Leveraging the Science and Technology of Internet Mapping for Homeland Security

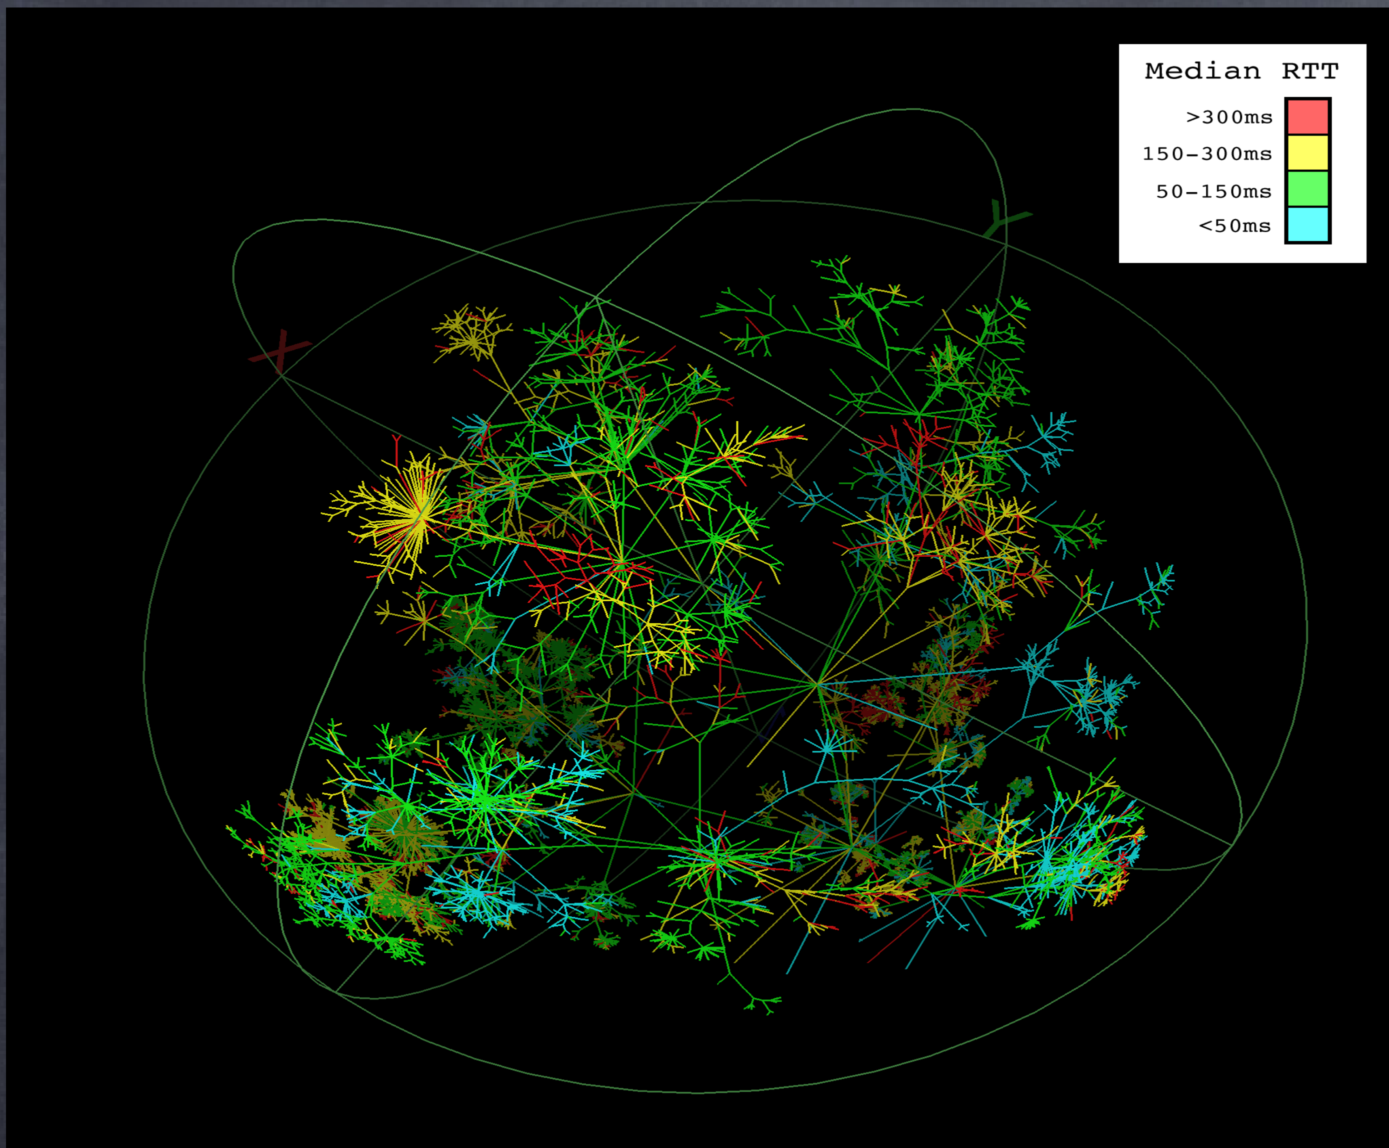

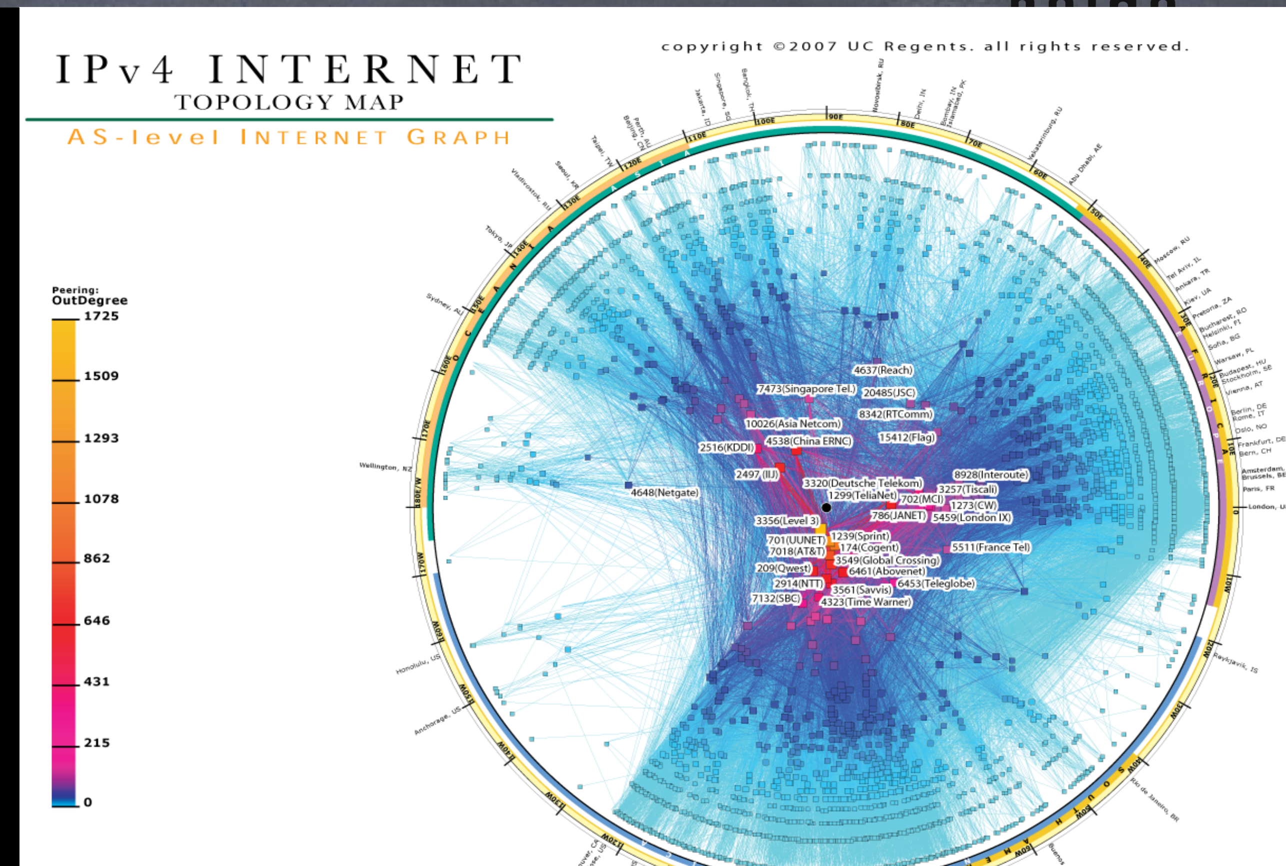

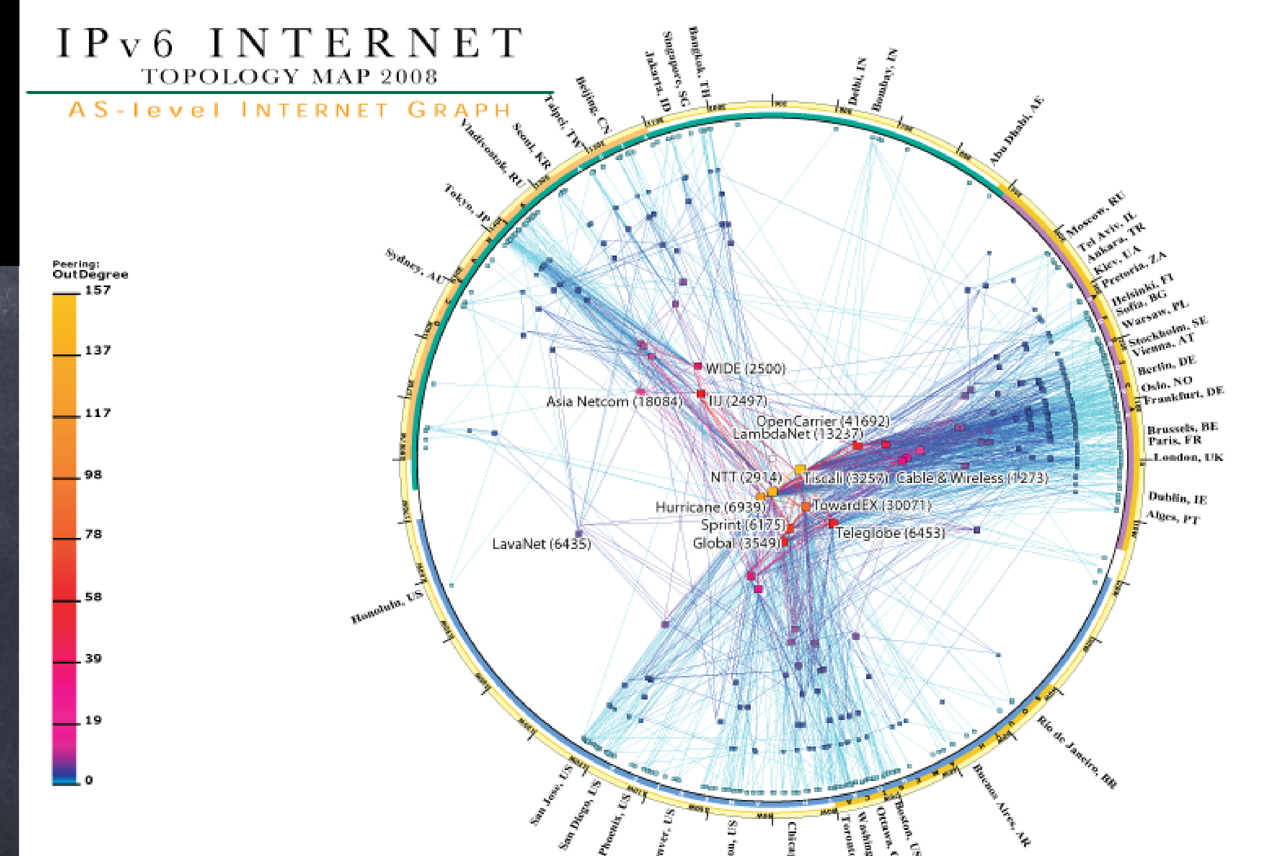

Young Hyun, Ken Keys, Amogh Dhamdhere, Bradley Huffaker, Joshua Polterock, Marina Fomekov, Dima Krioukov, kc claffy

CAIDA

DHS – PI meeting
SRI Roslyn, VA
9 March 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
- 43 active IPv4 probers
    - 15 in US
- 11 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 (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

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

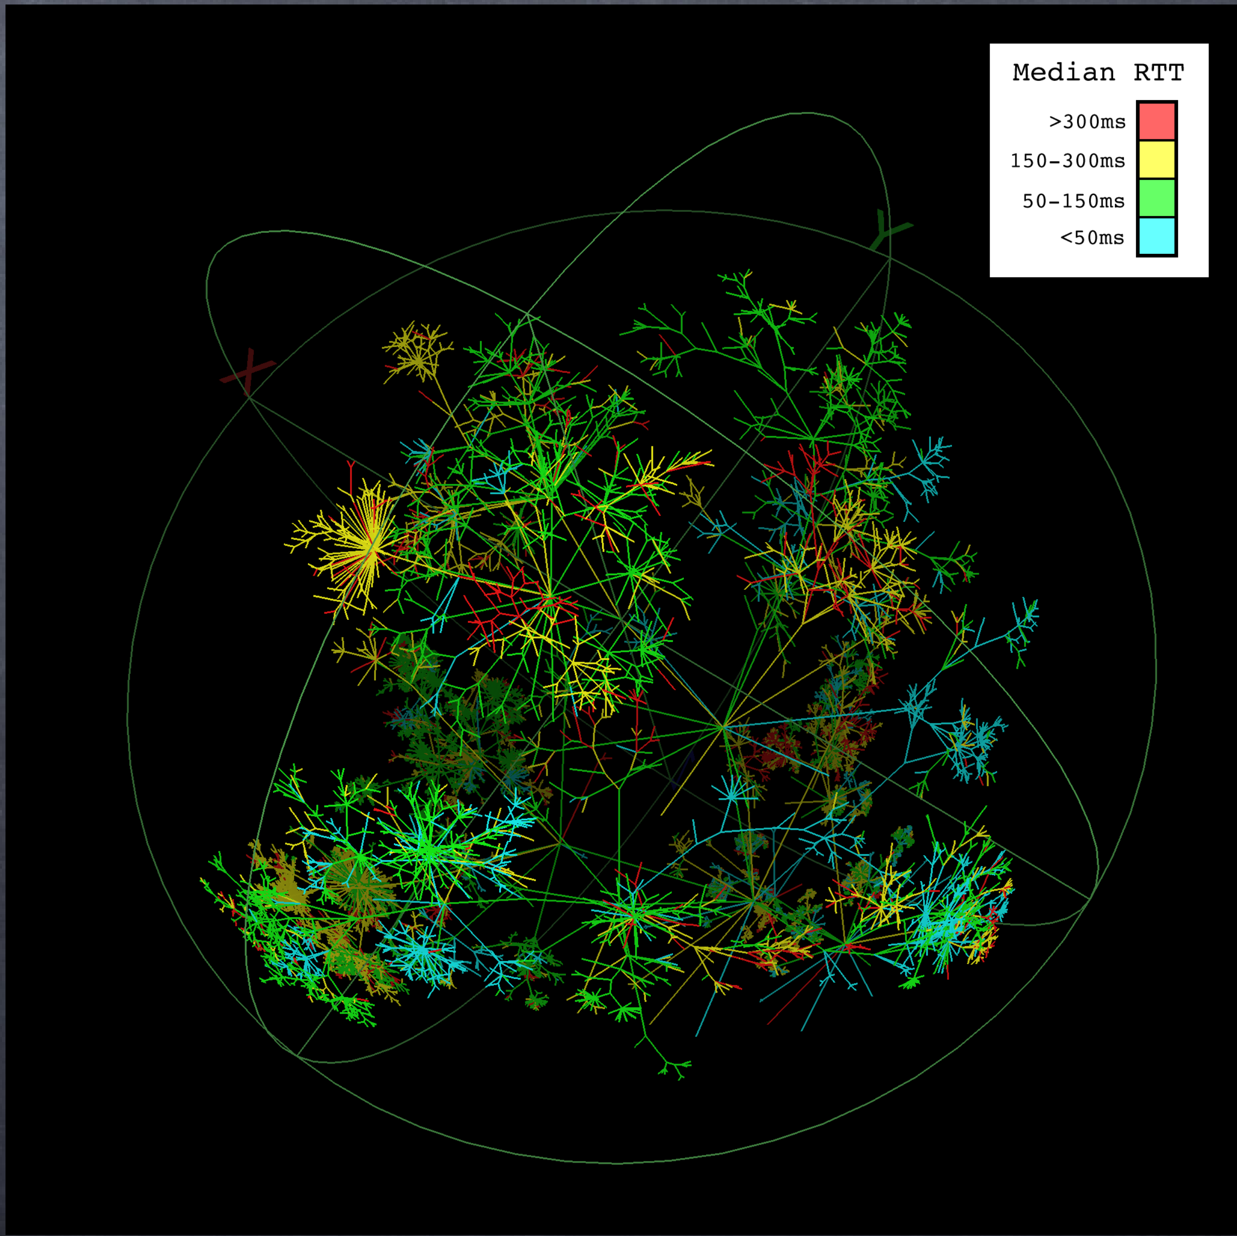

Methodology insights enabled

- Probing technique performance comparision
- Macroscopic vulnerability assessment: filtering
- Understanding Internet topology: theory and method

http://www.caida.org/publications/papers/2010/alias_resolution/

  - compare performance and accuracy of known alias resolution techniques used at Internet scale
  - develop enhancements
    - kapar (improved 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 remains validation (not tech problem)

2009-10 technical (infra.) accomplishments

- 43 monitors now active, 11 probing IPv6
- IPv4 topology data
  - 2.4TB data served by PREDICT, data.caida.org
  - collected from Sep 2007 to Jan 2010 (29 months):
    - 5.7 billion traceroutes; 2.3TB data
    - ~800 cycles
  - collecting every month now:
    - ~290 million traceroutes; ~120 GB data
  - IPv4 topology data is key input into other datasets e.g., AS links and alias resolution
  - Currently each cycle of each team collects traces from 8.25 million /24s
- IPv6 topology data

Ark monitors/data over time

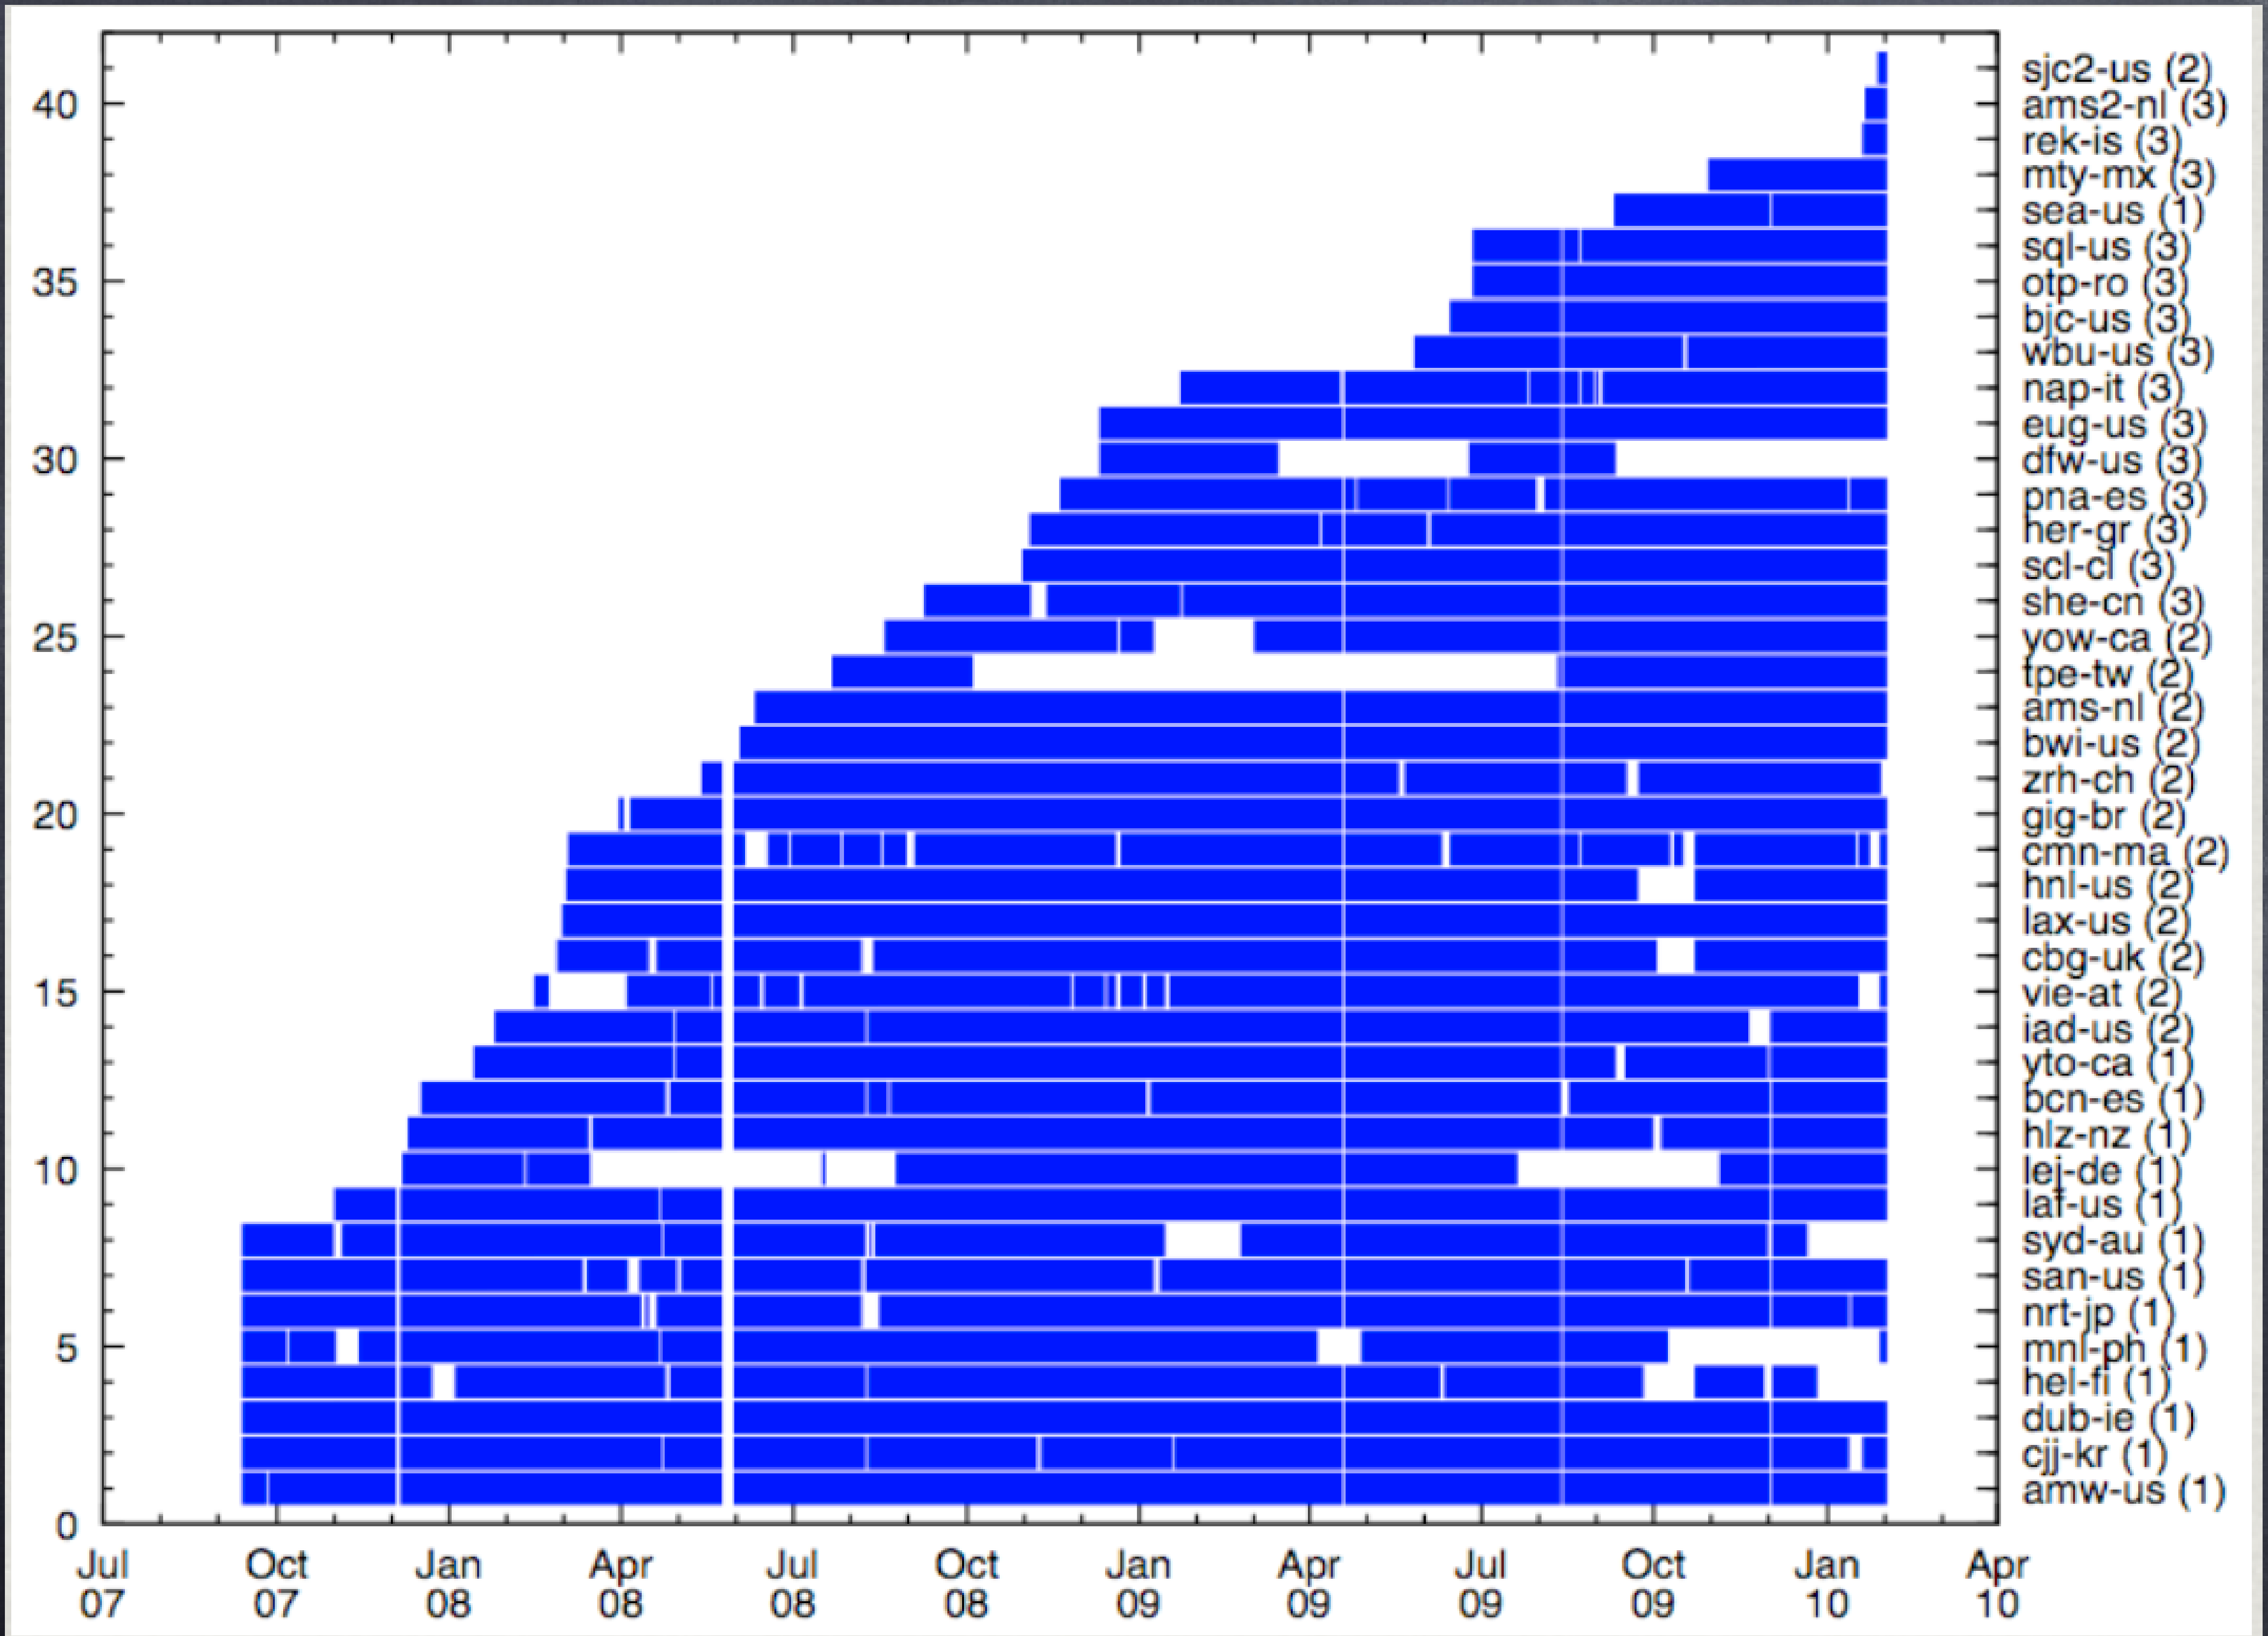

data availability per monitor (row)

2009-10 technical accomplishments

- AIMS workshop (Feb 2010) {w/PREDICT}
- Data for IP-to-router resolution (Dec.)
- Ark-based AS-level and router-level graph (Feb.)

  http://www.caida.org/data/active/internet-topology-data-kit/
- Ark-based dual AS-router graph (June)
  - Preliminary dual graph in B. Huffaker, A. Dhamdhere, M.Fomenkov, kc claffy, “Towards Topology Dualism: Improving the Accuracy of AS Annotations for Routers”, to be published in the proceedings of the Passive and Active Measurement Conference (PAM) in 2010. http://www.caida.org/publications/papers/2010/as_assignment/
- Tool for calculating topology statistics – topostats (Feb.)

  http://www.caida.org/tools/utilities/topostats/
- Supporting software: mper, marinda, midar, kapar,

Internet Topology Data Kit (ITDK) process

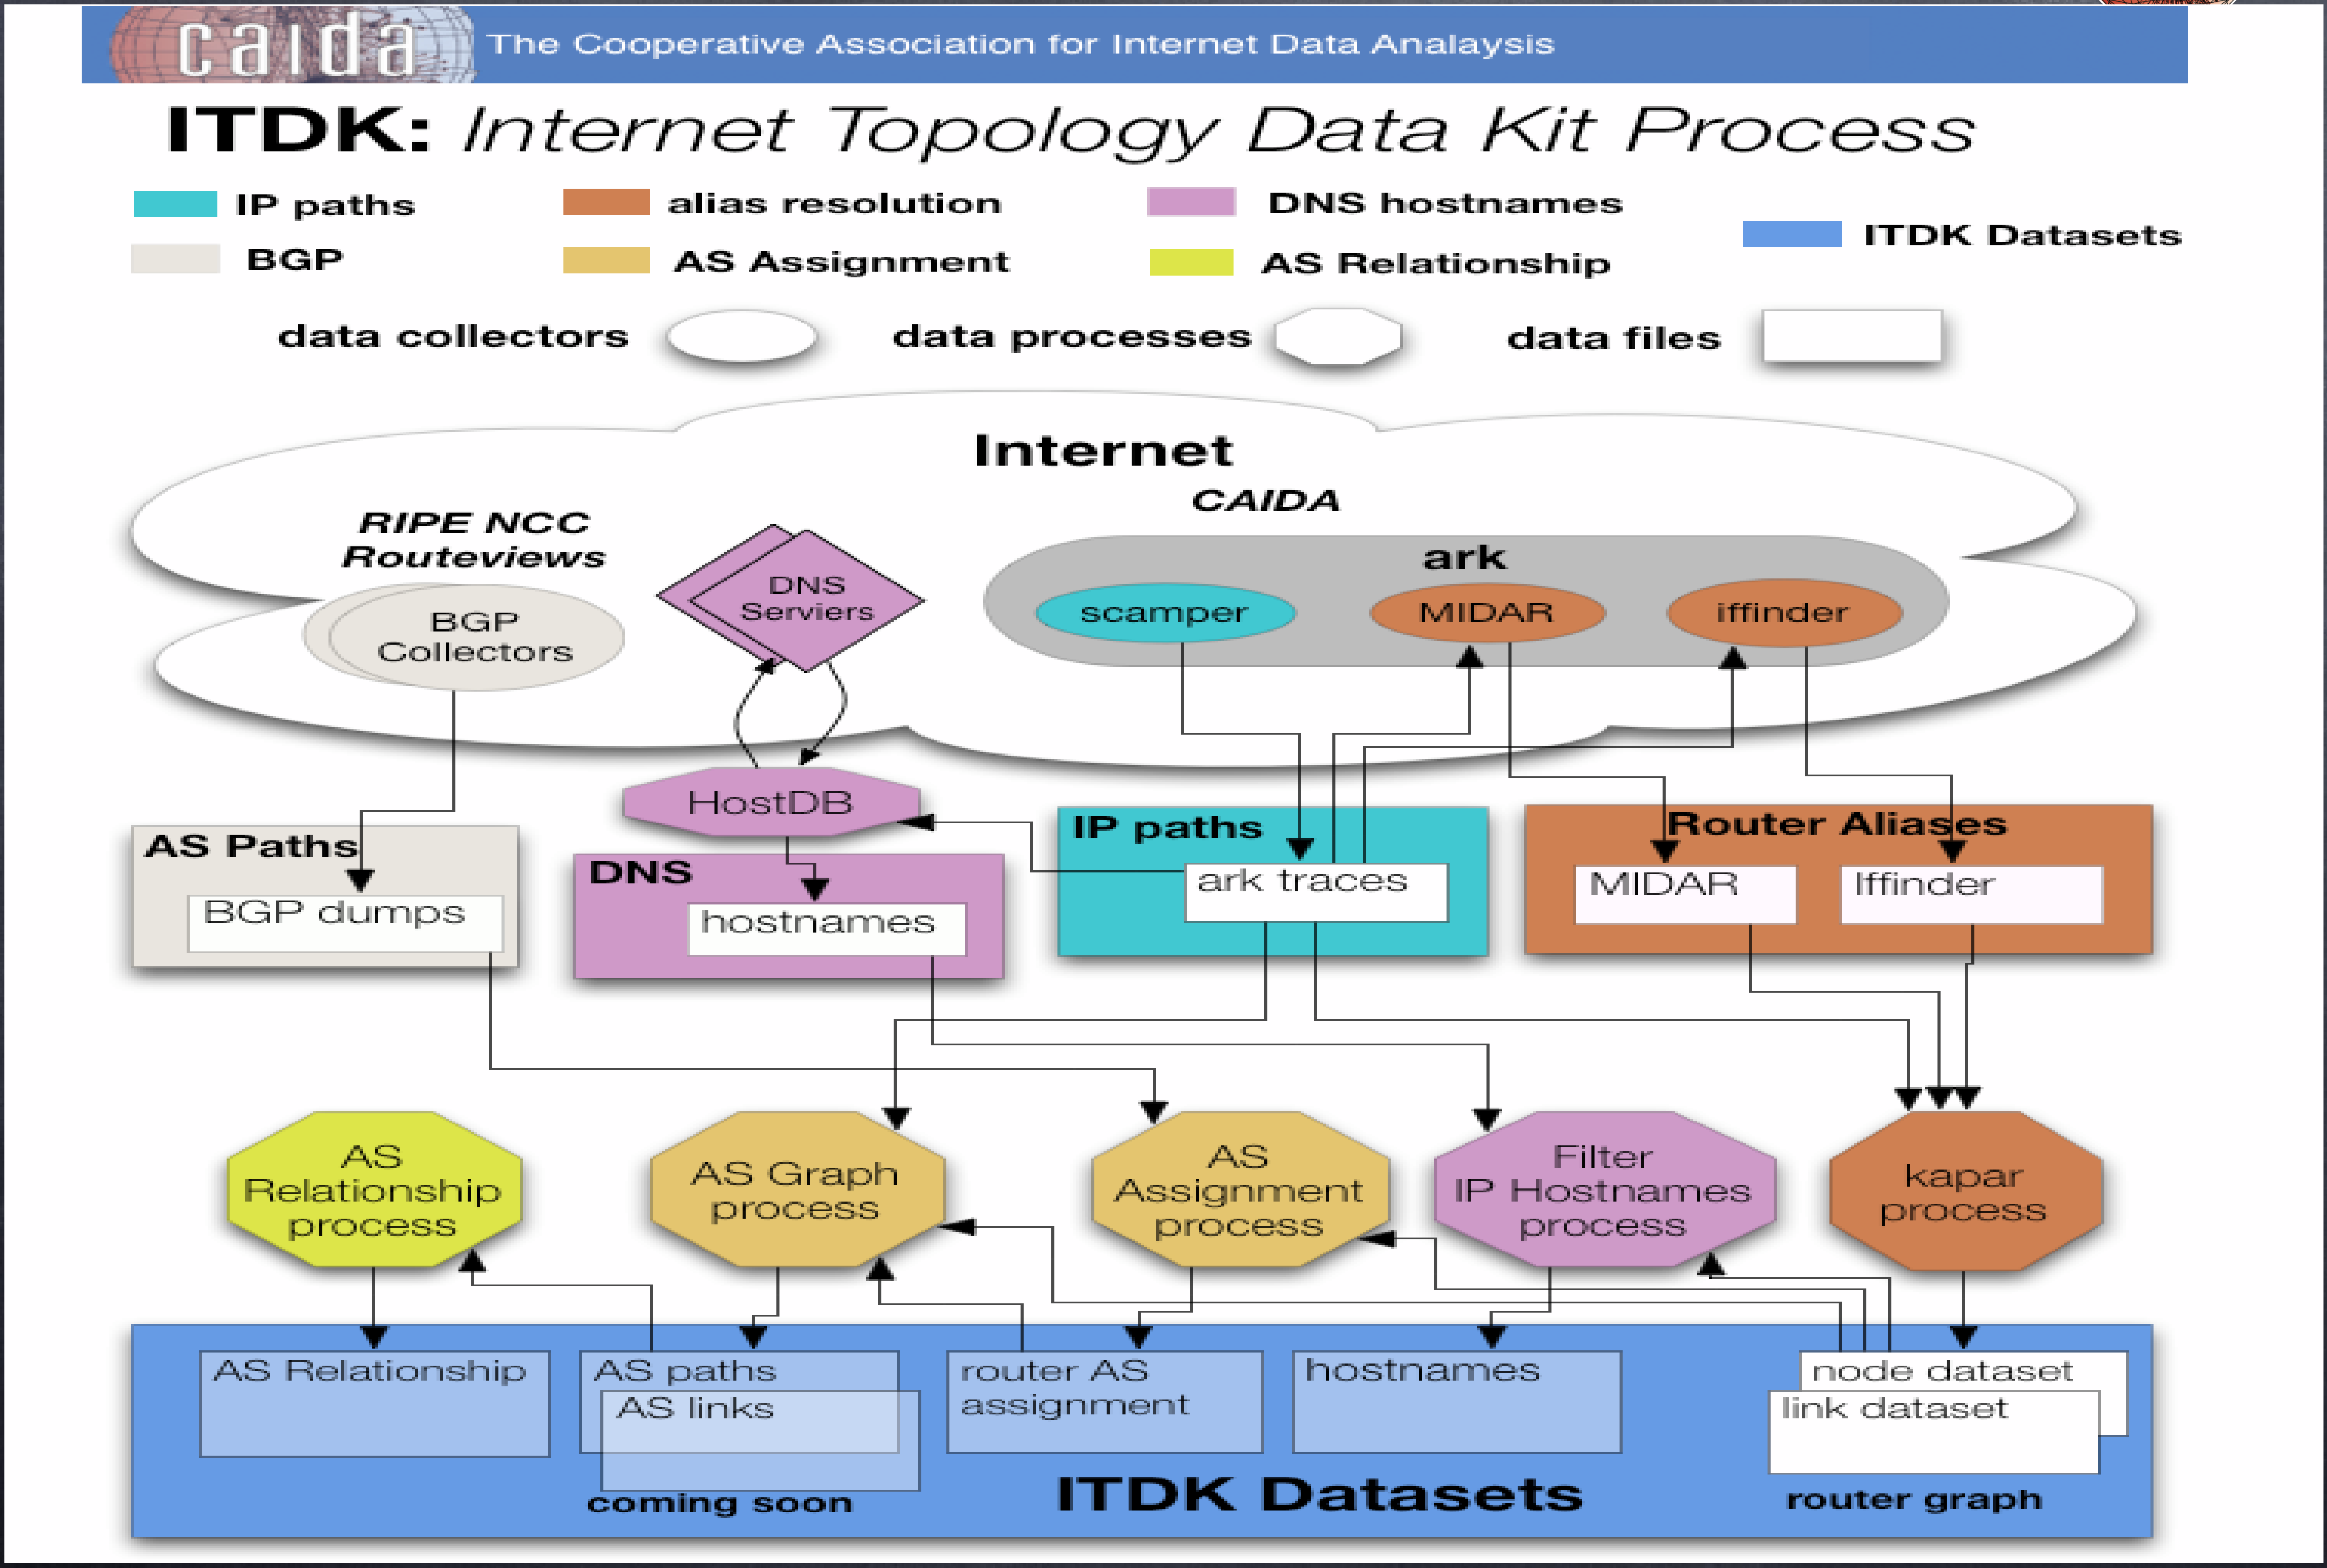

Measurement Big Picture

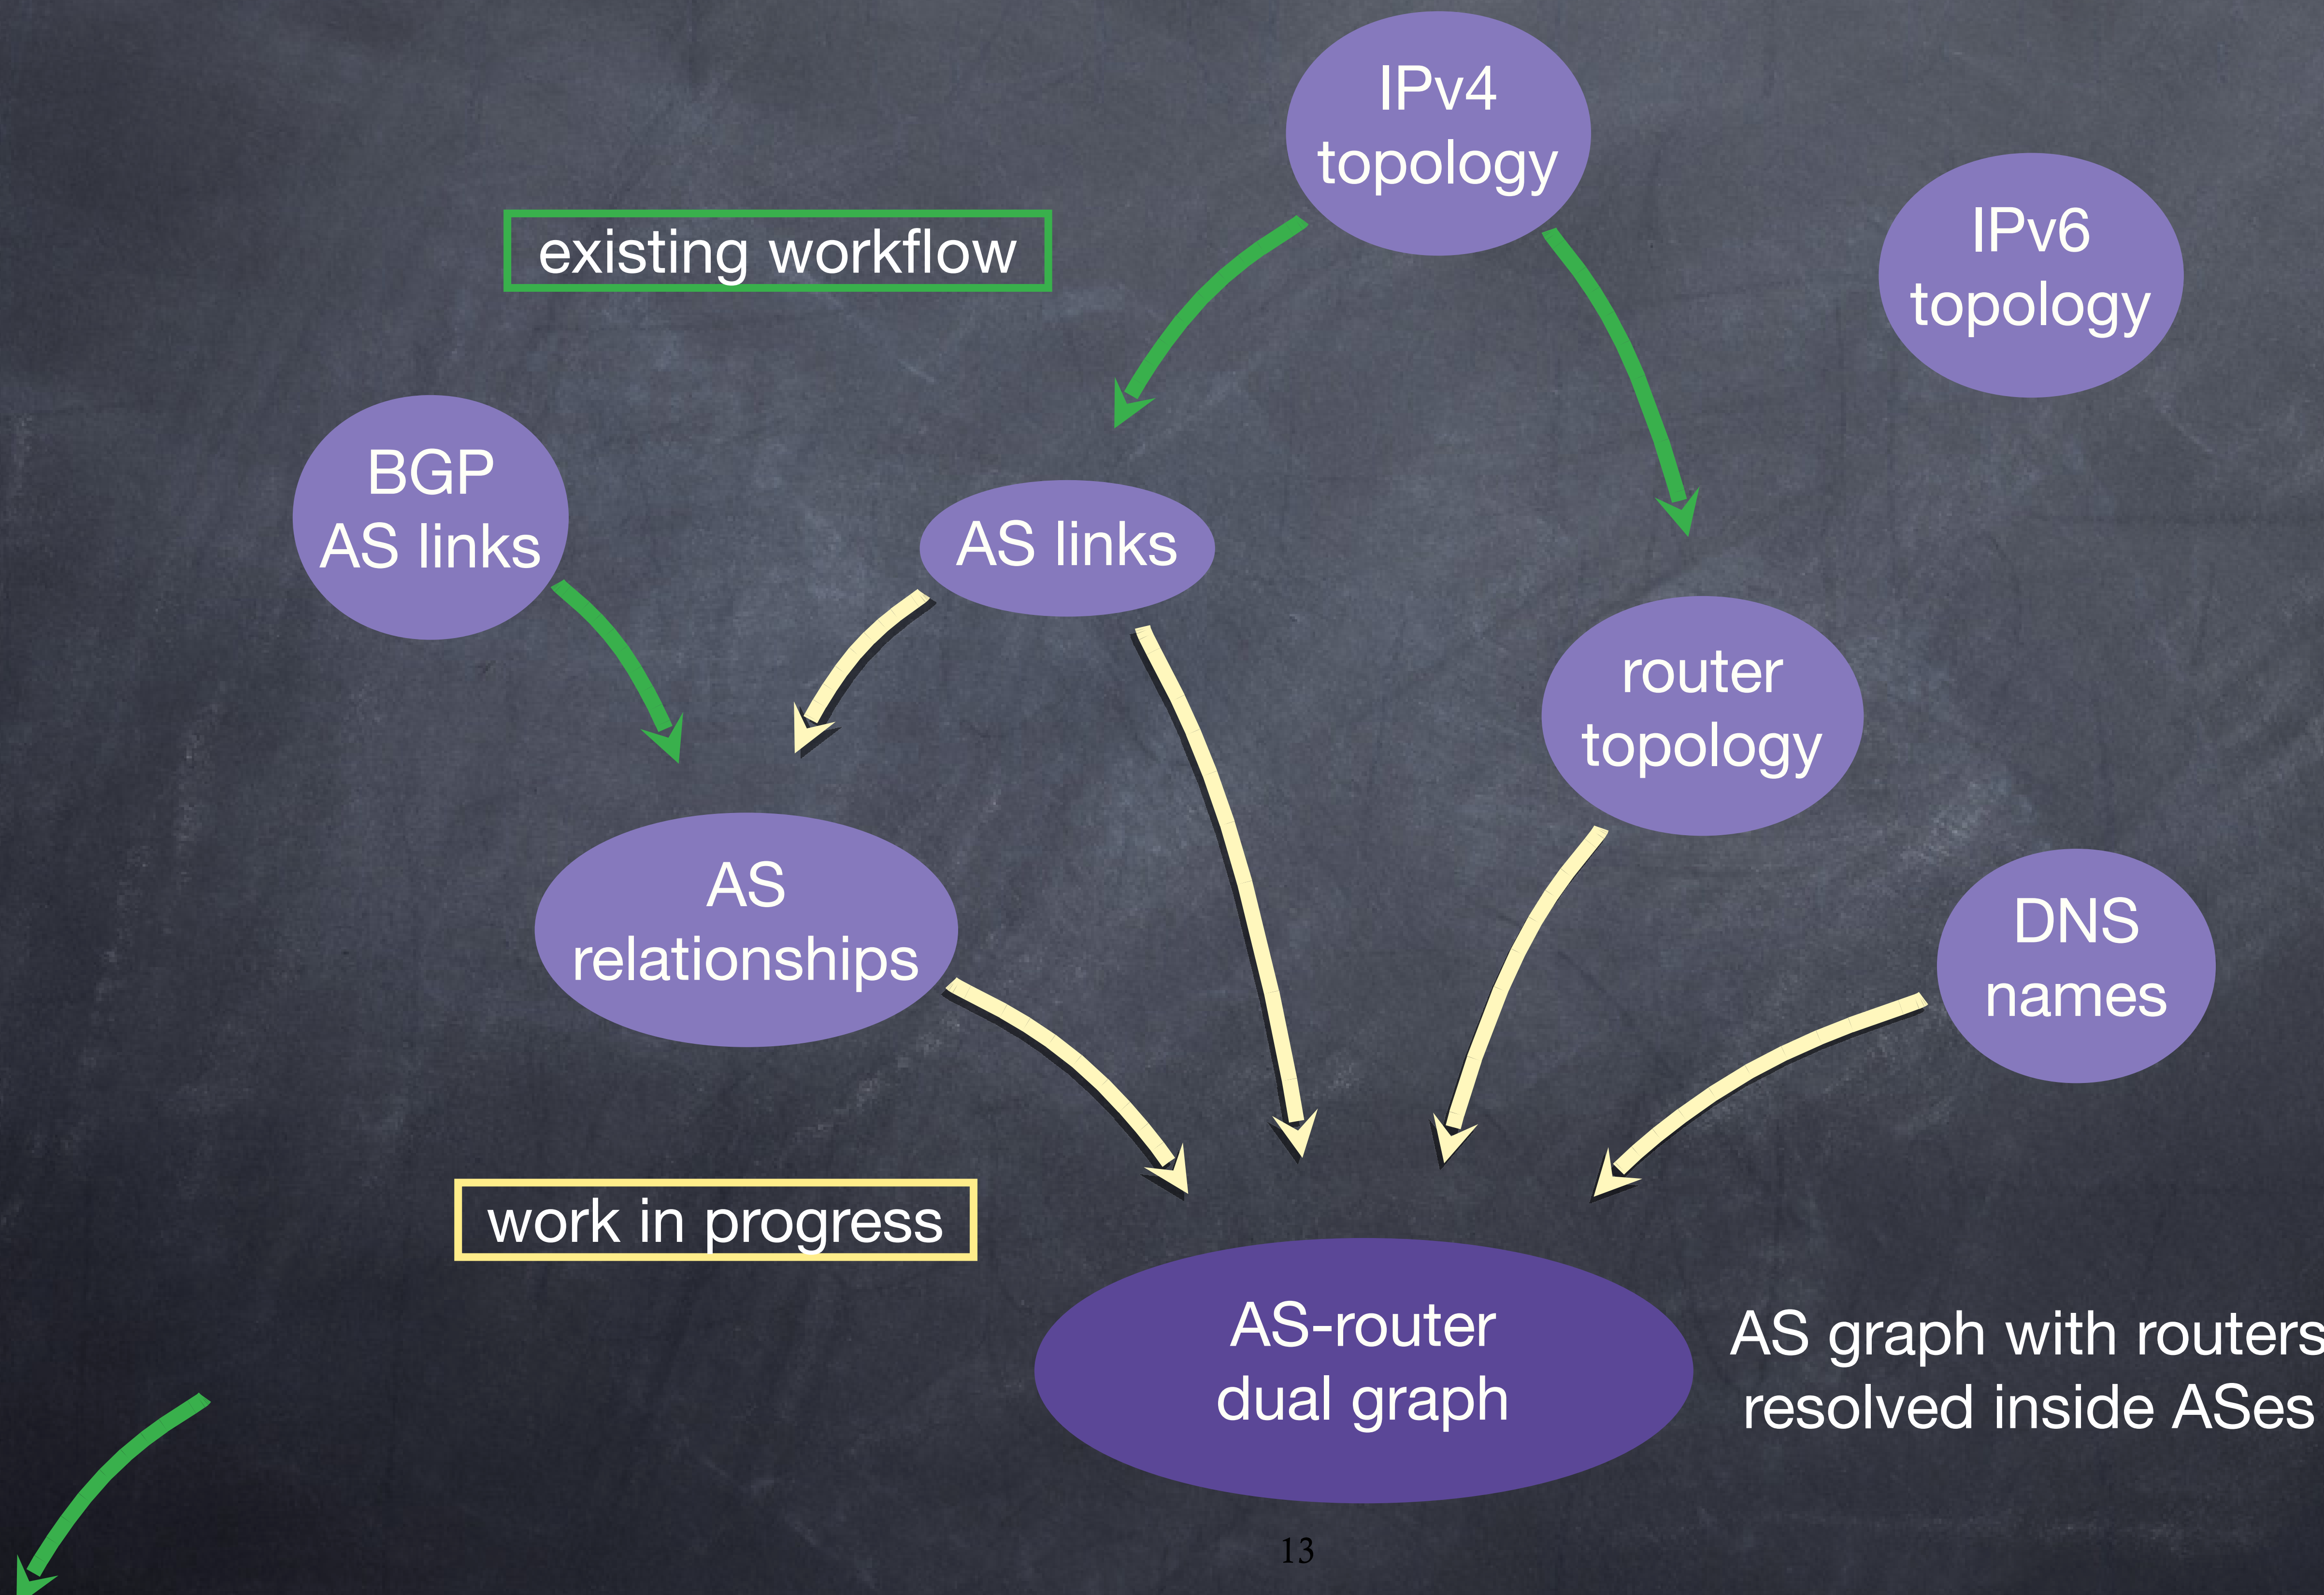

Measurements

- IPv4 Routed /24 Topology (and AS Links)
- IPv6 Topology
- DNS Names & Query/Response Traffic
- Alias Resolution

Data: IPv4 Routed /24 Topology

- ongoing large-scale topology measurements
  - ICMP Paris traceroute to every routed /24 (8.25 million)
    - about 126 /8-equivalents of routed space (as of Oct 2009)
  - running *scamper*
    - written by Matthew Luckie of WAND, University of Waikato
- dynamically divide up the measurement work among members of monitor teams
  - 3 teams active
  - 13-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 interfaces observed in traceroute paths into routers
  - toward a router-level map of the Internet
- earlier at CAIDA: iffinder, kapar (improved APAR)
- past year: MIDAR (Radargun++)
  - Intuition: two interfaces belonging to the same router will respond to probes in a similar way
  - specifically, IP ID values in response packets 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 the same router probed closely in time will return similar IP ID values; over time, similar time-series → use slope.

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*

RadarGun: nugget

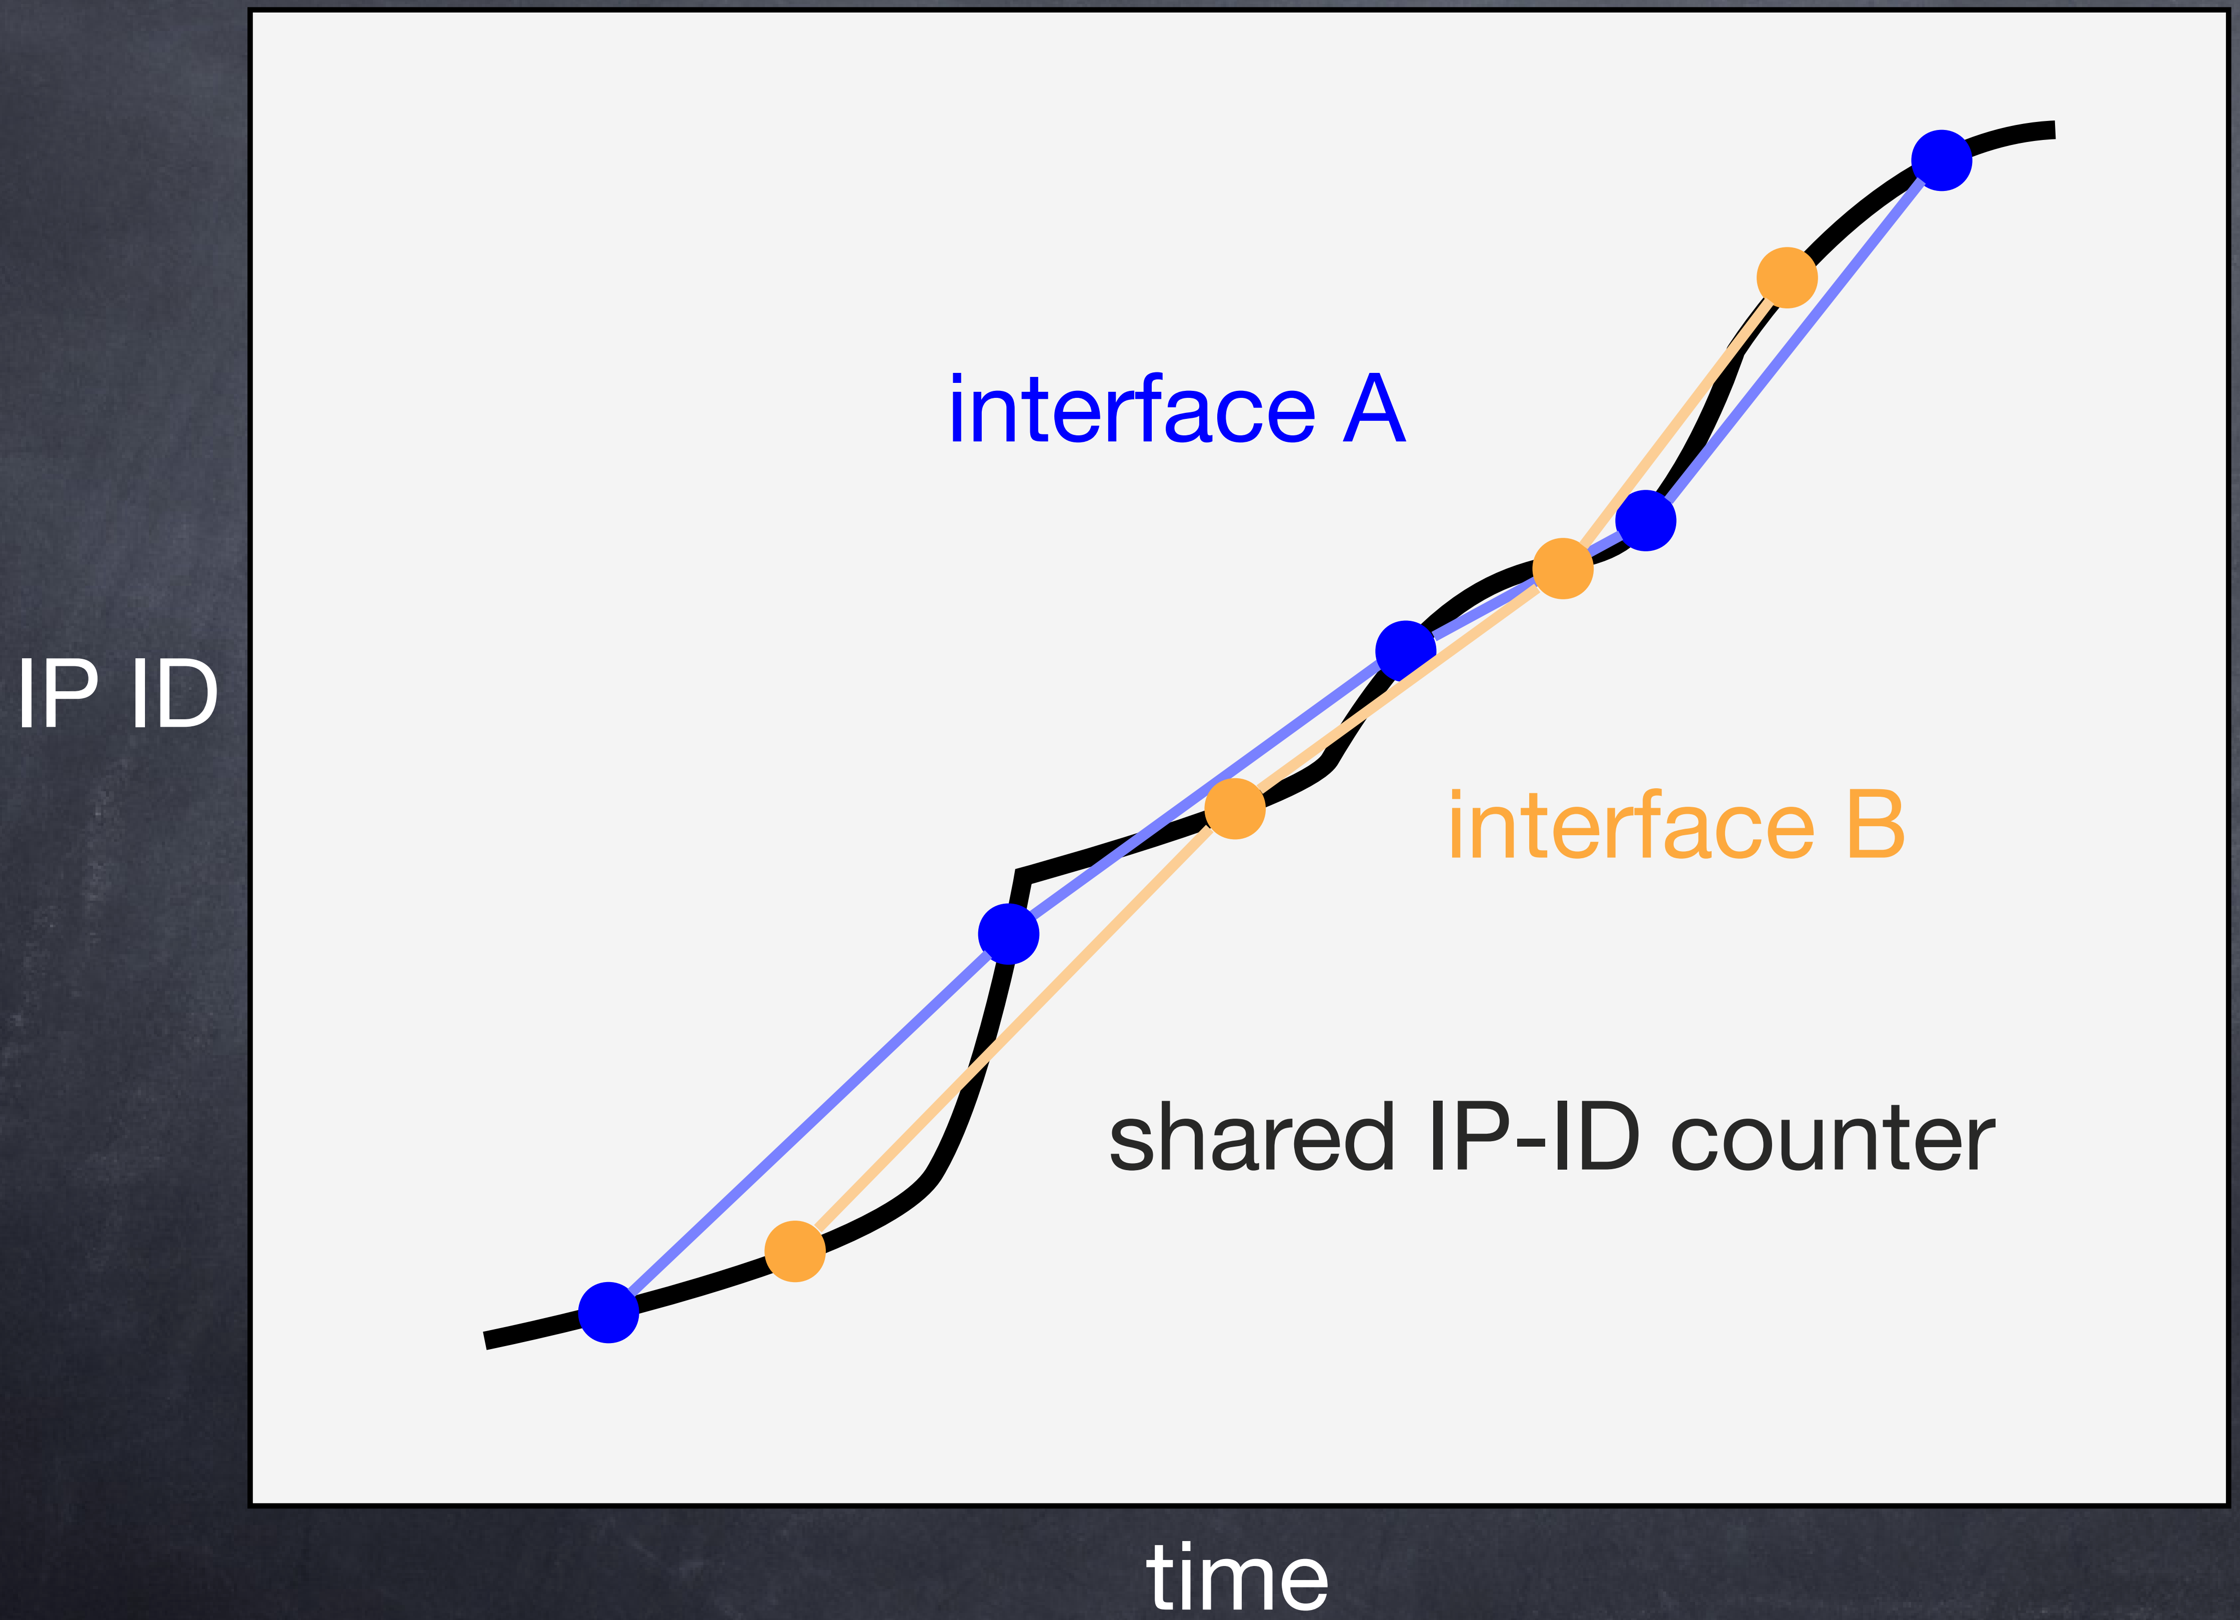

RadarGun's Distance Test

IP ID

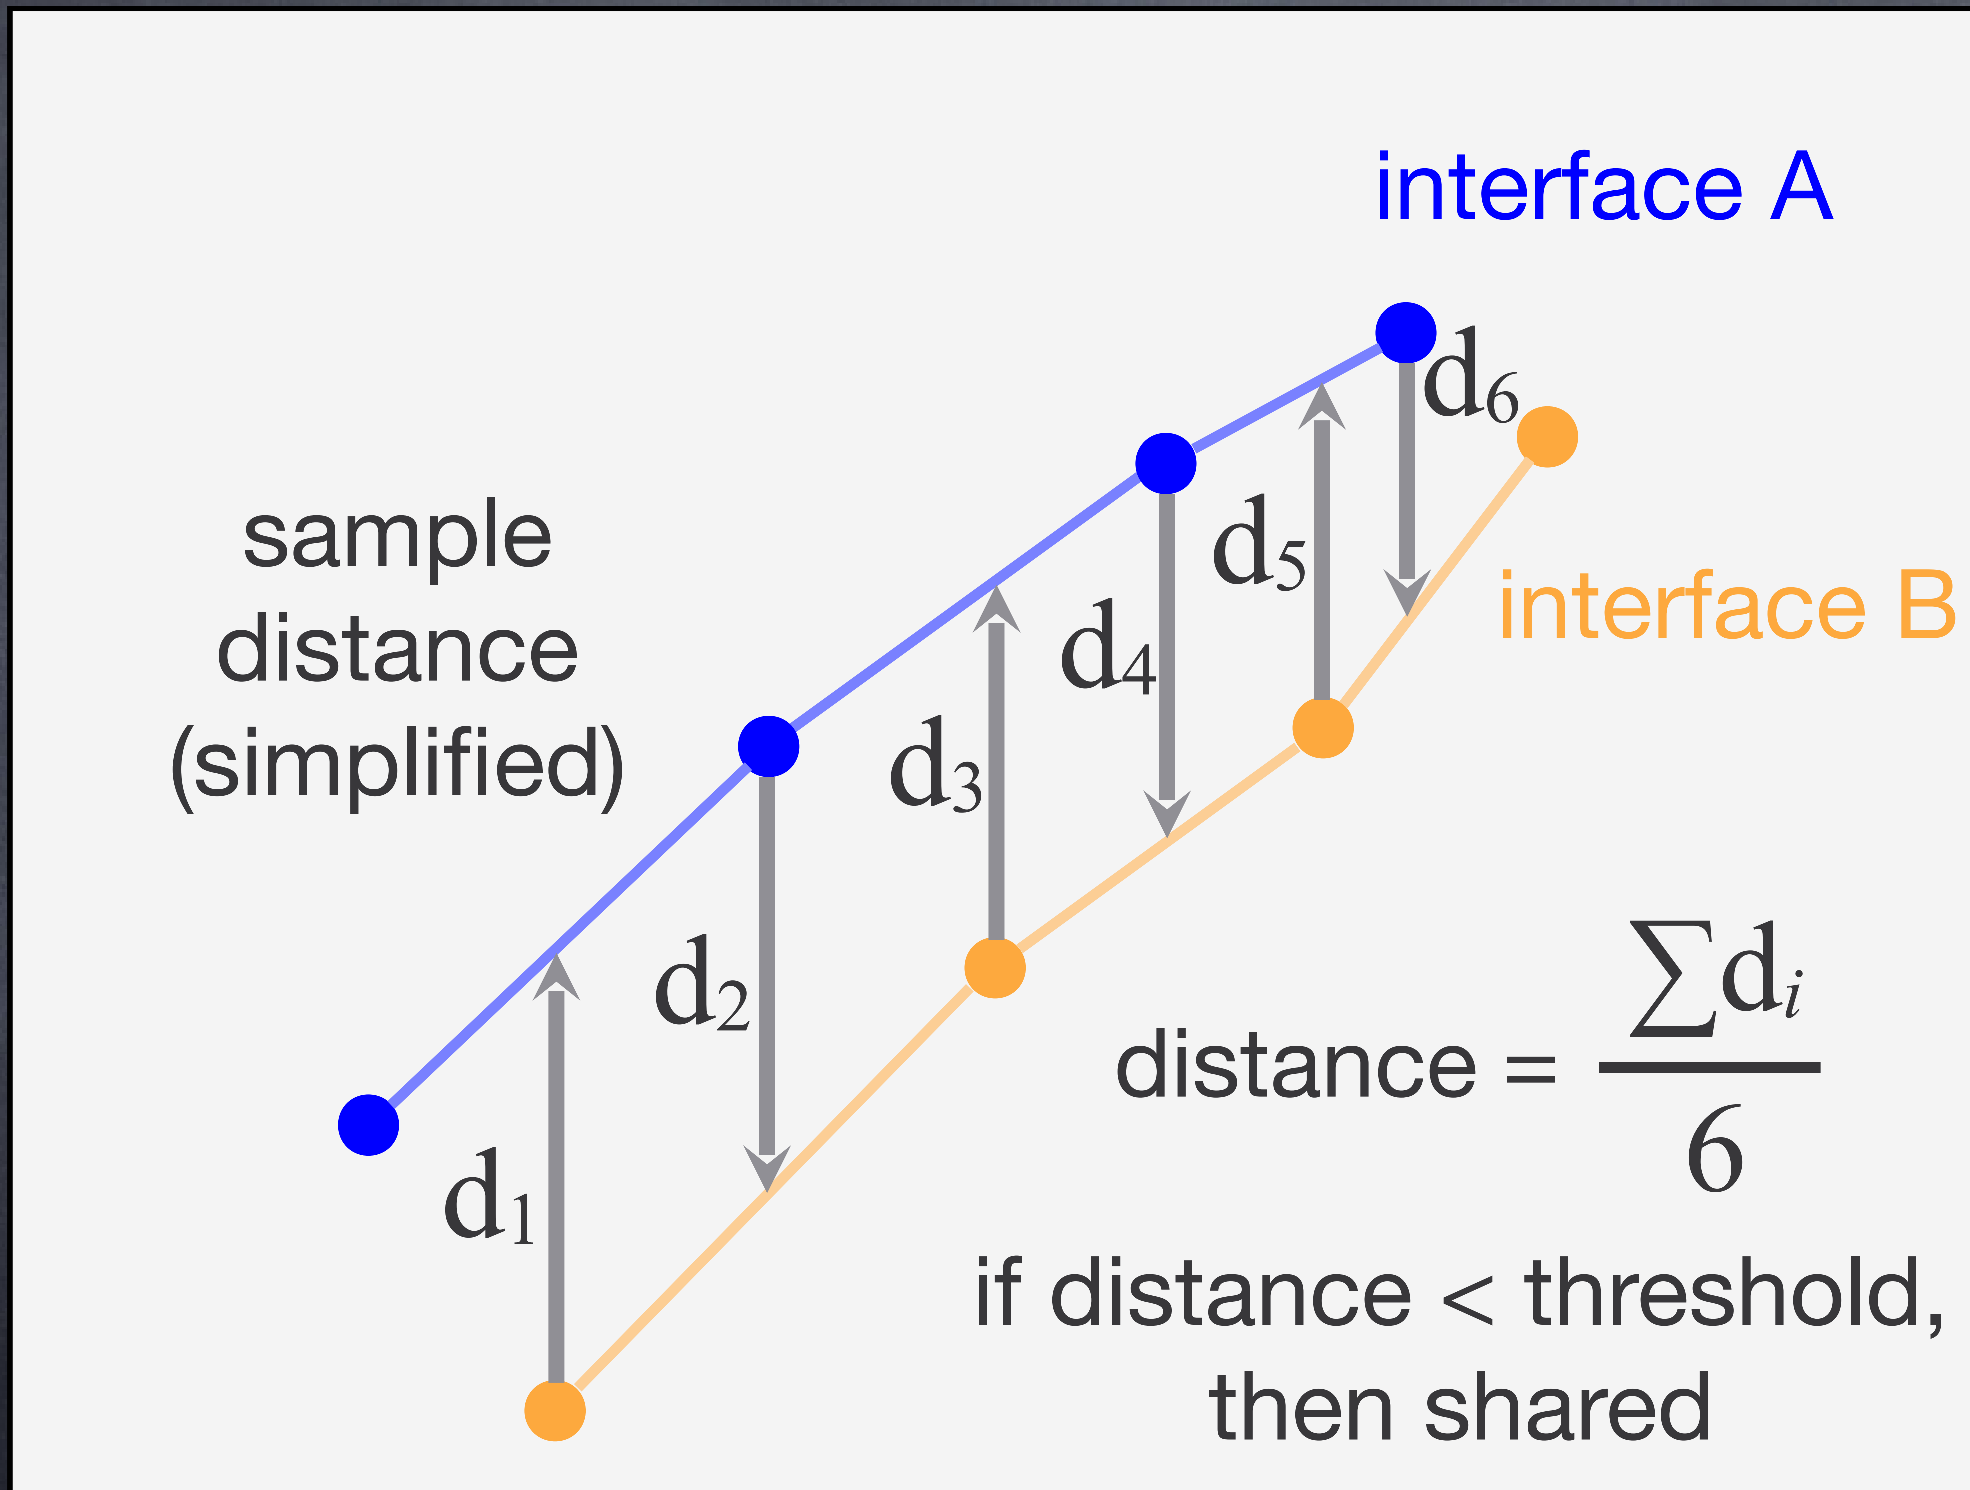

time

RadarGun Issues

* RadarGun is groundbreaking work but has both theoretical and practical issues
  * **the distance test for aliases is insufficient**
    * *threshold dependent on underlying dataset*
      * Bender, et al used traceroutes between PlanetLab nodes
      * Ark traceroutes are taken to the entire routed space
        * distance distribution noticeably different
    * *threshold doesn't account for velocity*
      * RadarGun *velocity* is the slope of the IP-ID time series
      * setting the threshold high enough to allow high-velocity aliases allows false positives in low-velocity cases
    * *false positives can exist for **any** chosen threshold*
      * even for a very low threshold

RadarGun false positive for any chosen threshold

IP ID

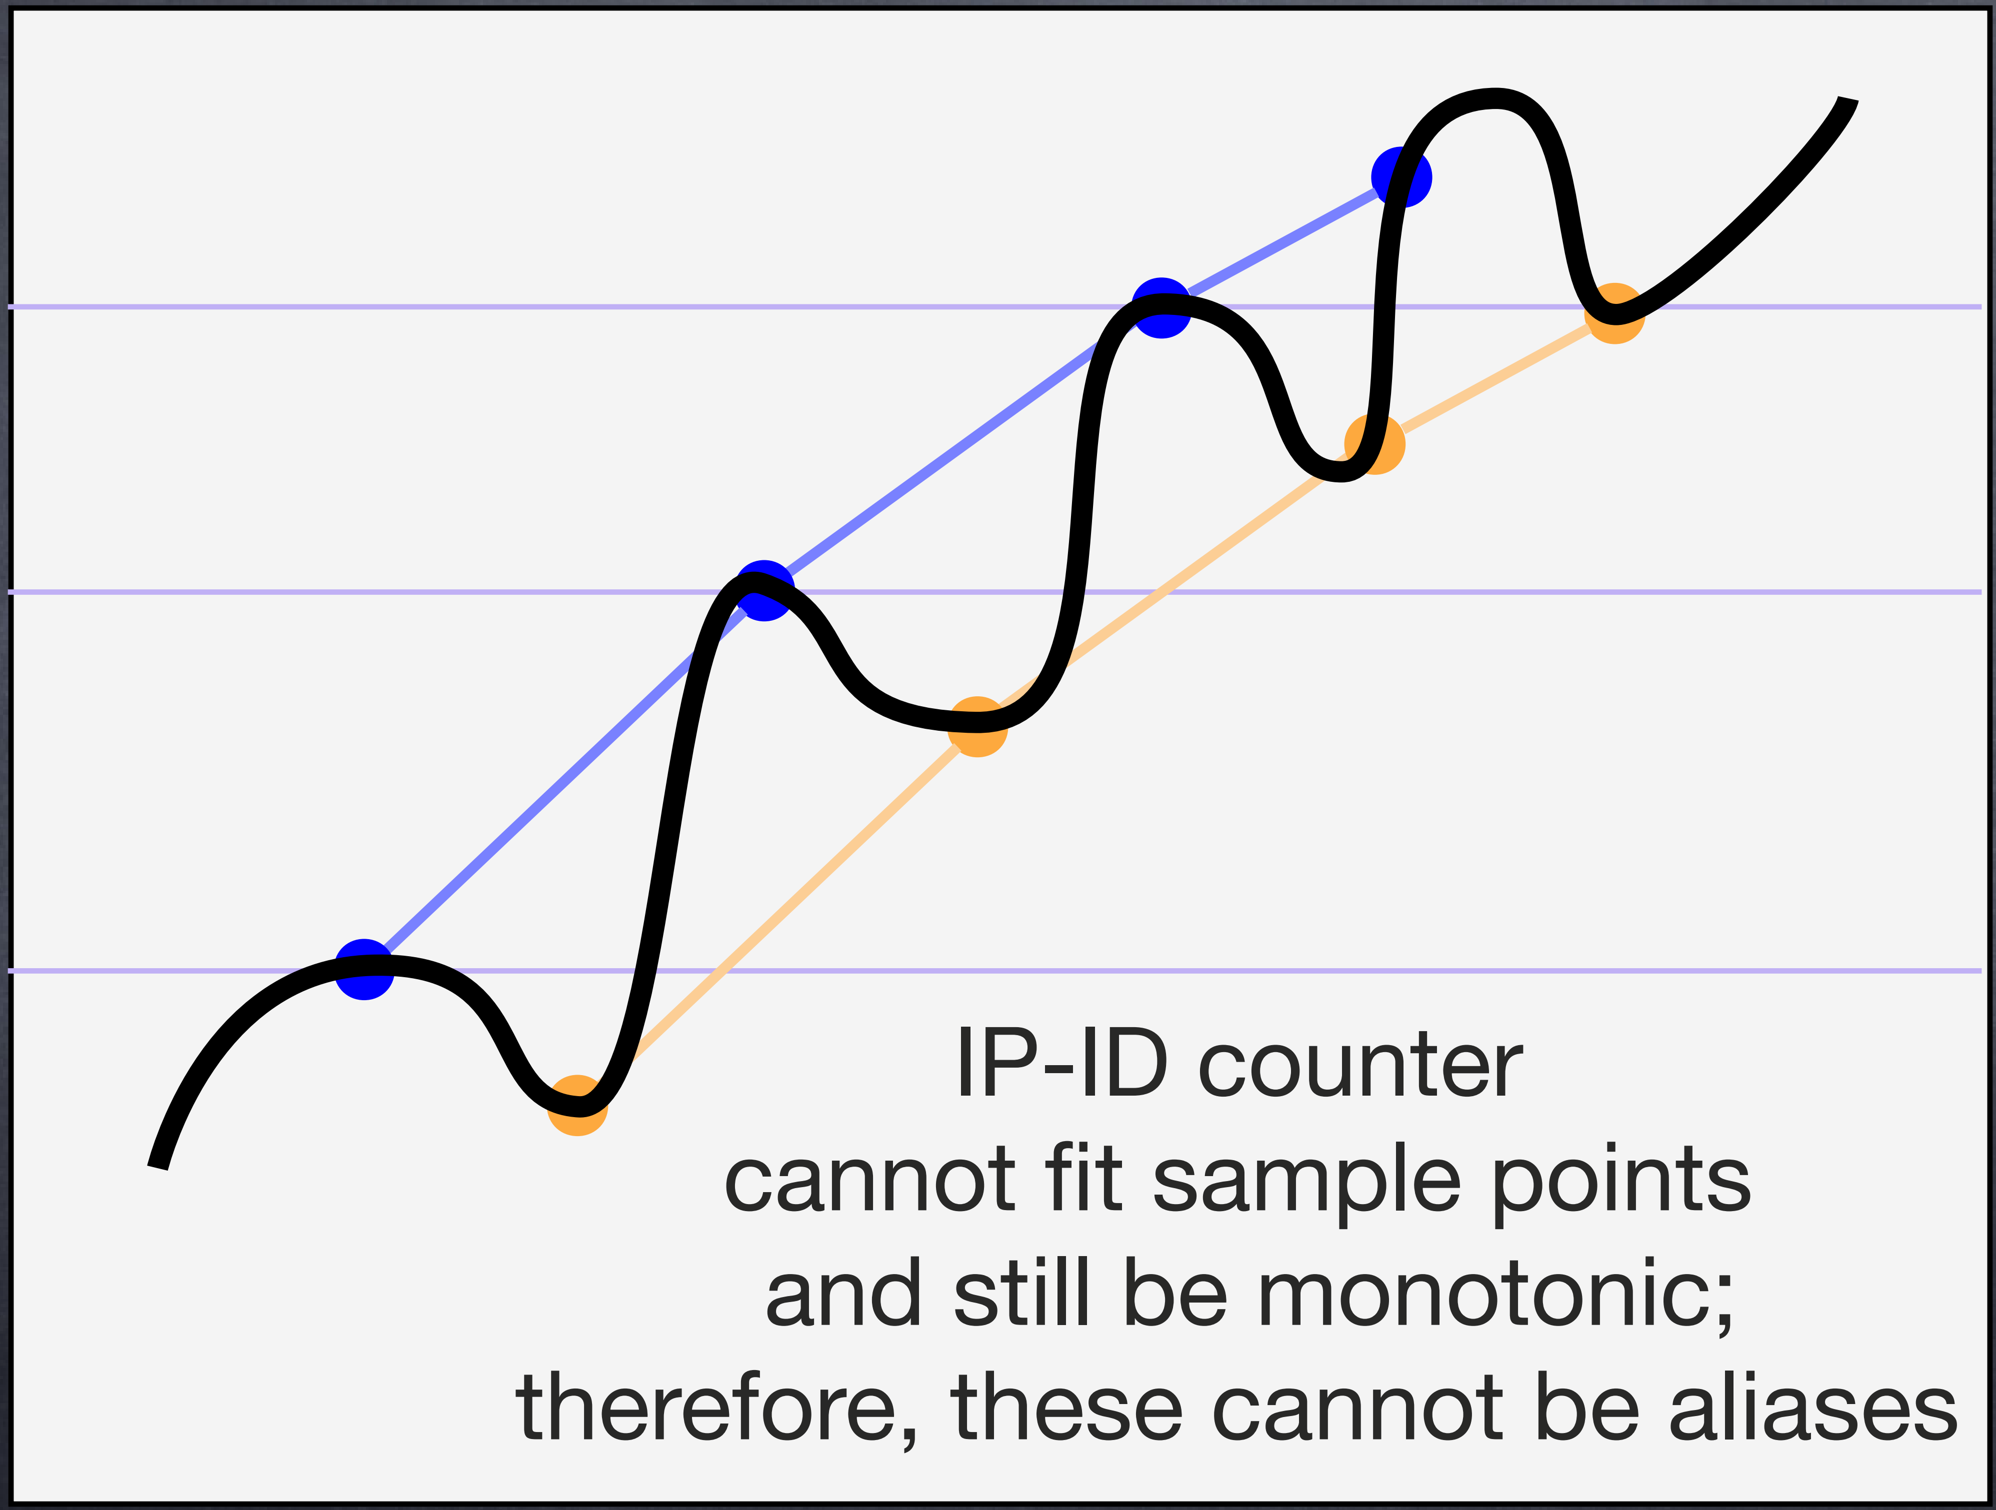

time

RadarGun Issues

- **applying RadarGun to 1 million addresses is problematic because RadarGun needs overlapping IP-ID time series for all targets in a short period of time**
  - looks like DDoS attack
  - triggers rate limiting
  - requires high probing rate or large number of machines

RadarGun Issues

$$\frac{\text{interface set size}}{\text{probing rate}} = \text{round duration}$$

or

$$\frac{\text{interface set size}}{\text{round duration}} = \text{probing rate}$$

- probing rate must increase if ...
  - interface set size increases
  - round duration decreases

MIDAR

- ***M**onotonic **ID**-Based **A**lias **R**esolution* (MIDAR) is our extension of the RadarGun approach
  - *monotonic bounds test* for accurate testing of pairs
  - *sliding window* for scaling up probing
  - 4 probing methods
  - multiple monitors
  - two stages: discovery (estimation, sliding window), corroboration (hours later)

MIDAR Results

- discovery stage (sliding window):
  - probed 1.0 million addresses
  - 486 **billion** pairs compared
  - shared pairs found: 1.6 million (0.00093%)
  - 55k alias sets containing 497k addrs
- corroboration stage:
  - shared pairs found: 428k (26% of discovery stage)
    - not actually 1.2 million false positives; inflated by human error
  - 69k alias sets containing 186k addrs
    - stable across multiple corroboration runs

MIDAR Results

- consistency check: out of 69k sets,187k addrs, 428k pairs after corroboration ...
  - every pair inferred by transitive closure was tested with the monotonic bounds test at least once and passed every time
  - all but 80 pairs were tested at least twice and passed every time
  - only 12 sets (49 addrs) contained transitive *closure conflicts:*

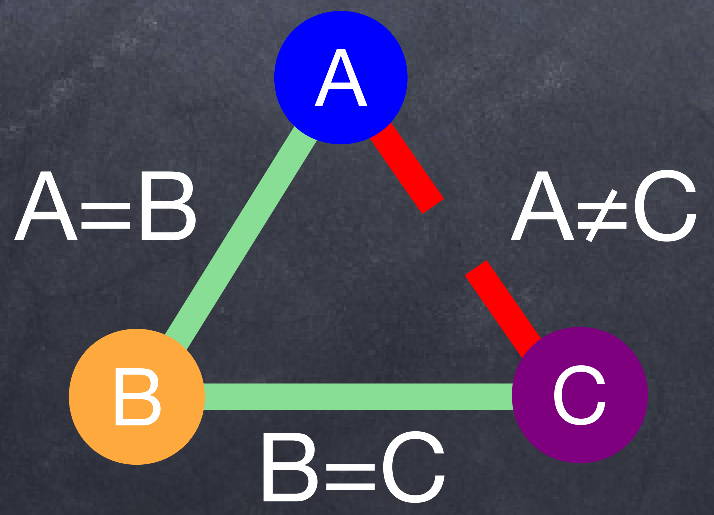

We suspect real network change caused these conflicts and not false positives.

MIDAR Validation

- we compared MIDAR results to ground truth for a tier 1 ISP
  - for comparison, we only consider routers that appear with multiple interfaces in Ark traces
    - observed multi-interface routers (OMIRs)
- **0 false positives**

	full ISP topology	OMIRs	MIDAR
routers	1,986	983	434
addresses	24,429	4,008	1,284
pairs	611,407	16,900	2,133

Topology mapping: future work

- MIDAR improvements
  - adapt corroboration spacing to responsiveness
- MAARS: Multi-Approach Alias Resolution System
  - combine MIDAR, kapar, iffinder (and others?)
- AS-router Dual graph, including regular updates
- Release supporting tools under GPL
- Support additional collaborators' experiments

Statistics Pages

- per-monitor analysis of IPv4 topology data

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

AS dispersion by AS hop

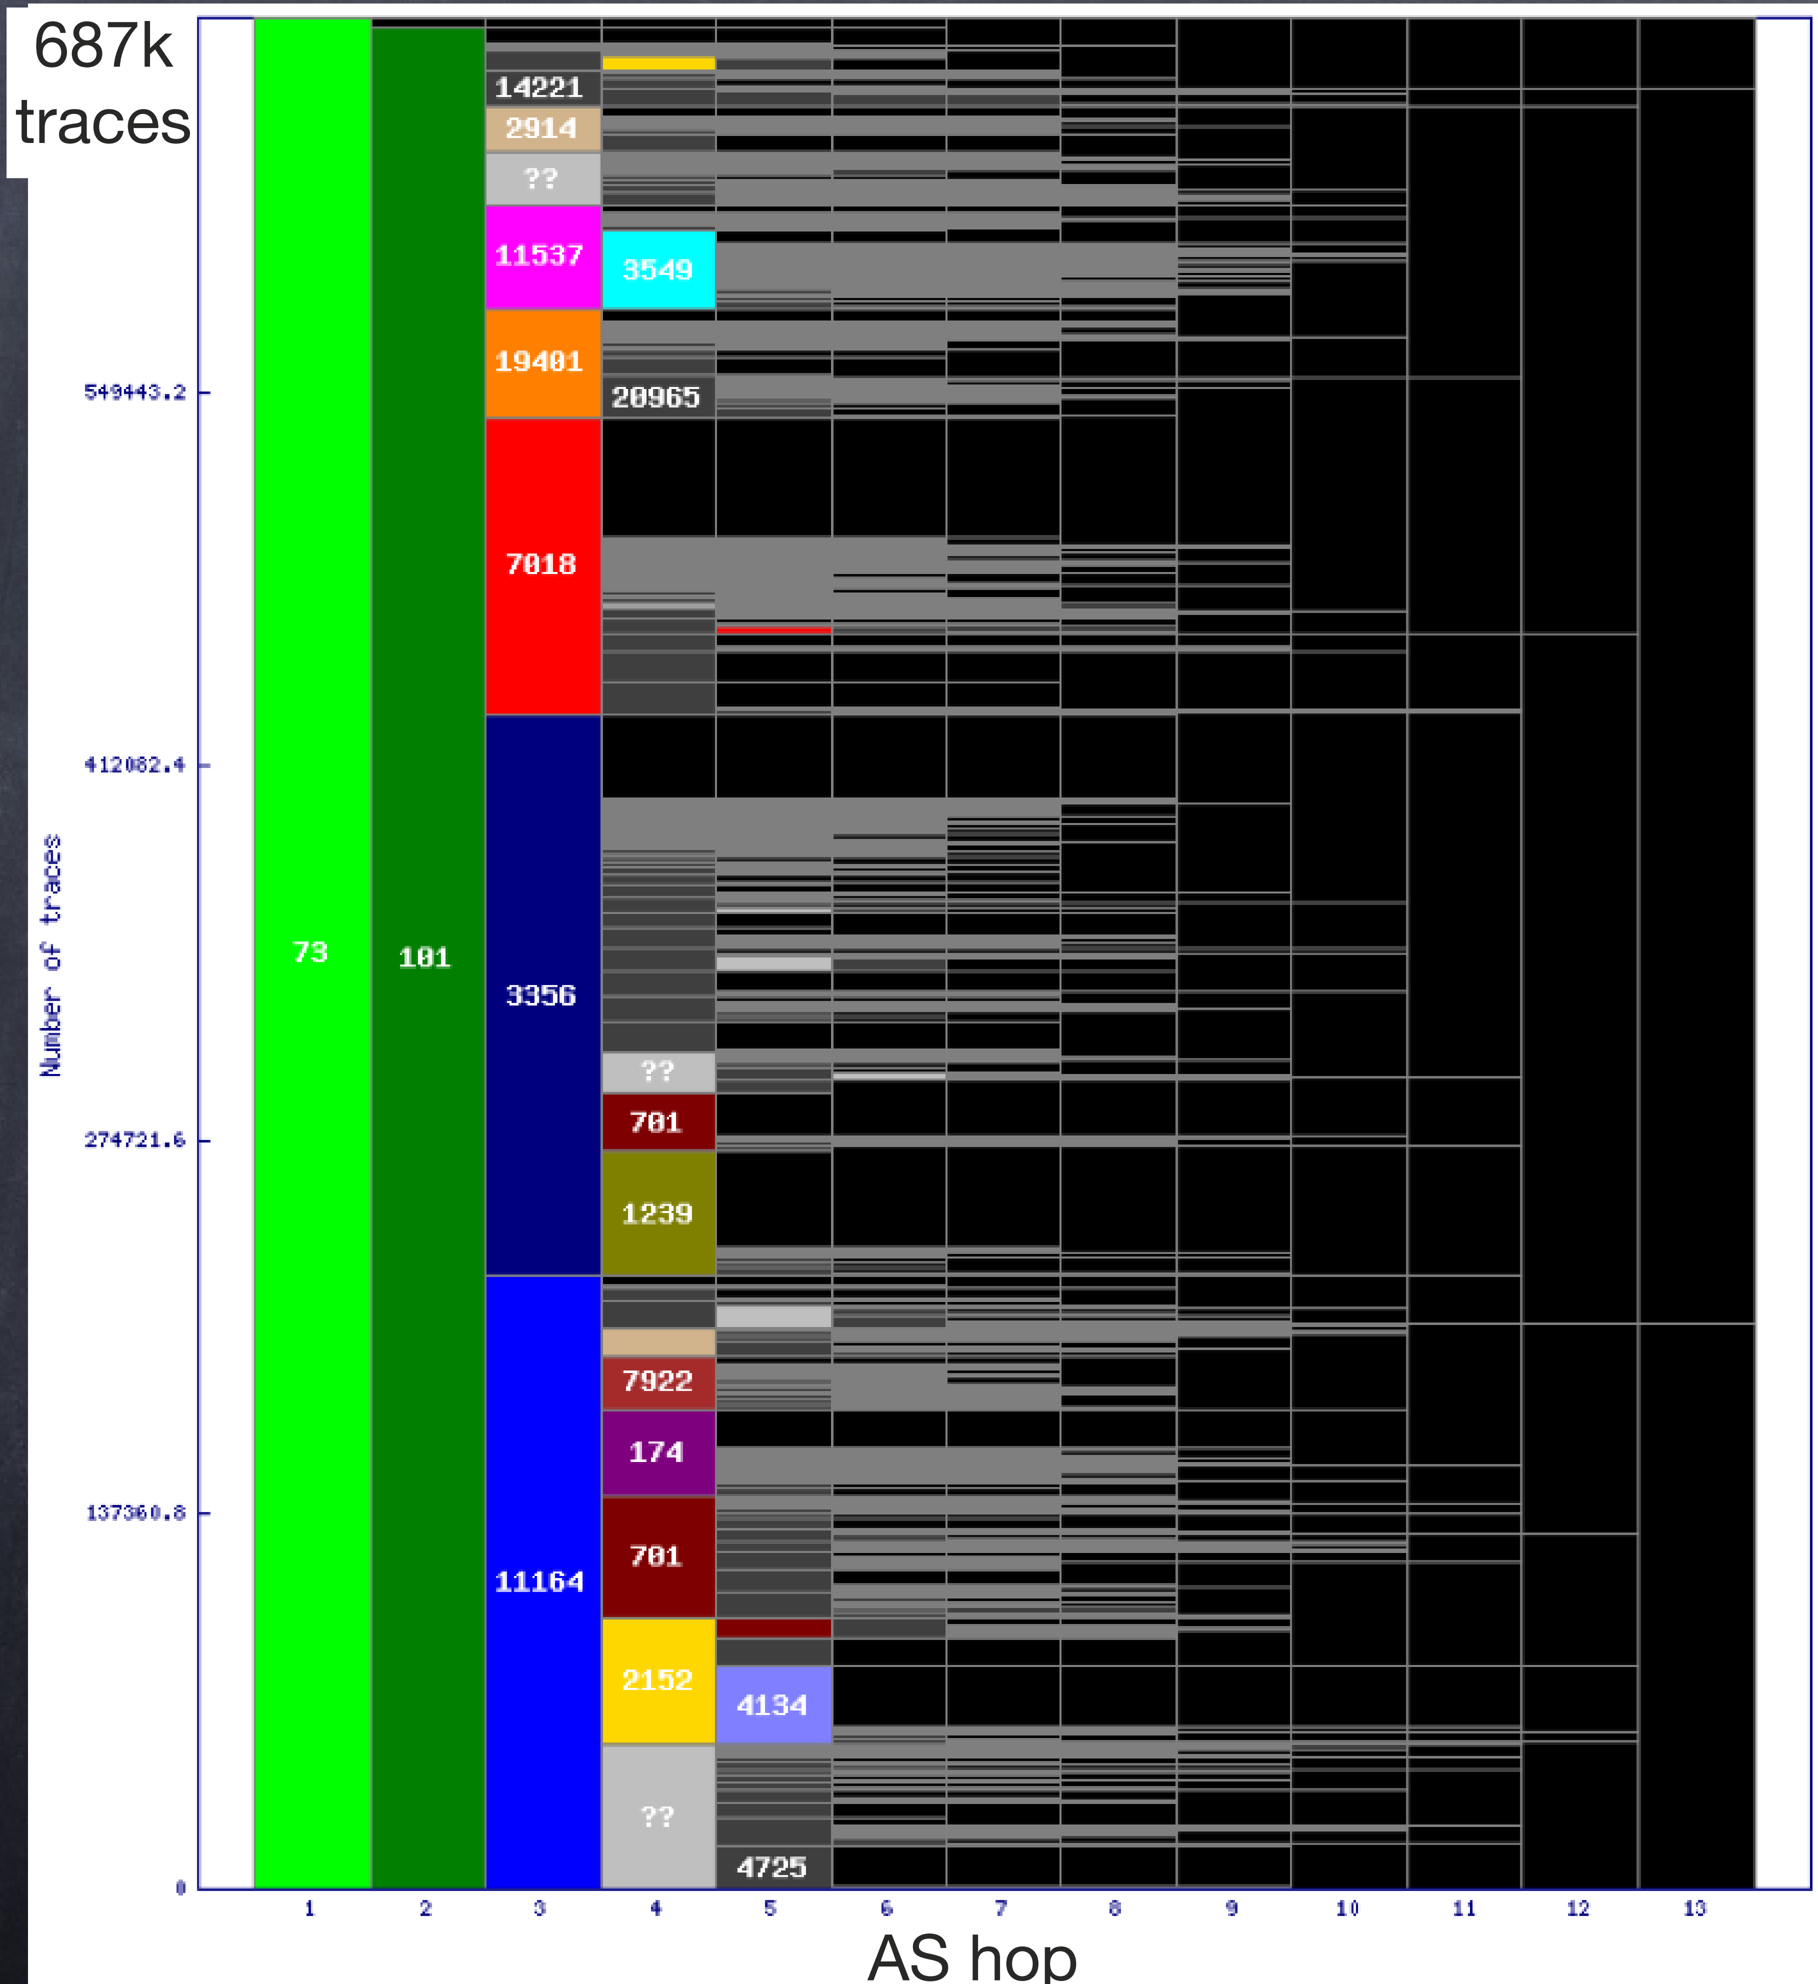

- 73 WASHINGTON-AS - University of Wash
- 101 WASH-NSF-AS - University of Washingt
- 11164 TRANSITRAIL - National LambdaRail, L
- 3356 LEVEL3 Level 3 Communications
- 7018 ATT-INTERNET4 - AT&T WorldNet Serv
- 701 UUNET - MCI Communications Service
- 2152 CSUNET-NW - California State Universi
- 1239 SPRINTLINK - 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

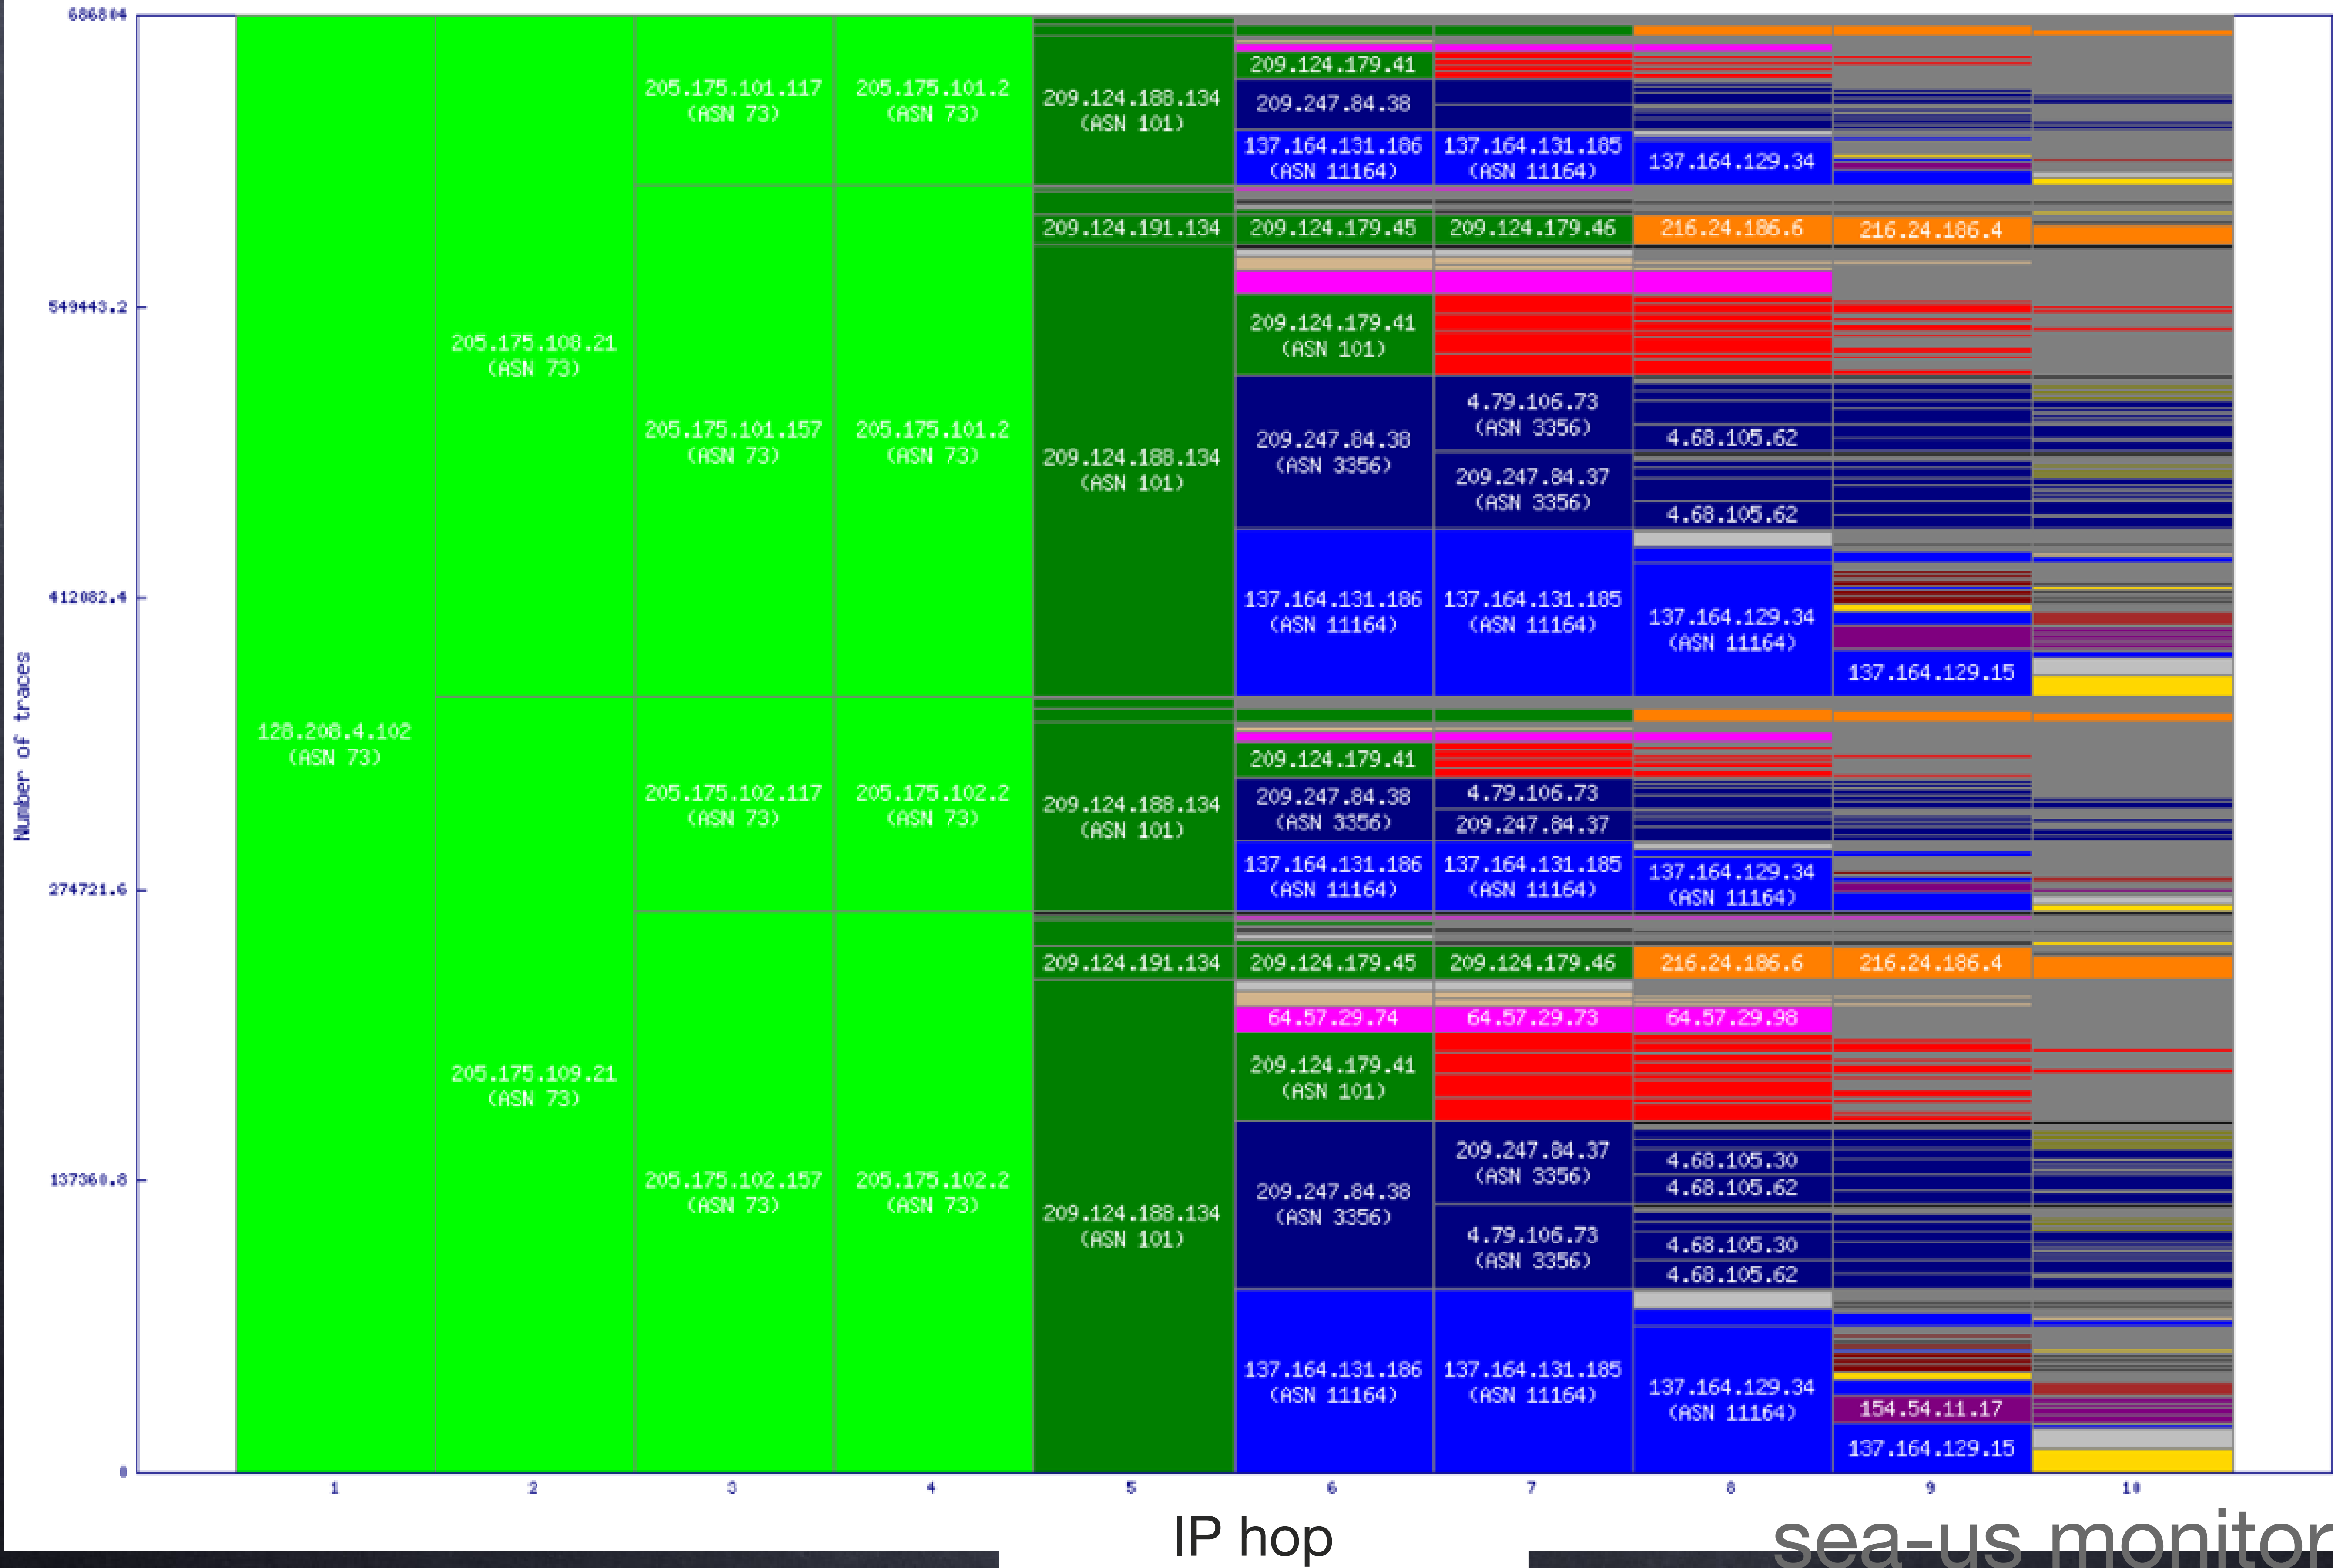

sea-us monitor

AS dispersion by IP hop: see load balancing

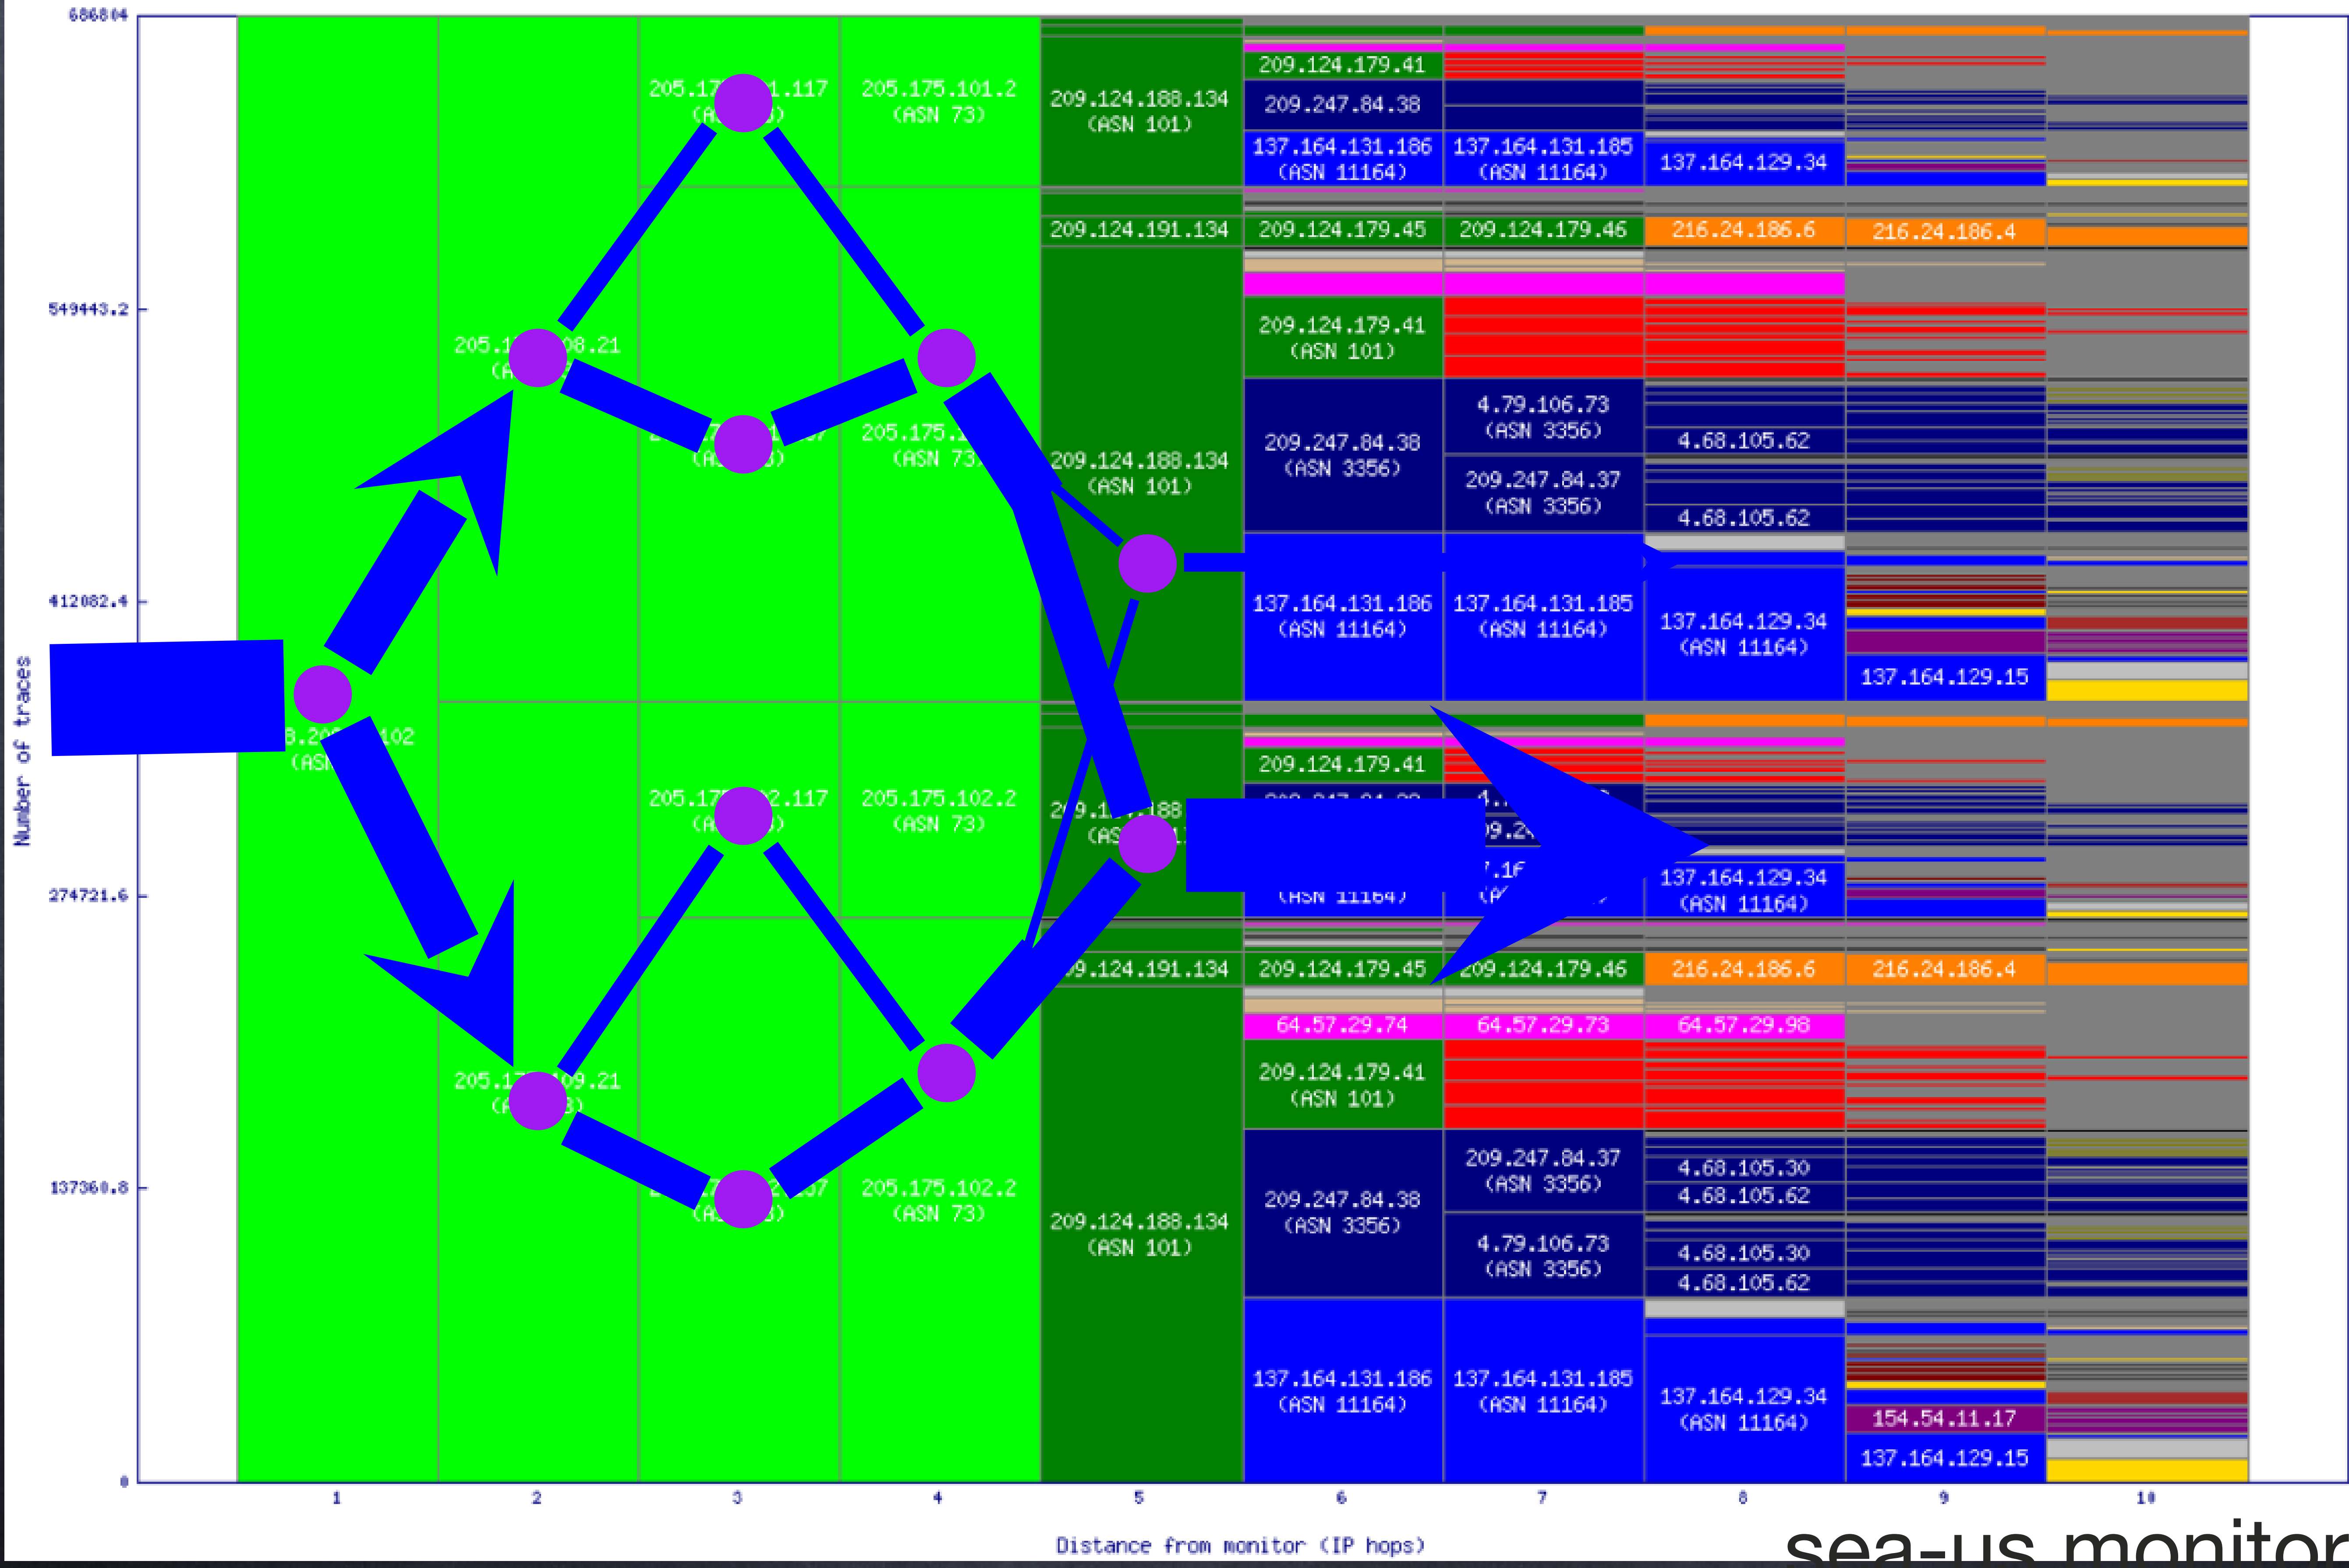

sea-us monitor

Statistics Pages

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

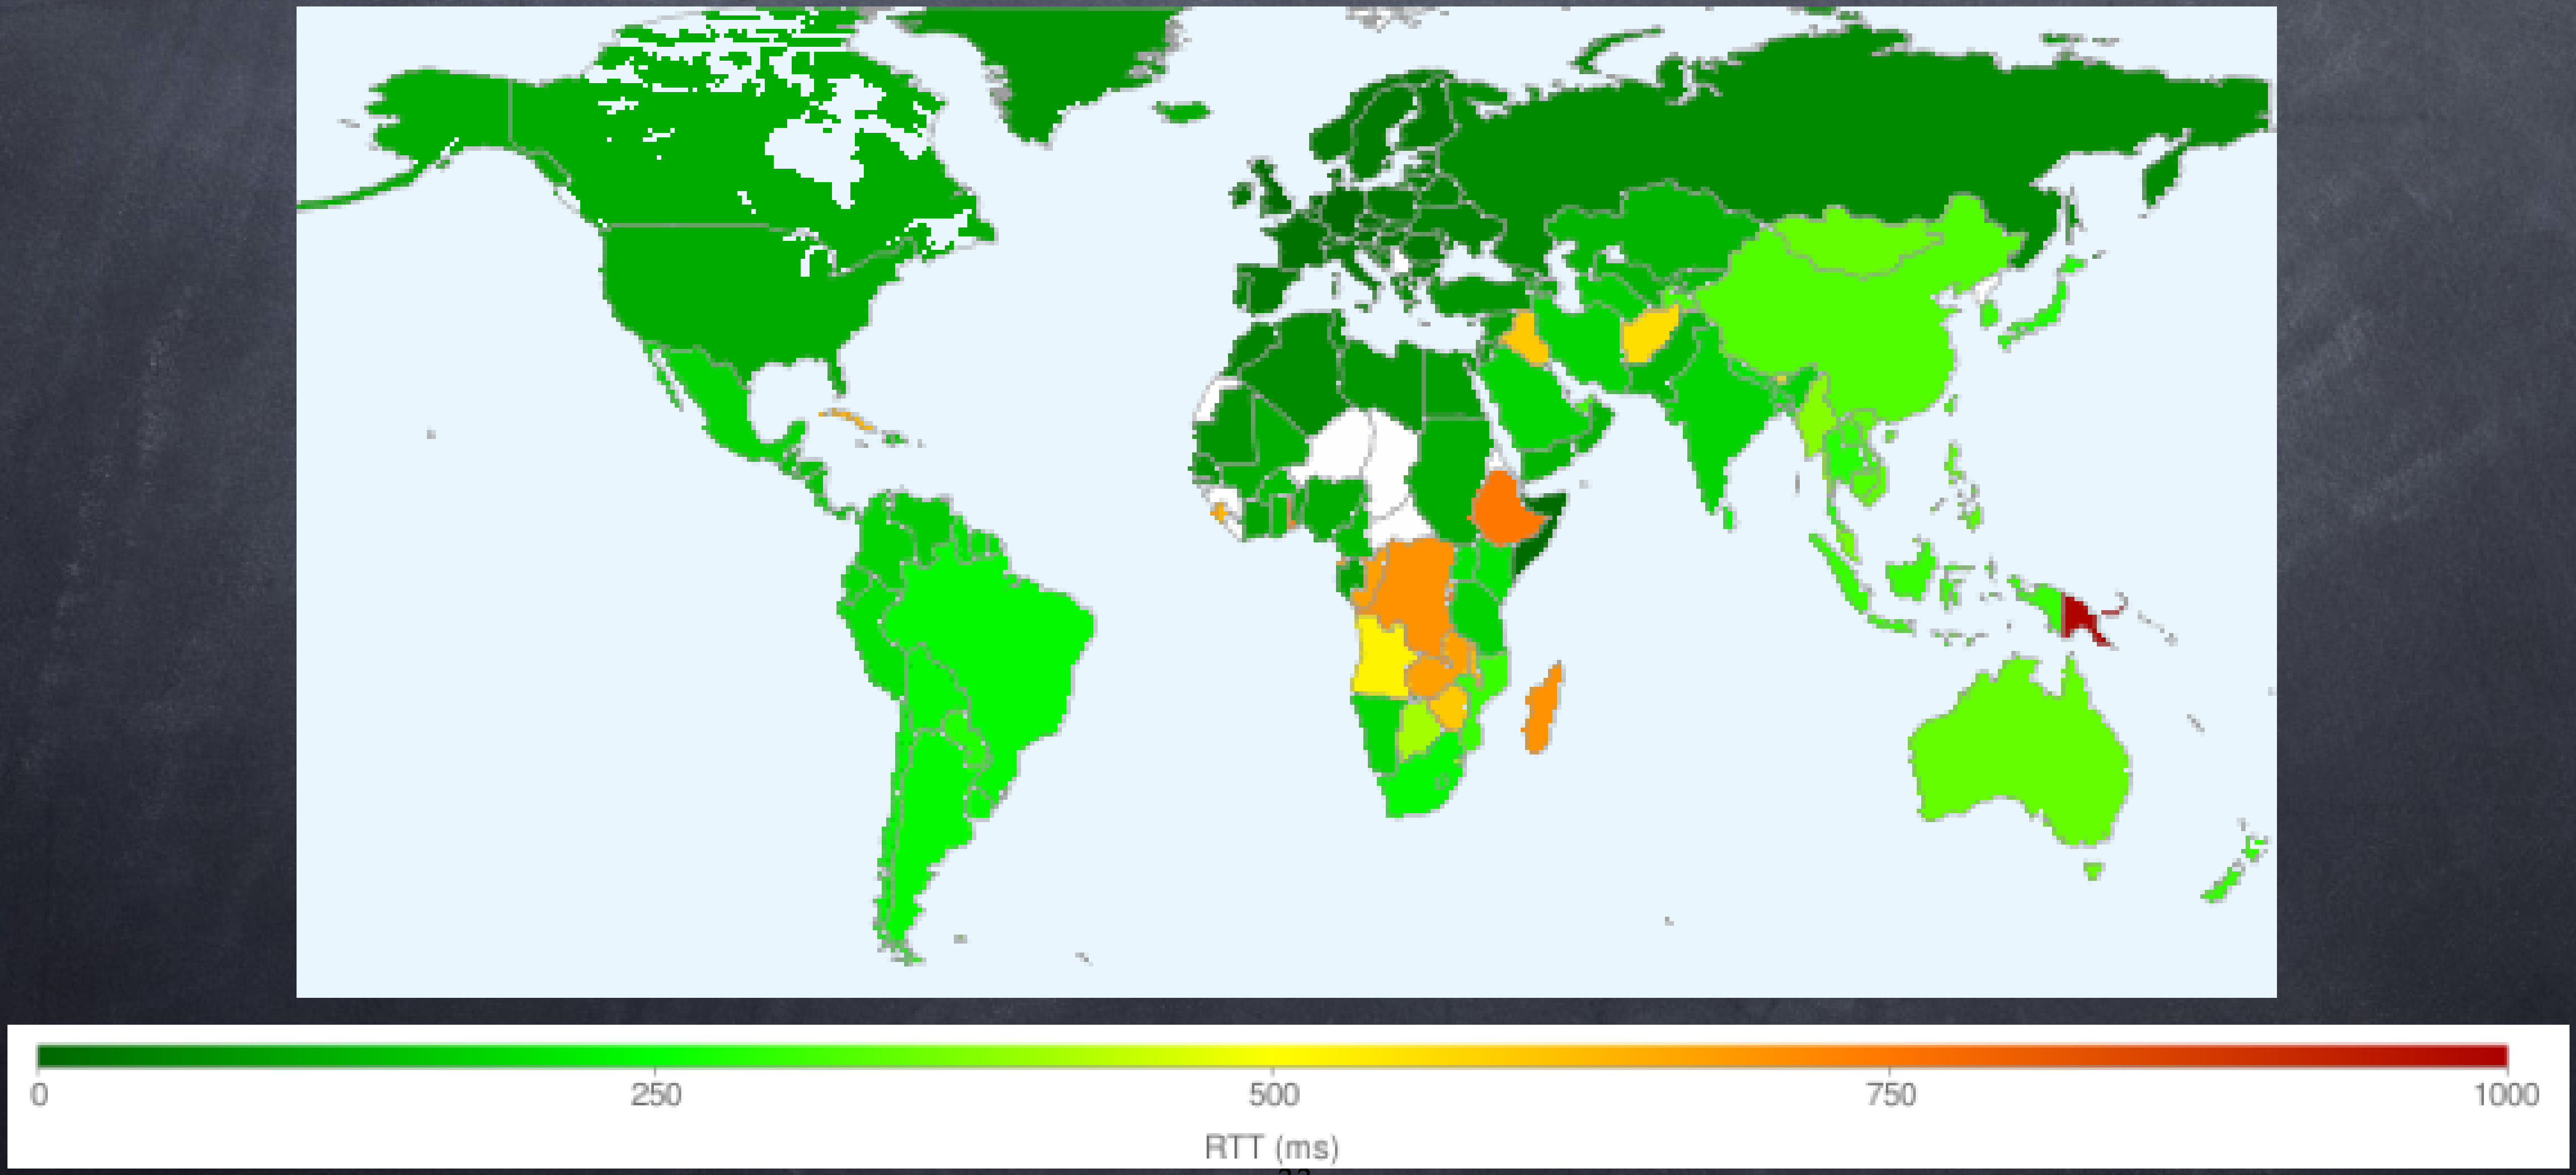

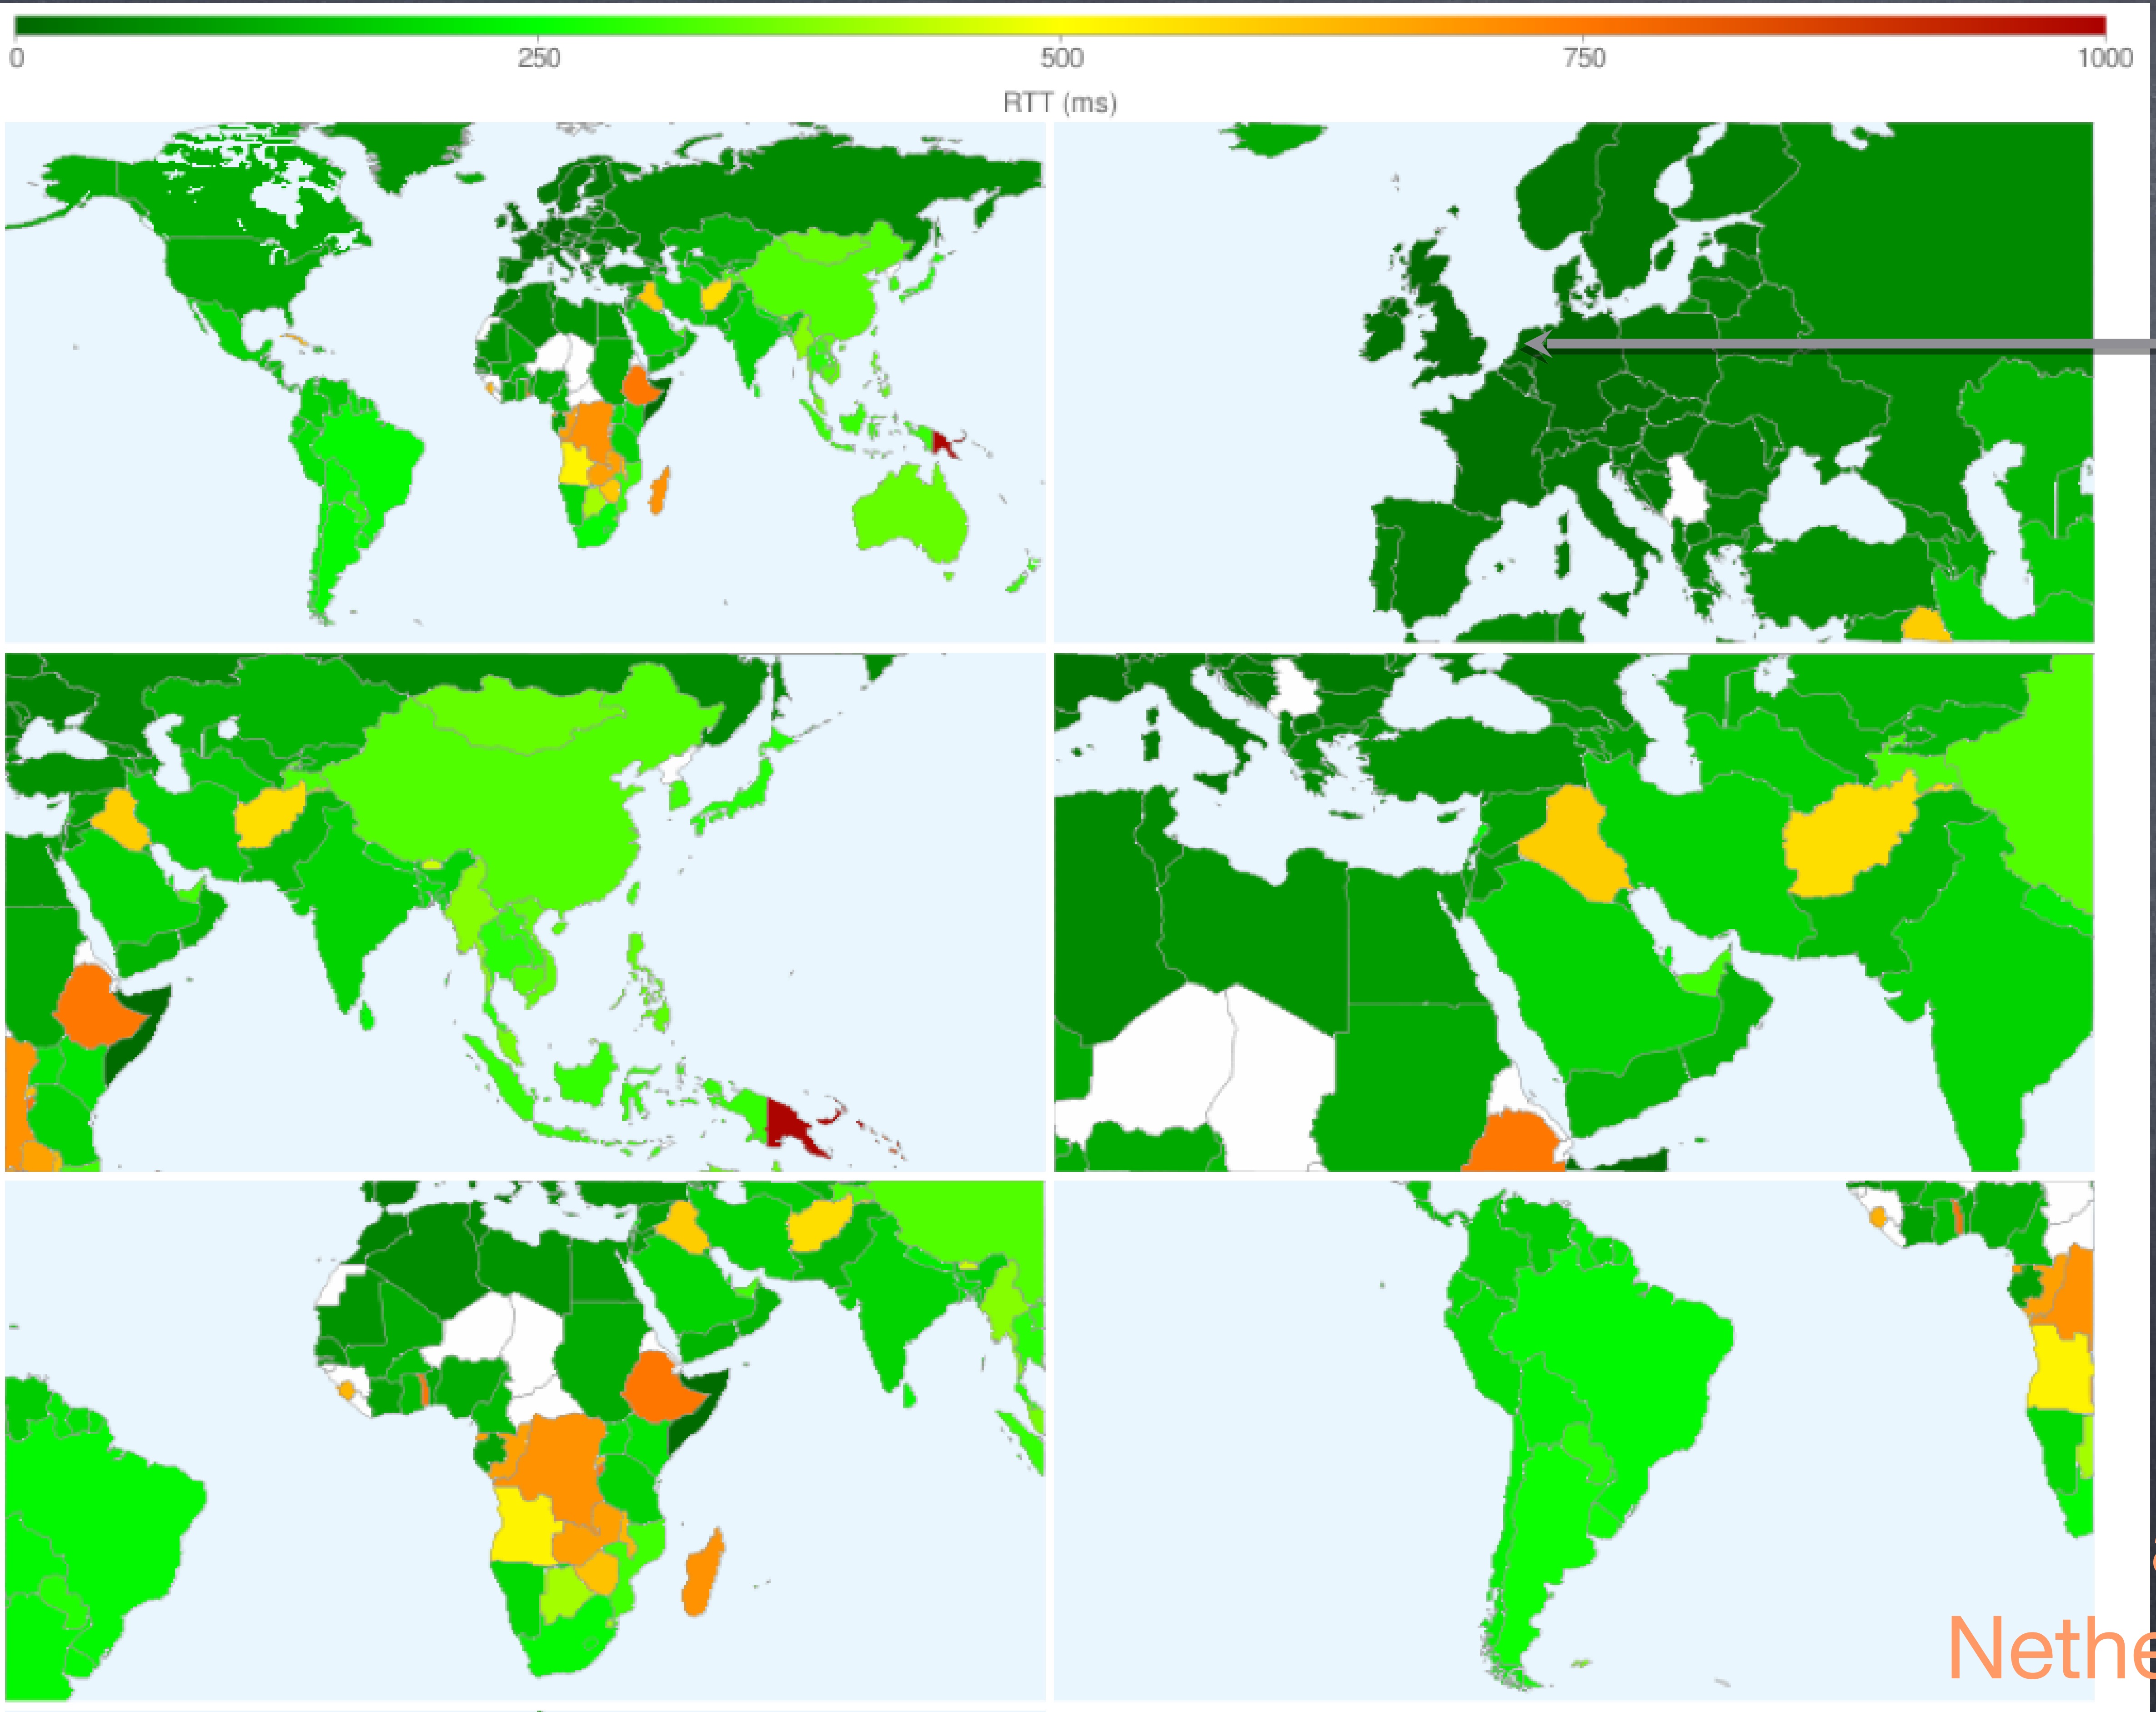

view from ams-nl Netherlands

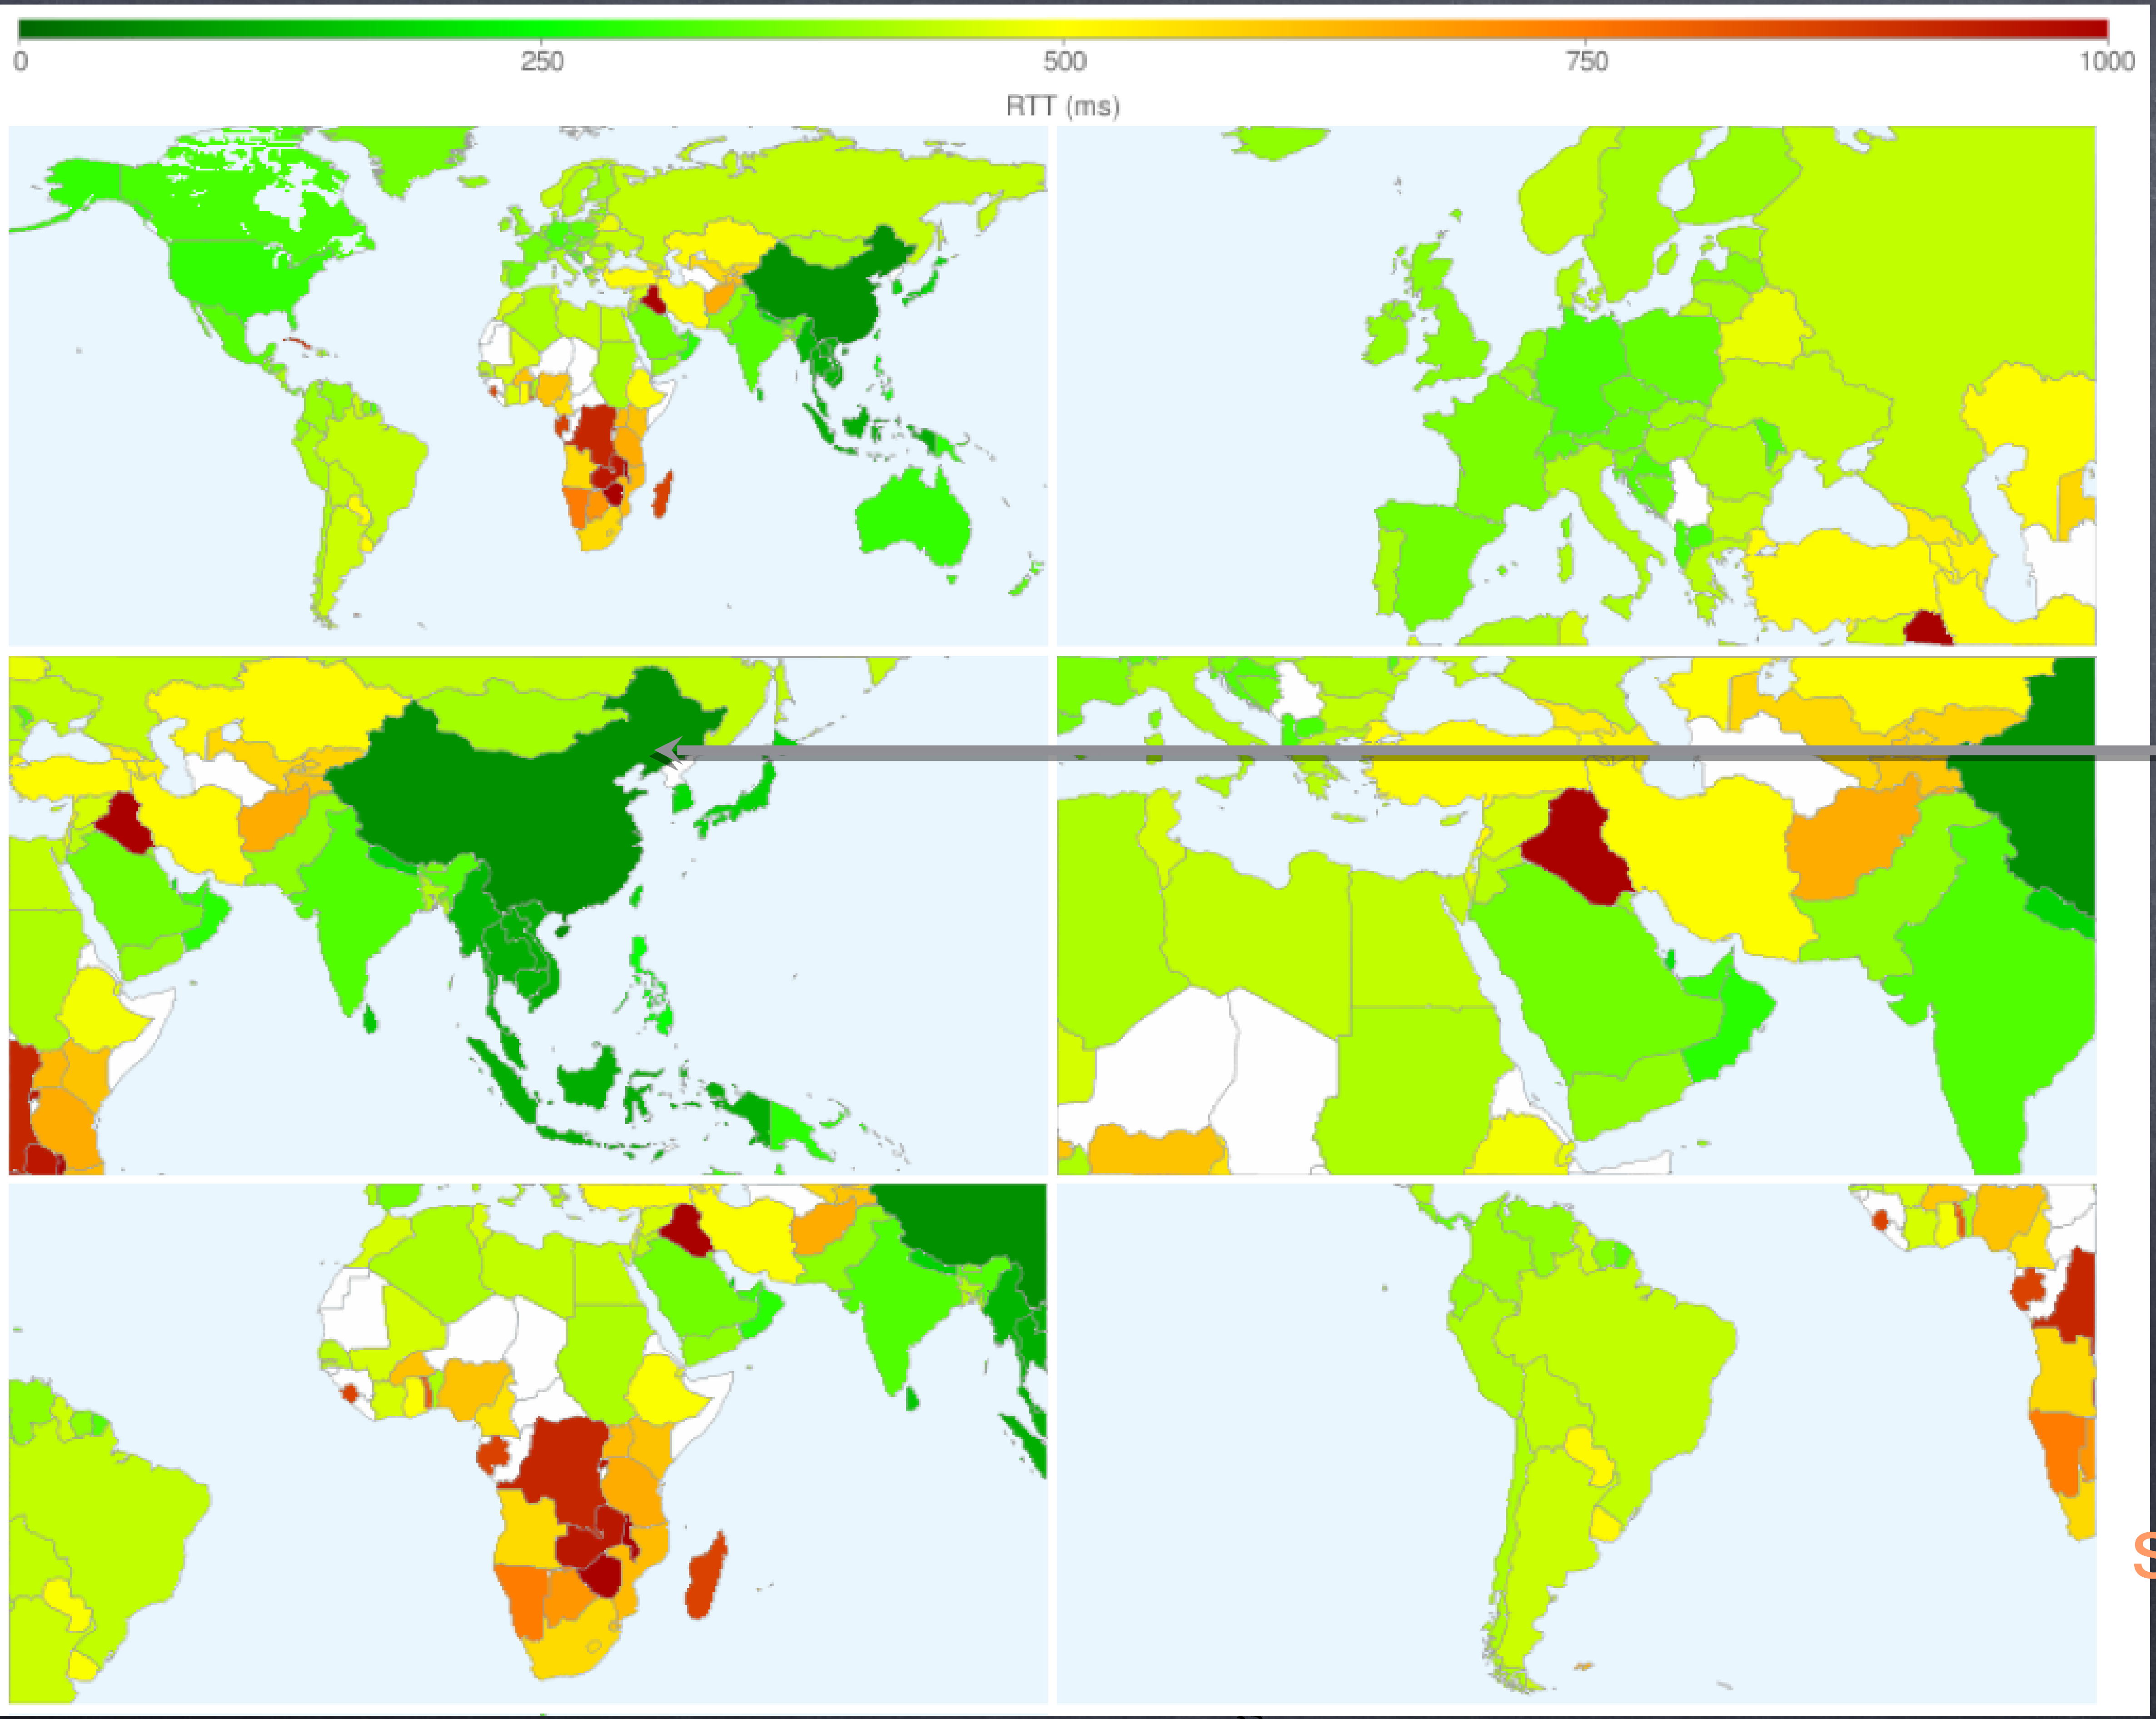

view from she-cn China

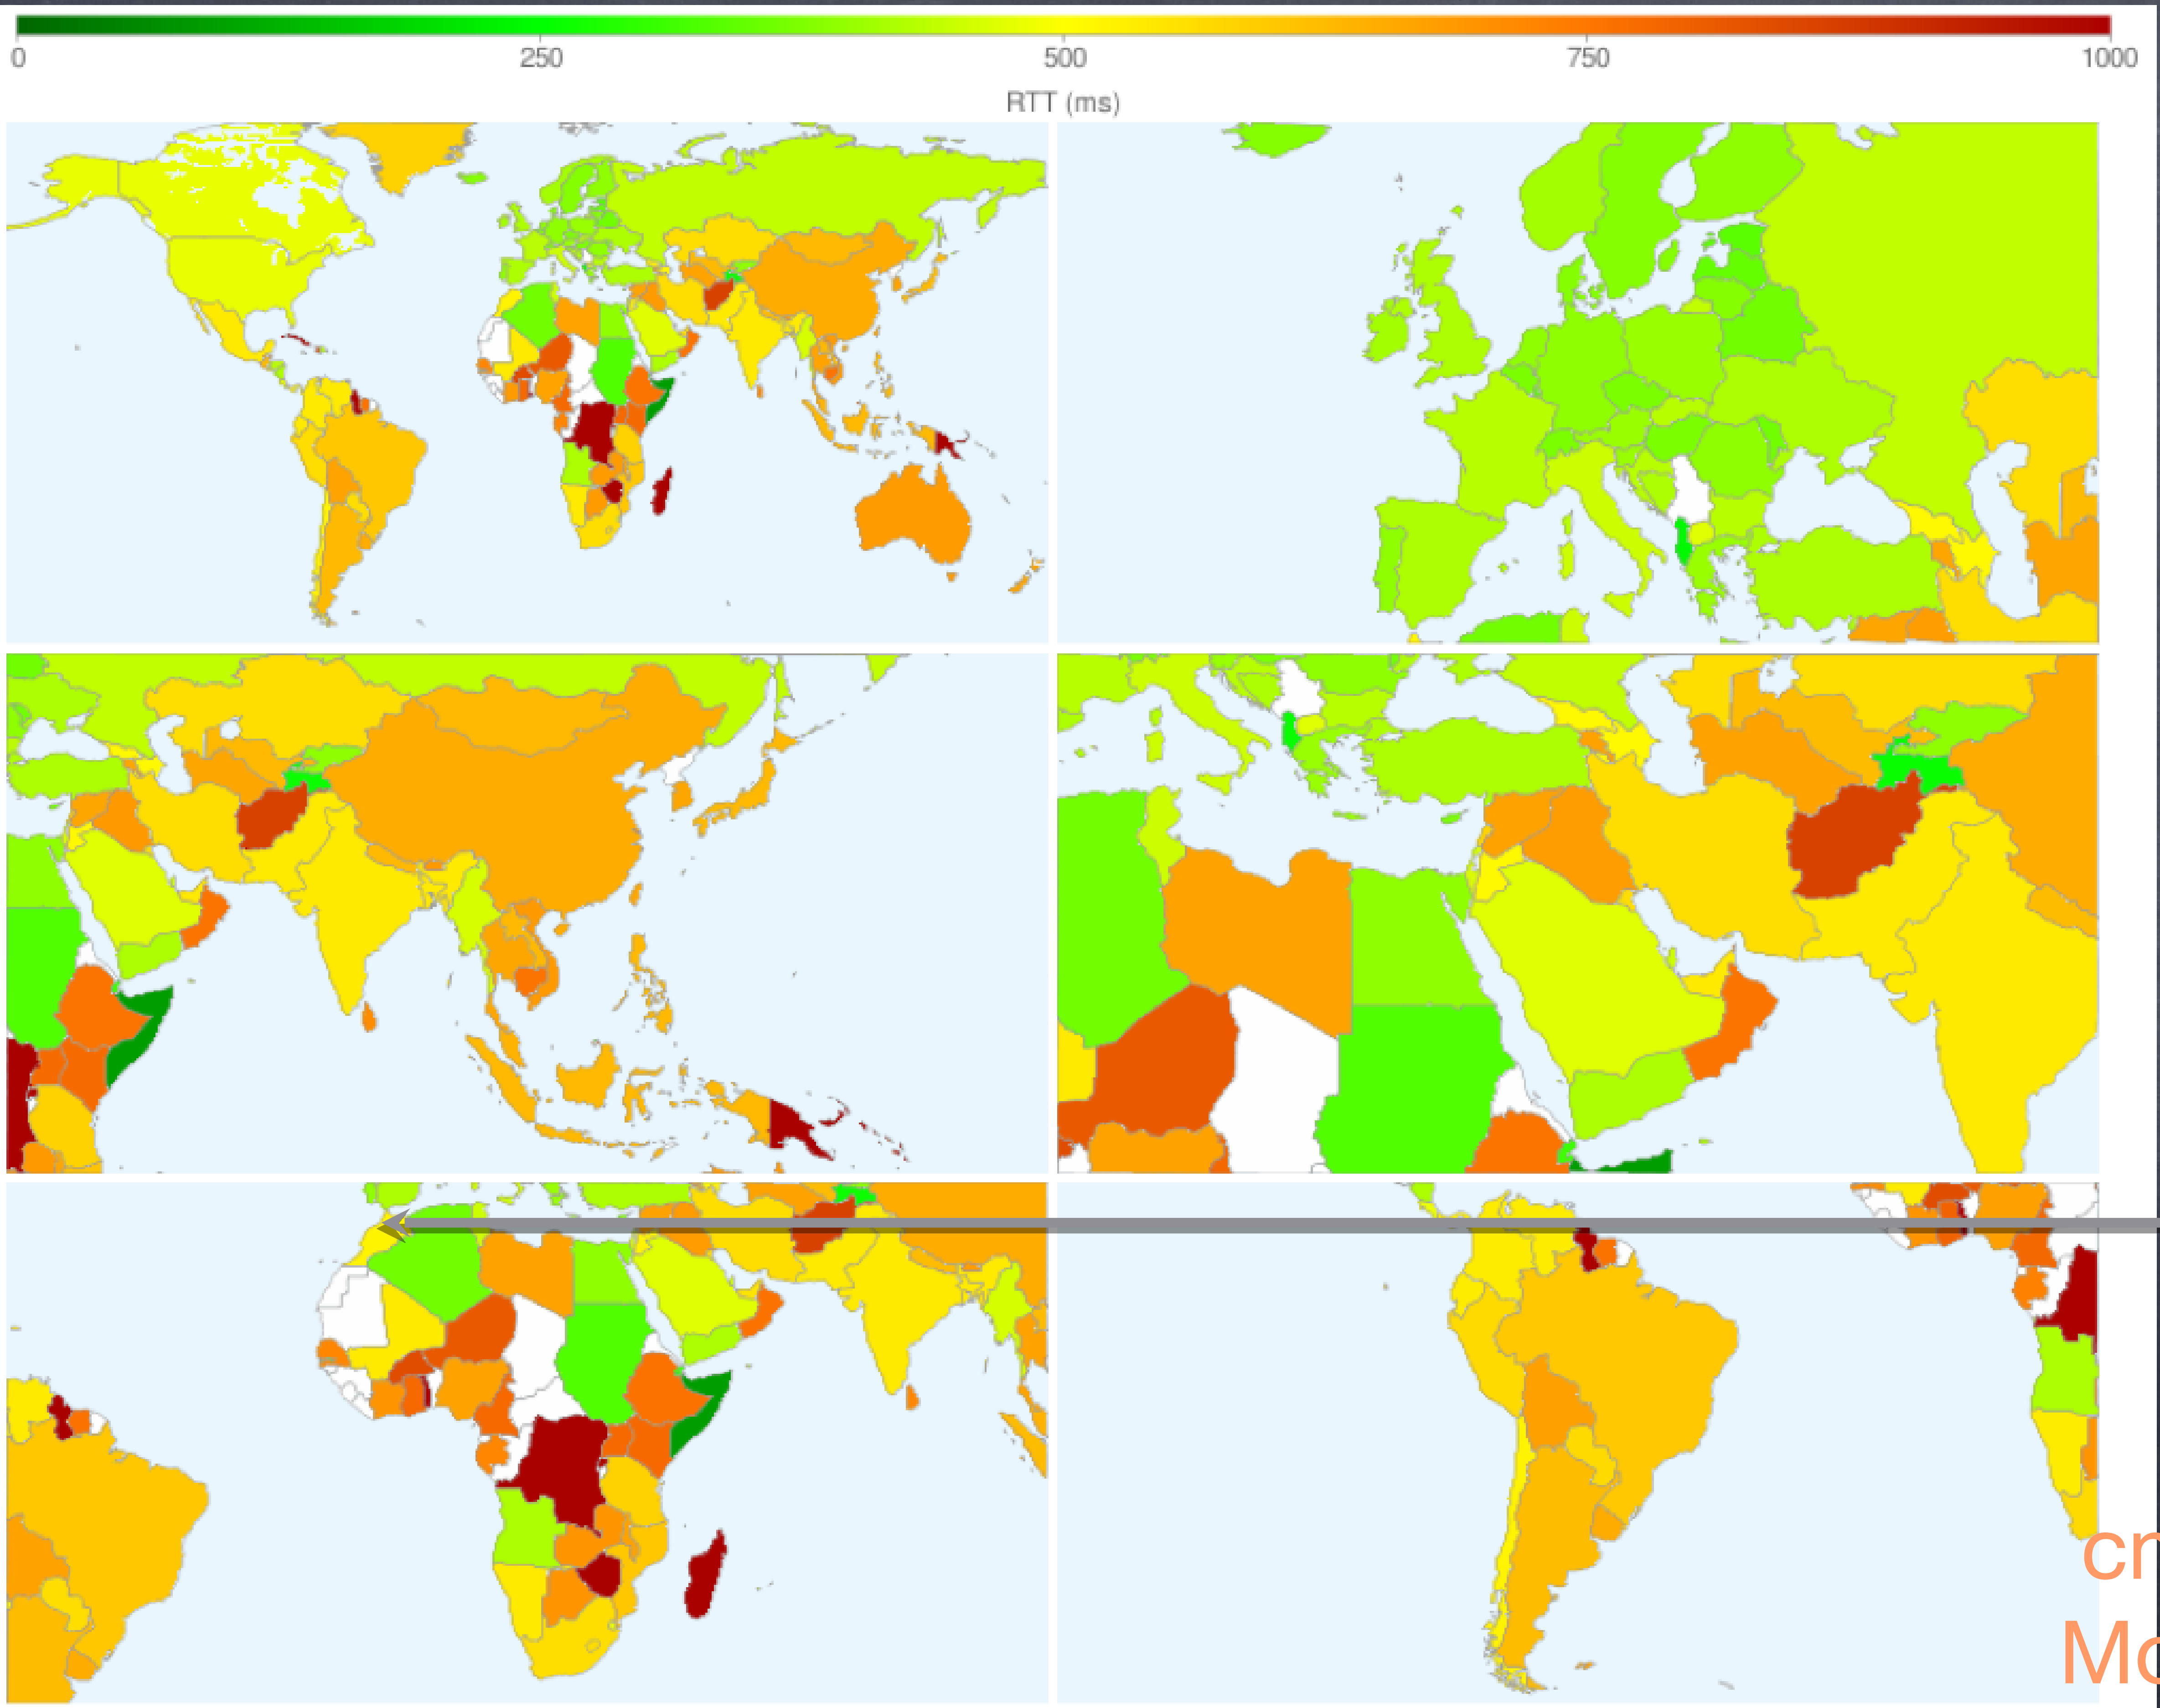

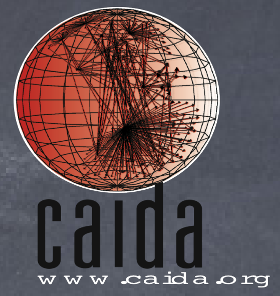

view from cmn-ma Morocco

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.

IP Dispersion by IP Hop

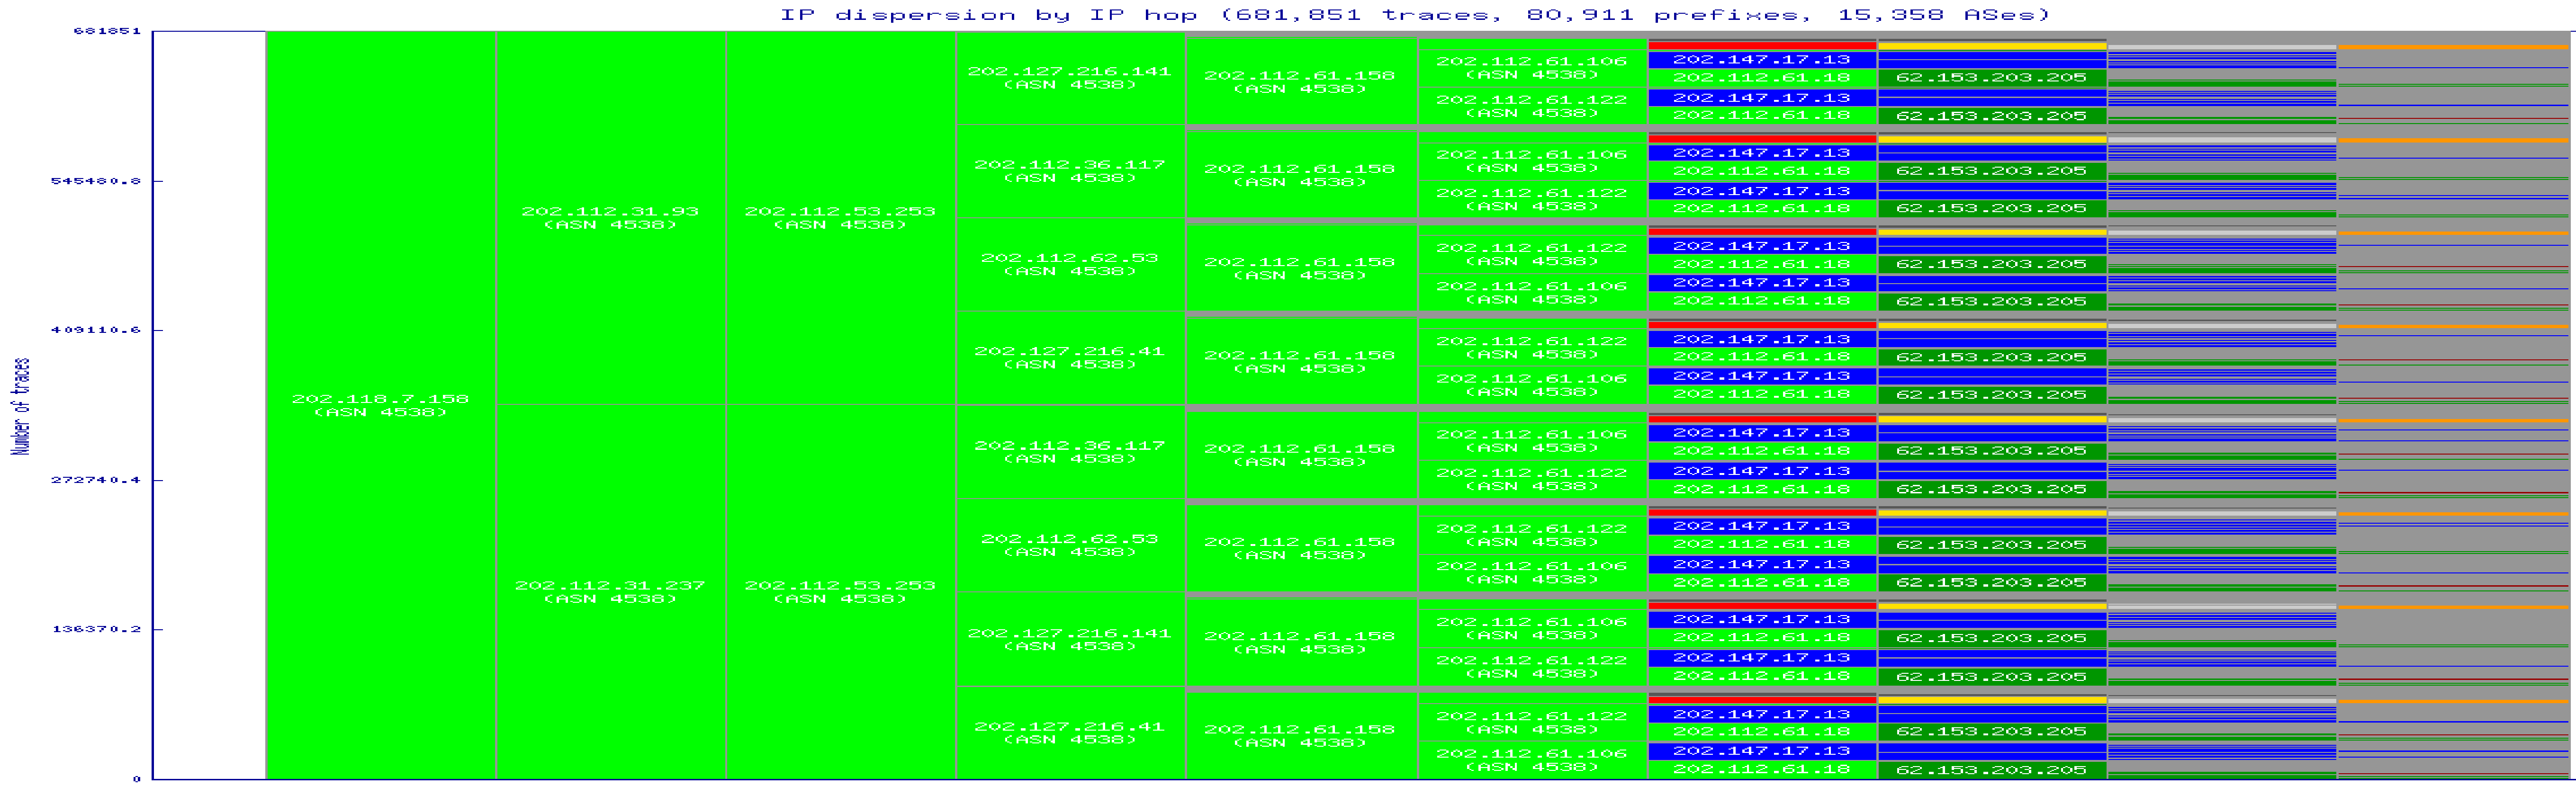

IP Dispersion by IP Hop

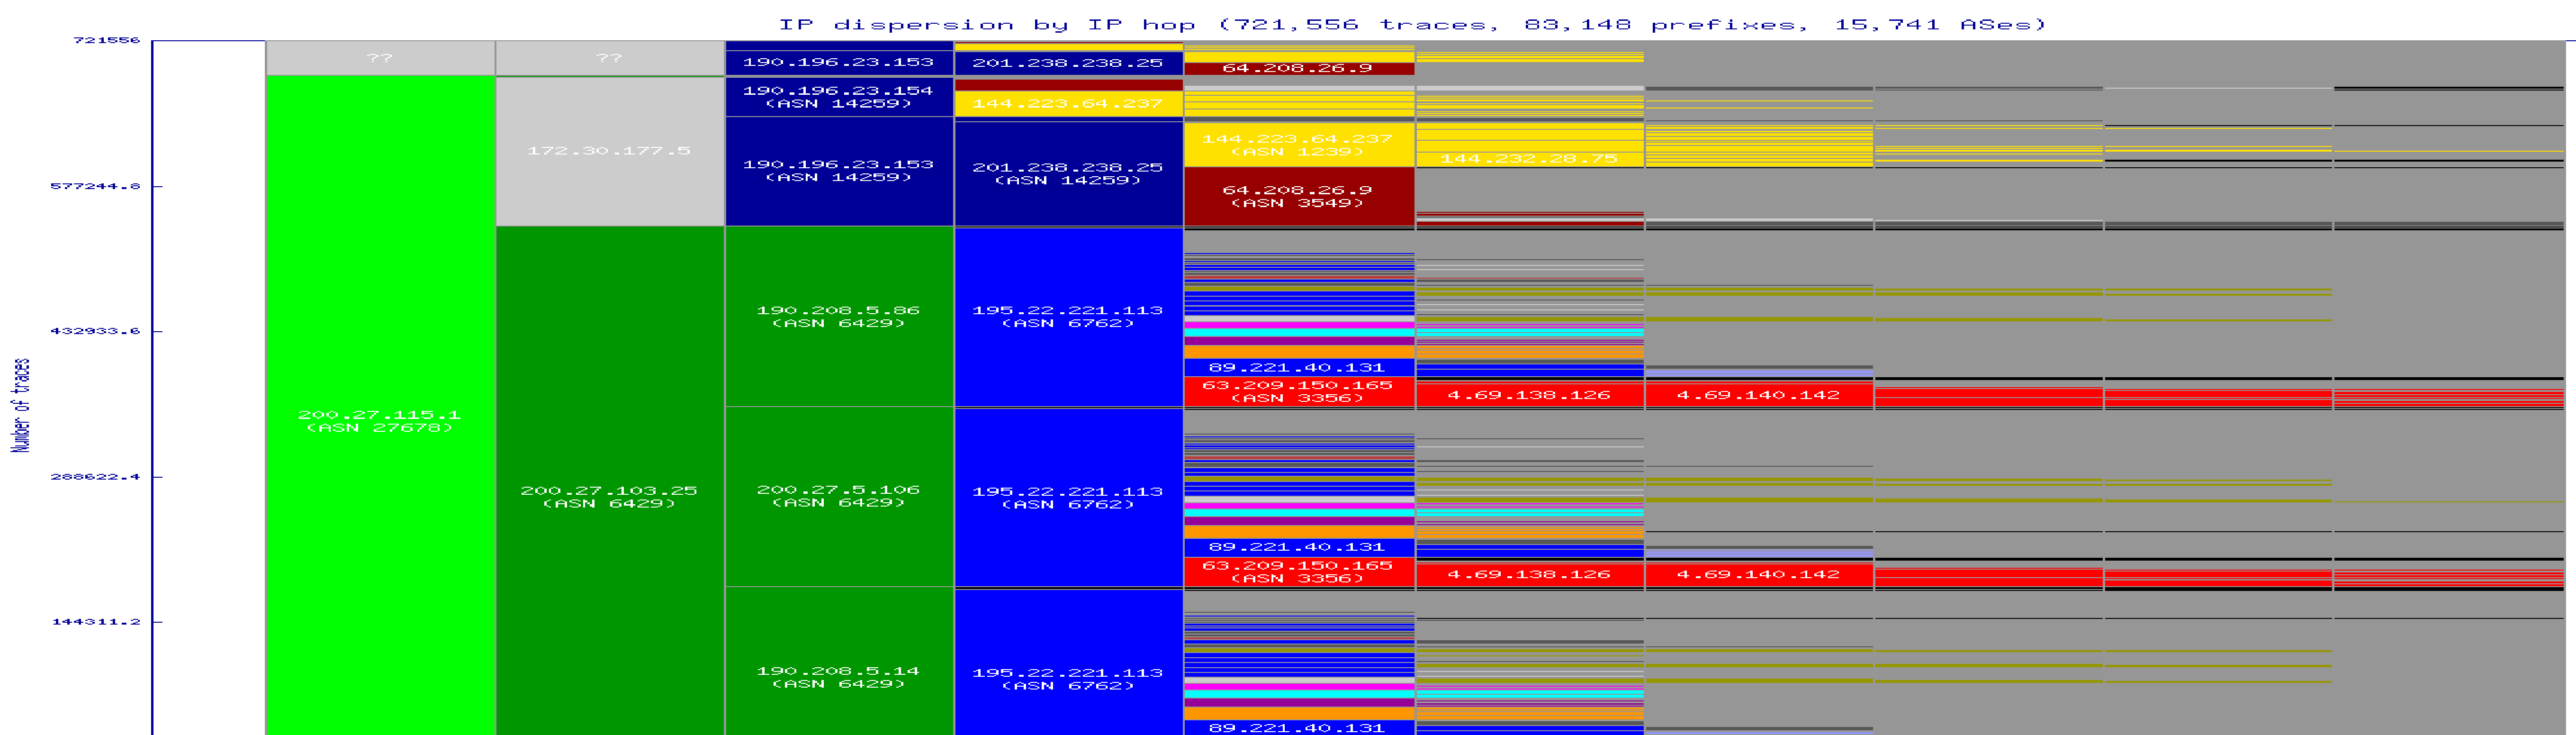

Schedule, Planned activities

- 1-2 monitors/month
- IPv4, IPv6 topology data
- Release dual-graph as part of ITDK
  http://www.caida.org/data/active/internet-topology-data-kit/
- Continue alias resolution study, regular updates
- Visualization (in support of)
- Validation against ground truth
- AIMS 2011
- Begin work on 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

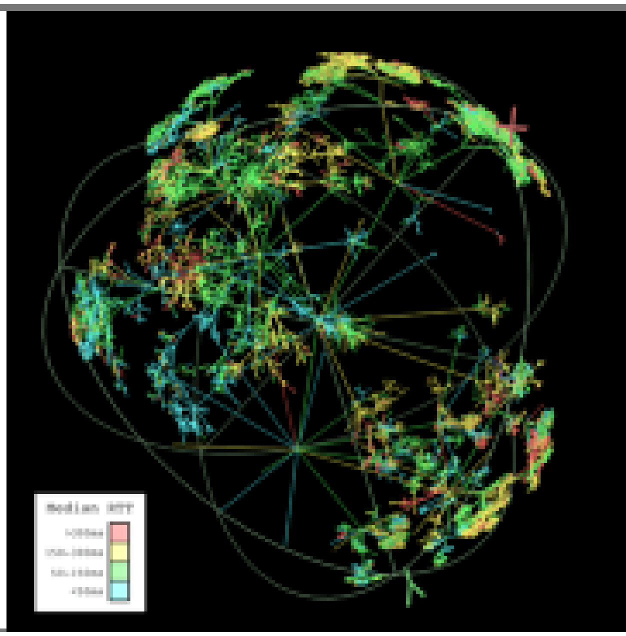

Walrus visualizations of round-trip time measurements made by CAIDA's macroscopic Internet topology monitor located in Herndon, VA, USA.

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/

- Macroscopic Internet Topology Data Kit (ITDK)
http://www.caida.org/data/active/internet-topology-data-kit/

- topostats
http://www.caida.org/tools/utilities/topostats/

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