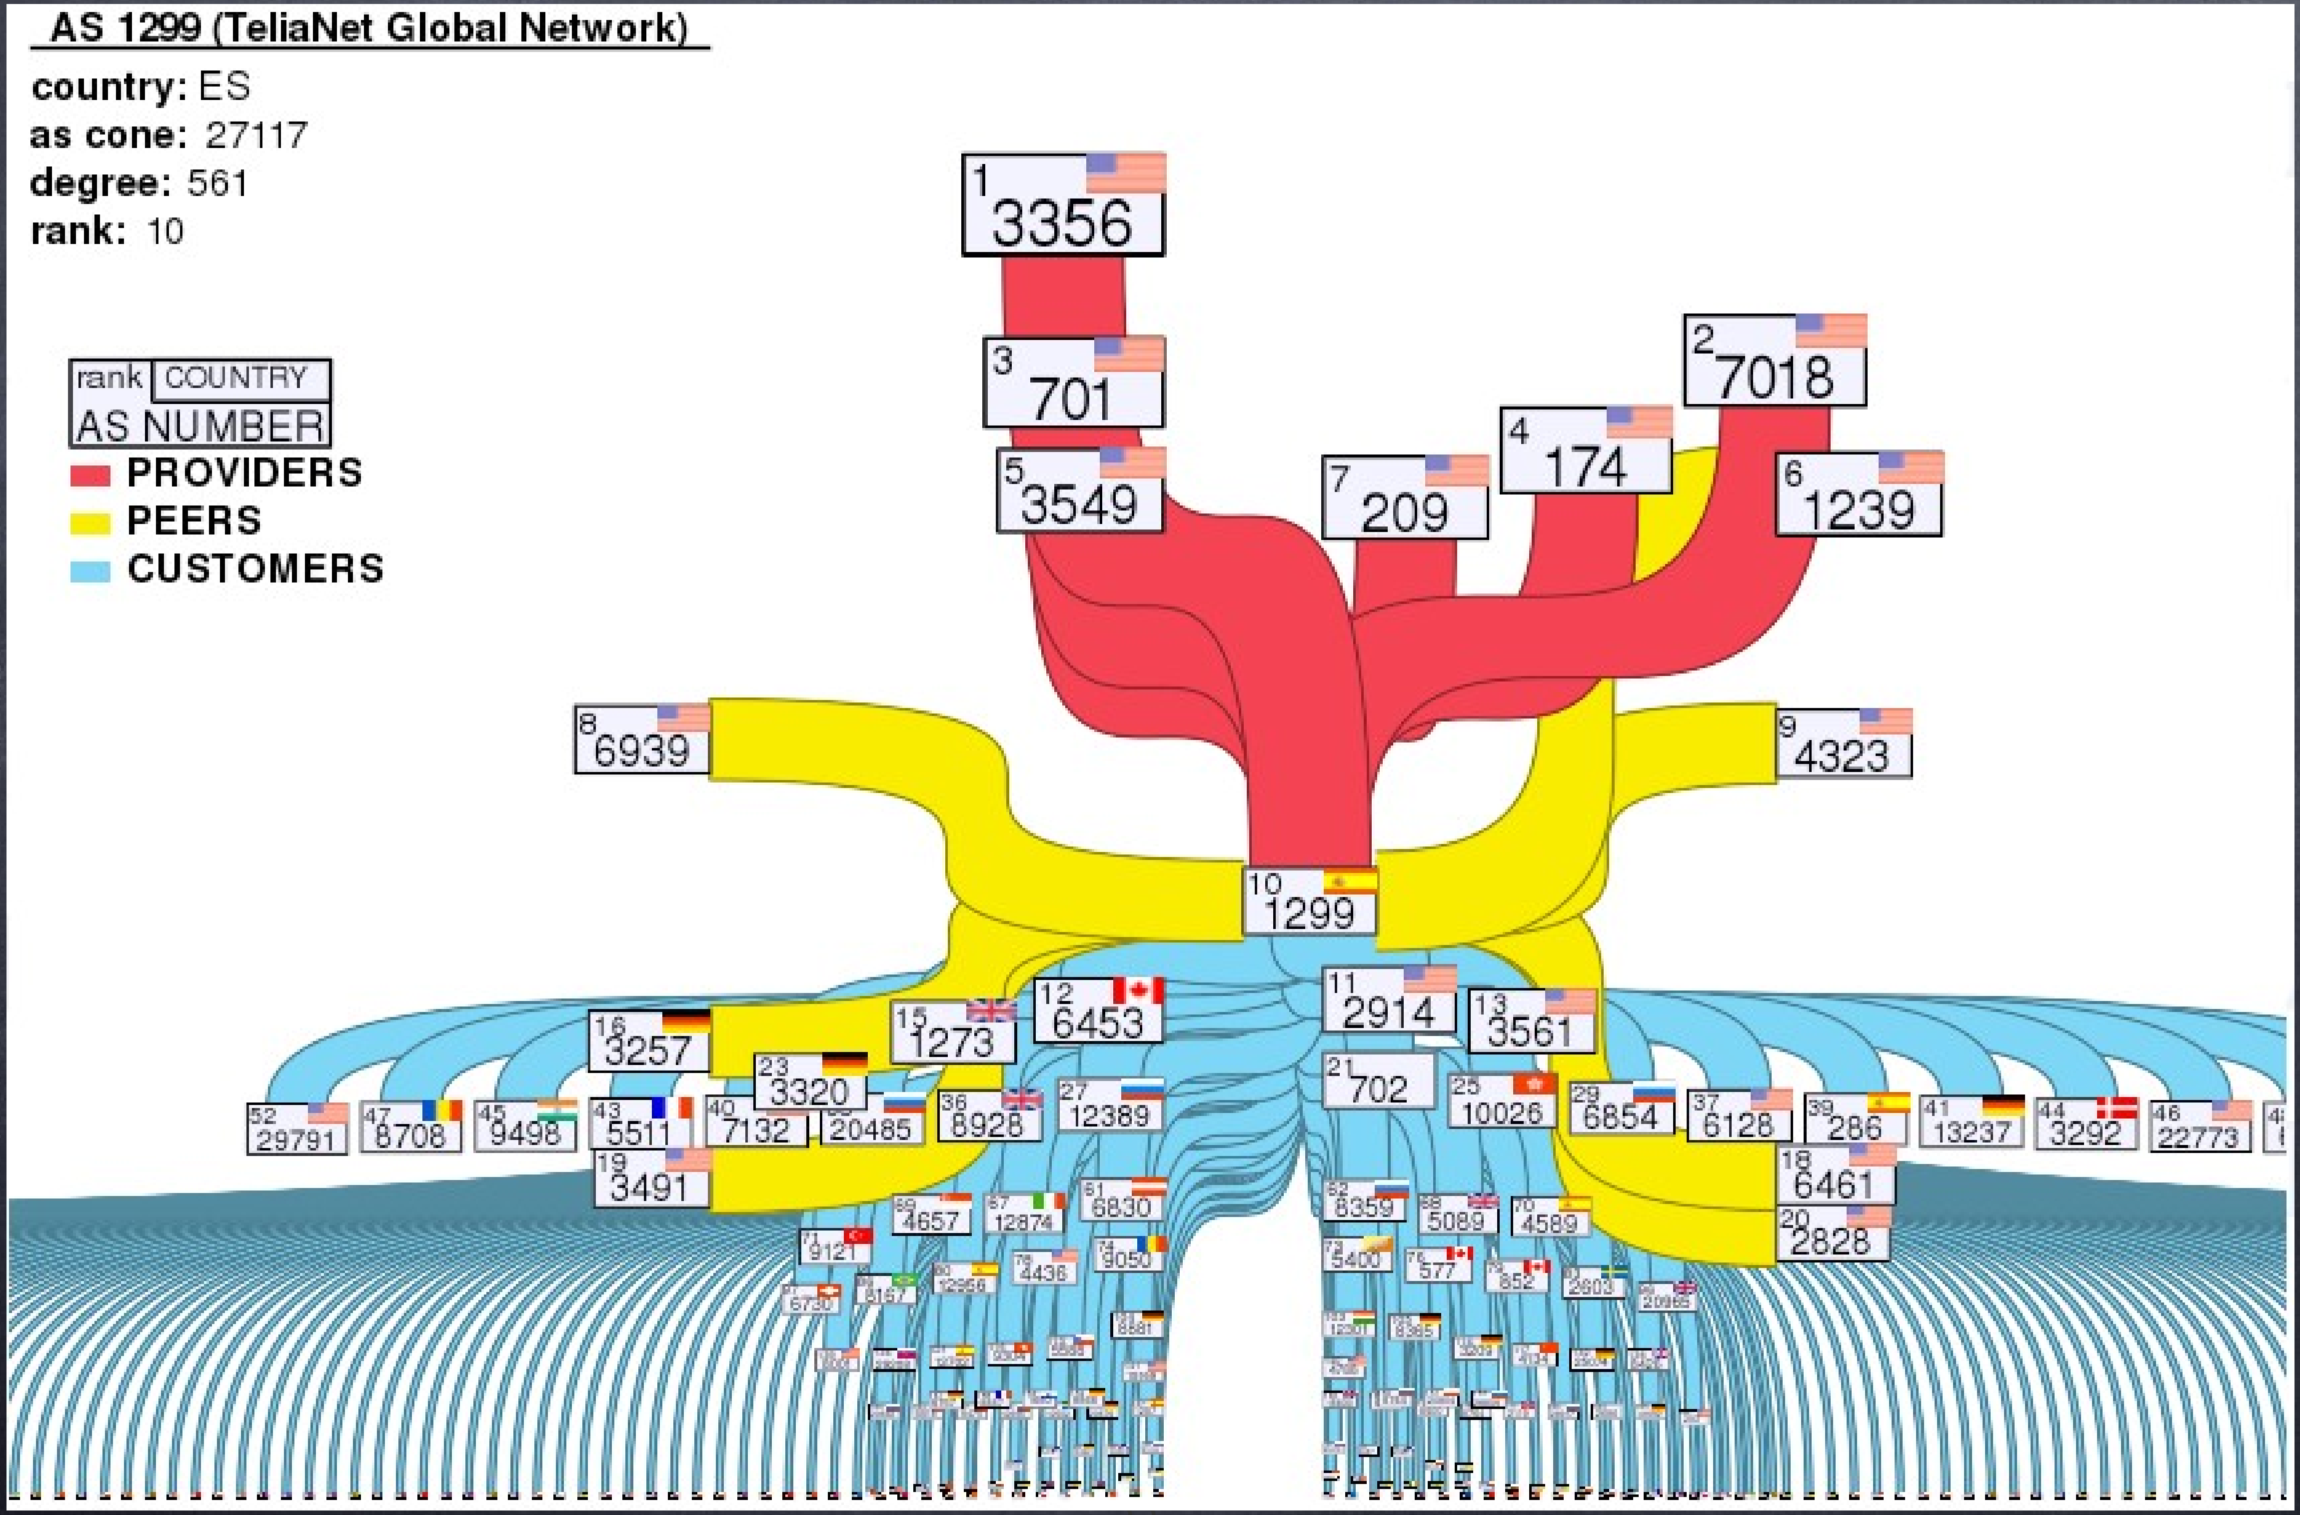

AS Rank validation

- Interface to provide corrections to relationships.

rank	neighbor AS	neighbor name	type	correction
4	174	Cogent/PSI	↑ provider	
6	1239	Sprint	↔ peer	provider
5	3549	Global Crossing Ltd.	↑ provider	
7	209	Qwest Communications	↑ provider	
8	6939	Hurricane Electric,	↔ peer	
9	4323	tw telecom holdings,	↔ peer	
11	2914	NTT America, Inc.	↓ customer	peer
12	6453	TATA Communications	↓ customer	
13	3561	Savvis	↓ customer	
15	1273	Cable and Wireless p	↓ customer	

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

Geolocation Tools Comparison

- Geolocation Service Evaluation Criteria
  - What geographic granularity does it provide?
  - Continent, country, state/prefecture, city, zip code.
  - What Internet identifier granularity does it support?
  - Internet Protocol (IP) address, network prefix, Autonomous System (AS).
  - Does the accuracy of the results vary by geographic region or by type of network?
  - With what frequency does a service update its database?
  - How many queries per second can clients execute?

Geolocation Tools Comparison

- Geolocation tools we hope to evaluate
  - Digital Envoy's Netacuity
  - MaxMind
    - Free and commercial versions
  - Akamai
  - Google
  - IP2Location
  - Quova
  - IPligence
  - HostIP.info

Schedule, Planned activities

- 1-2 monitors/month
- IPv4, IPv6 topology data
- Continue to release and refine ITDK
  http://www.caida.org/data/active/internet-topology-data-kit/
- Will publish alias resolution study
- Visualization (in support of)
- Validation against ground truth
- AIMS 2011
- Begin work on BGP data coupling to Ark
- AS Rank
- Geolocation Tools Comparison

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

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

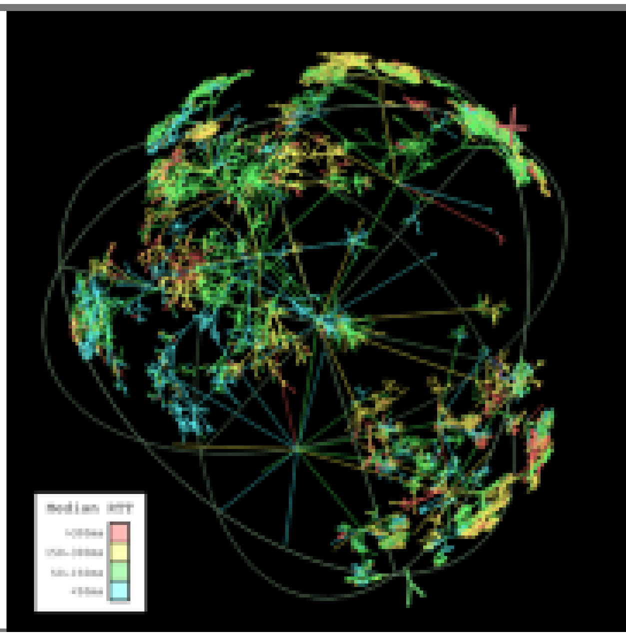

Walrus visualizations of round-trip time measurements made by CAIDA's macroscopic Internet topology monitor located in Herndon, VA, USA.

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