
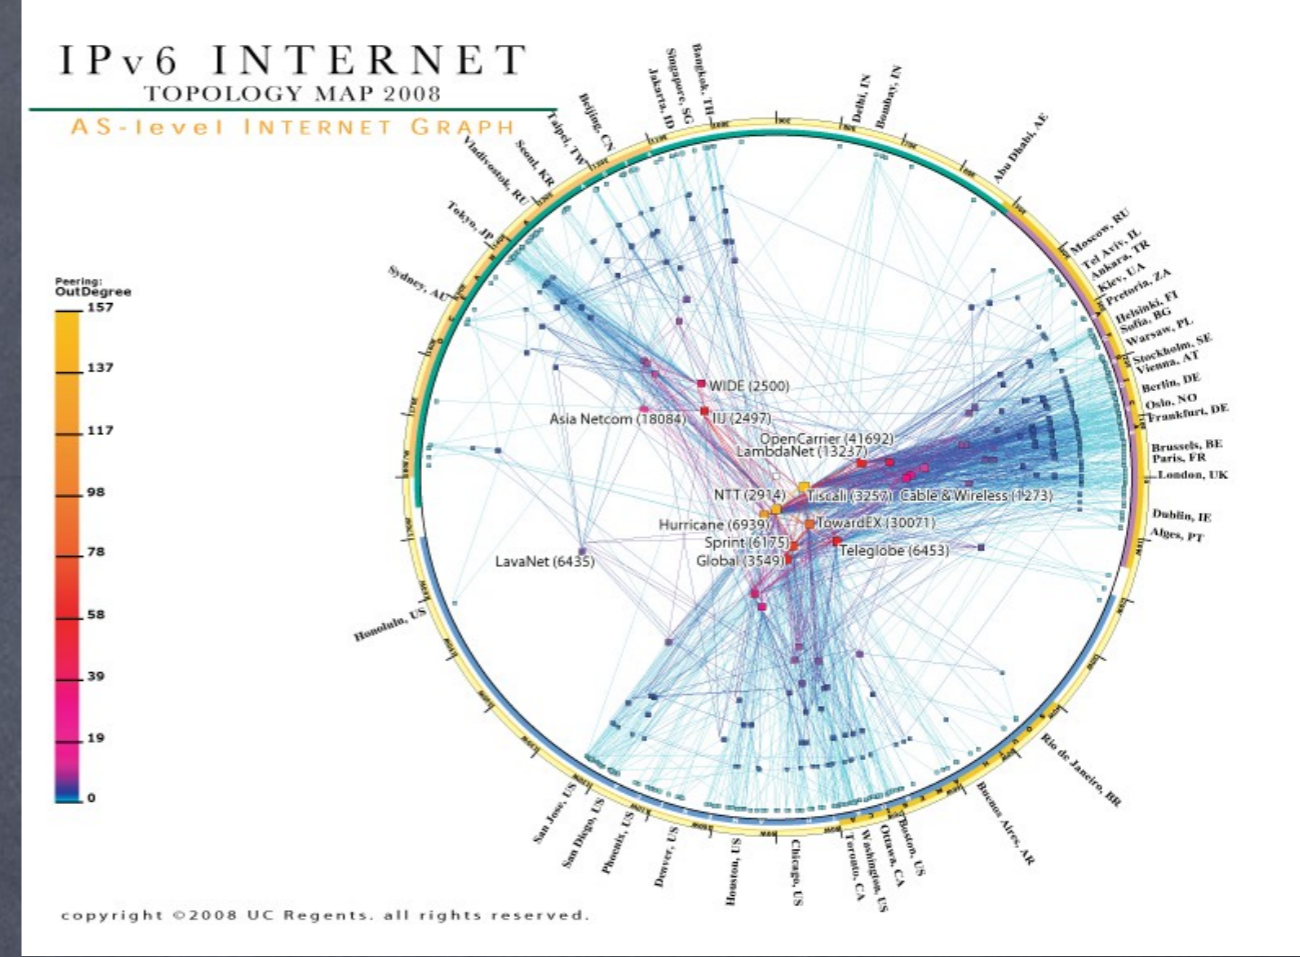



IRNC-SP: Sustainable data-handling and analysis methodologies for the IRNC networks



caida
www.caida.org



Principle Investigator: *kc claffy* <kc@caida.org>

Presenter: *Josh Polterock* <josh@caida.org>

CAIDA
NSF – IRNC Workshop
Arlington, VA
13 July 2010

Overview



To help make operational network data available to the research community, we propose three concrete contributions to the IRNC community's measurement efforts:

- (1) to foster and distill discussion of how to best make IRNC data and statistics available,
- (2) to adapt two CAIDA measurement technologies for IRNC community needs, and
- (3) to experiment with two innovations in data-handling procedures applied to existing IRNC measurements.

CAIDA IRNC-SP Plans



We plan to:

(1) Participate in IRNC series of workshops to discuss measurement priorities and to identify how CAIDA and other researchers can support them.

(a) IRNC Kickoff Meeting (today)

(b) IRNC PI Meeting

(c) 2-day annual meetings (AIMS) dedicated to measurement activities/strategies and how IRNC community can make better use of perfSONAR, metadata, and other data-handling and data-protection technologies

CAIDA IRNC-SP Plans (cont)



(2) Improve two CAIDA technologies we already know could better serve the community.

(a) Upgrade Coralreef to handle IPv6, DNSSEC, read data formats such as netflow.

(b) Install IPv4/v6 capable ark monitors at IRNC locations or downstream customers.

CAIDA IRNC-SP Plans (cont)



(3) Apply two innovations in data-handling procedures to existing IRNC measurement data:

(a) a recently proposed framework for privacy-sensitive data sharing, to apply to data not appropriate for public posting, but explicitly requested through designated channels to use in clearly defined research, and

(b) we propose to illustrate our community building effort with a landmark reporting deliverable: a prototype of a "Bureau of Internet Statistics" report, hopefully inspiring other network infrastructure communities to join in this effort.

CoralReef Improvements

- add IPv6 support to `crl_flow` (an app that counts packets, bytes, and flows, and is used as the back end to the report generator)
- add IPv6 support to `crl_anf` (a faster alternative to `crl_flow` that samples packets)
- improved decoding (printing) of IPv6 headers
- | IP header in an ICMP error message)

Coralreef (cont)



- add IPv6 address anonymization
 - prefix preserving anonymization with generalized Crypto-PAn algorithm
 - or zero-out some or all bits of address
 - option to apply IPv4 anonymization policy to IPv4 addresses embedded within IPv6 addresses (IPv4-mapped, SIIT, Teredo, 6to4, 6over4, ISATAP)
 - option to leave multicast addresses intact
 - anonymizes IP addresses not just in the top IP header, but also in nested headers (e.g., IPIP, or the original IP header in an ICMP error message)

Coralreef (cont)



Planned CoralReef Changes

- IPv6 support in report generator
- DNSSEC support in `crl_flow`, `crl_anf`, report generator?
- additional stats in `crl_flow`, `crl_anf`, report generator?
- netflow import to report generator?

Coralreef (cont)



Passive Network Monitors

Chicago monitor A Chicago monitor B (current) San Jose monitor A San Jose monitor B SDNAP monitor

Quick links:

- [Application bits/second](#)
- [Application packets/second](#)
- [Application flows/second](#)
- [Source country bits/second](#)
- [Destination country bits/second](#)
- [Latest day application tables](#)
- [Latest day application pie charts](#)
- [Latest day source country pie charts](#)
- [Latest day destination country pie charts](#)

Choose a view:

Rows:

Timescales

Columns:

Data sources

Data sources:

- Protocol
- Application
- Source Country
- Destination Country

Graphs:

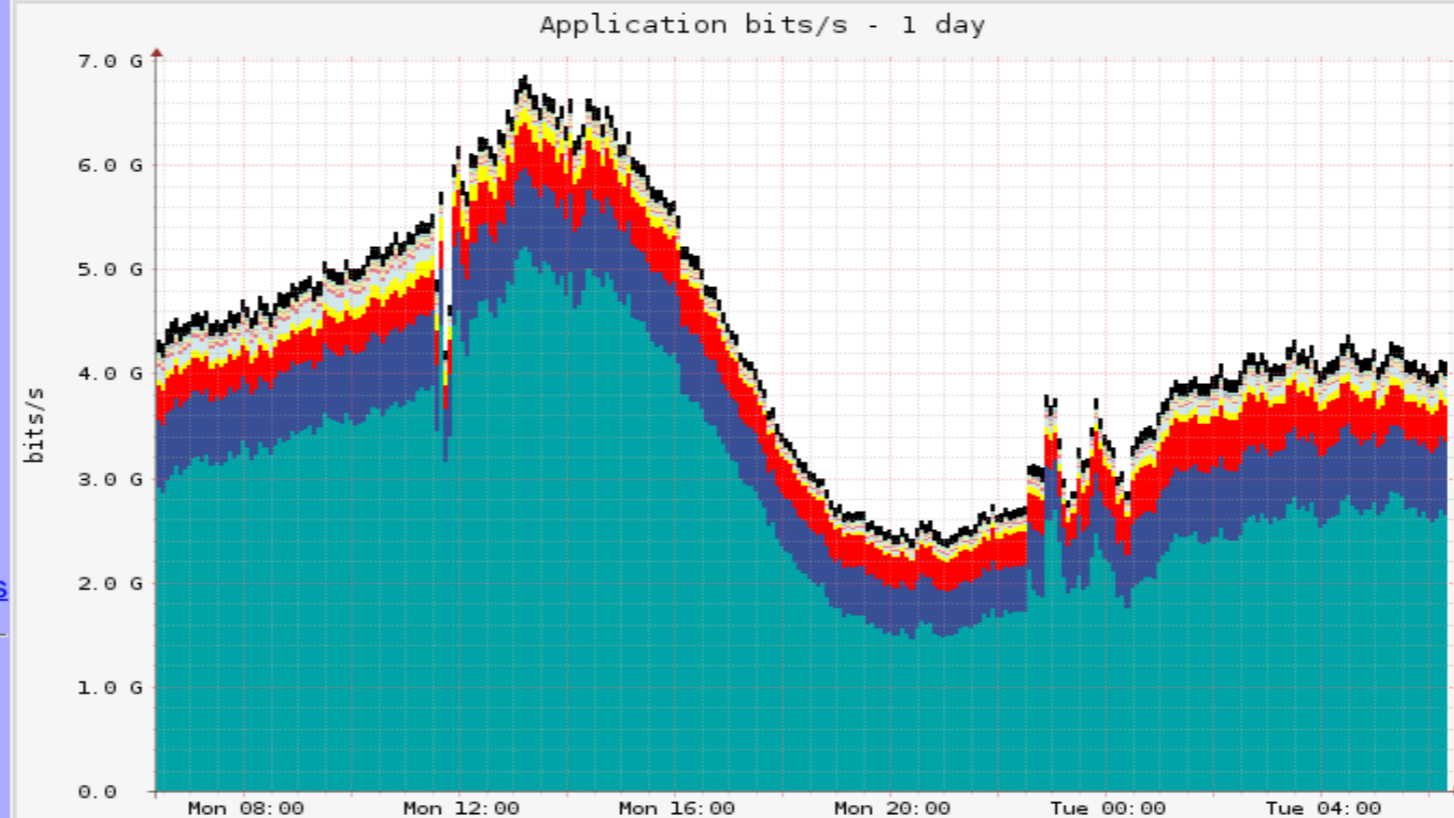
timeseries graphs (abs)

Counters:

bits/sec

Timescales:

1 day



Application	Min	Avg	Max
HTTP	1.45G	2.98G	5.22G
UNKNOWN_TCP	428.72M	616.20M	772.13M
UNKNOWN_UDP	224.37M	373.55M	498.98M
RTMP	24.90M	76.99M	215.15M
HTTPS	20.55M	69.69M	180.07M
ABACAST	5.82M	20.26M	39.58M
RemoConChubo	1.96M	15.96M	48.20M
QUAKE	8.84M	14.05M	22.24M
RTSP	3.55M	13.18M	33.68M
GNUTELLA	5.20M	12.80M	20.93M
SMTP	6.52M	12.33M	23.00M
IPSEC	5.22M	11.91M	26.38M
NNTP	0.00	247.79k	4.61M
other	65.64M	107.48M	195.44M

generated 2010-06-02 05:26 UTC
created with CAIDA's CoralReef (c) 2010 UC Regents

Archipelago (Ark)



- CAIDA's measurement infrastructure
- Built on decade of achievements, from SIGCOMM to MOMA
- Launch 12 Sept 2007
- 45 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



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)

Topology mapping: future work

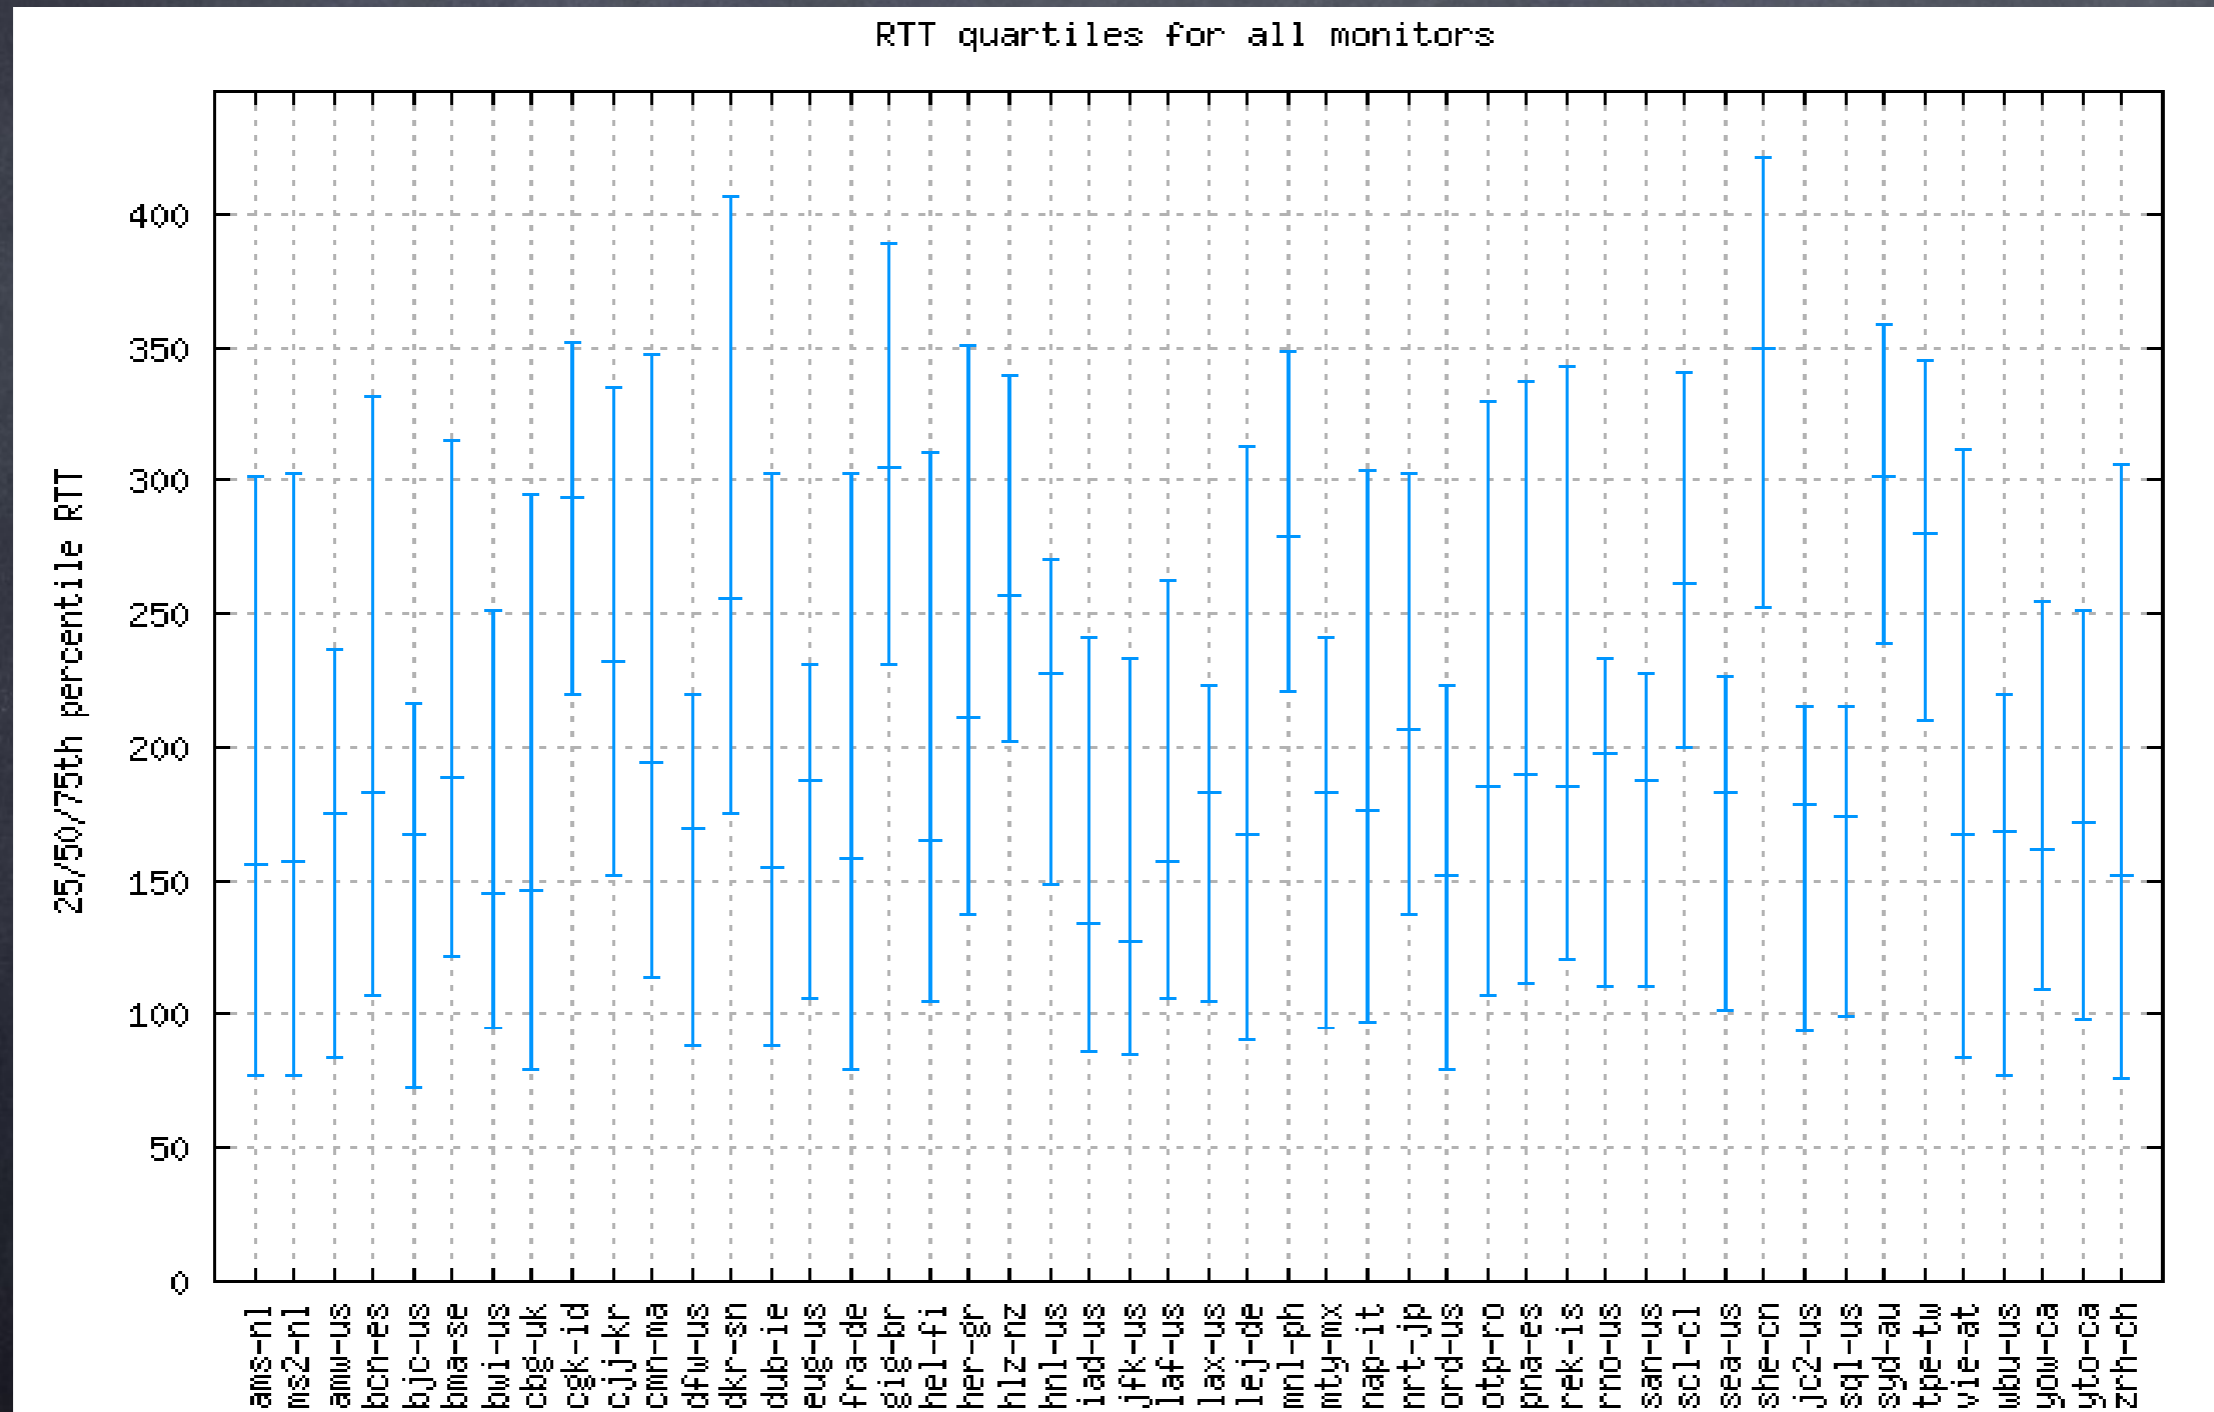


- MIDAR improvements
 - adapt corroboration spacing to responsiveness
- MAARS: Multi-Approach Alias Resolution System
 - combine MIDAR, kapar, iffinder (and others?)
- AS-router Dual graph, including regular updates
- Release supporting tools under GPL
- Support additional collaborators' experiments

Ark Monitor Statistics Pages



- per-monitor analysis of IPv4 topology data






caida
www.caida.org

AS dispersion by AS hop



687k
traces




- 73 WASHINGTON-AS - University of Wash
- 101 WASH-NSF-AS - University of Washing
- 11164 TRANSITRAIL - National LambdaRail, L
- 3356 LEVEL3 Level 3 Communications
- 7018 ATT-INTERNET4 - AT&T WorldNet Serv
- 701 UUNET - MCI Communications Services
- 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 WASHINGTON-RD-AS - University of

sea-us monitor

AS dispersion by IP hop




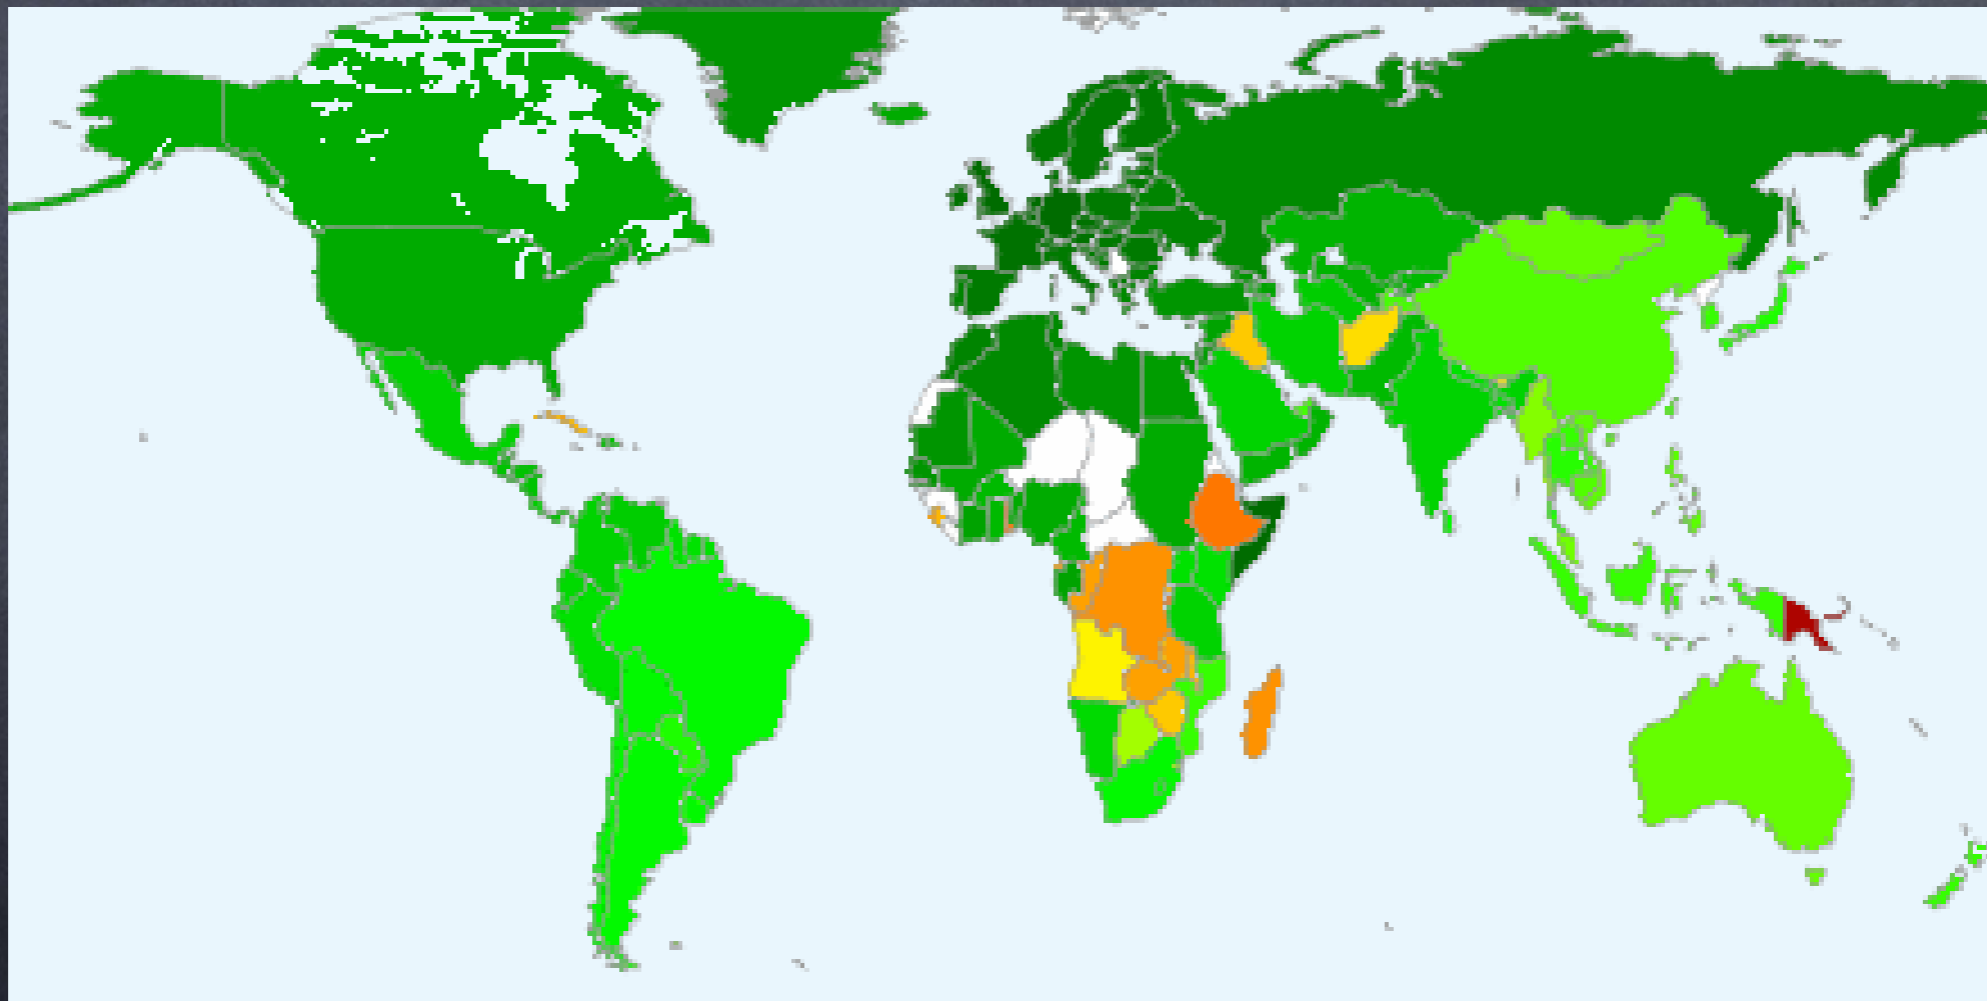
AS dispersion by IP hop: see load balancing


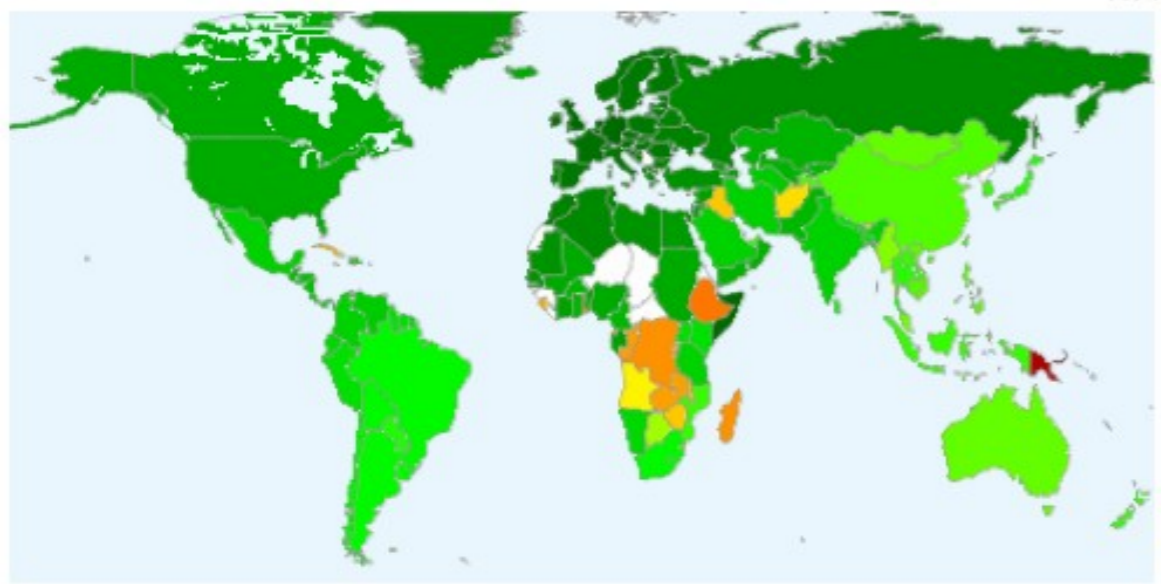




Statistics Pages

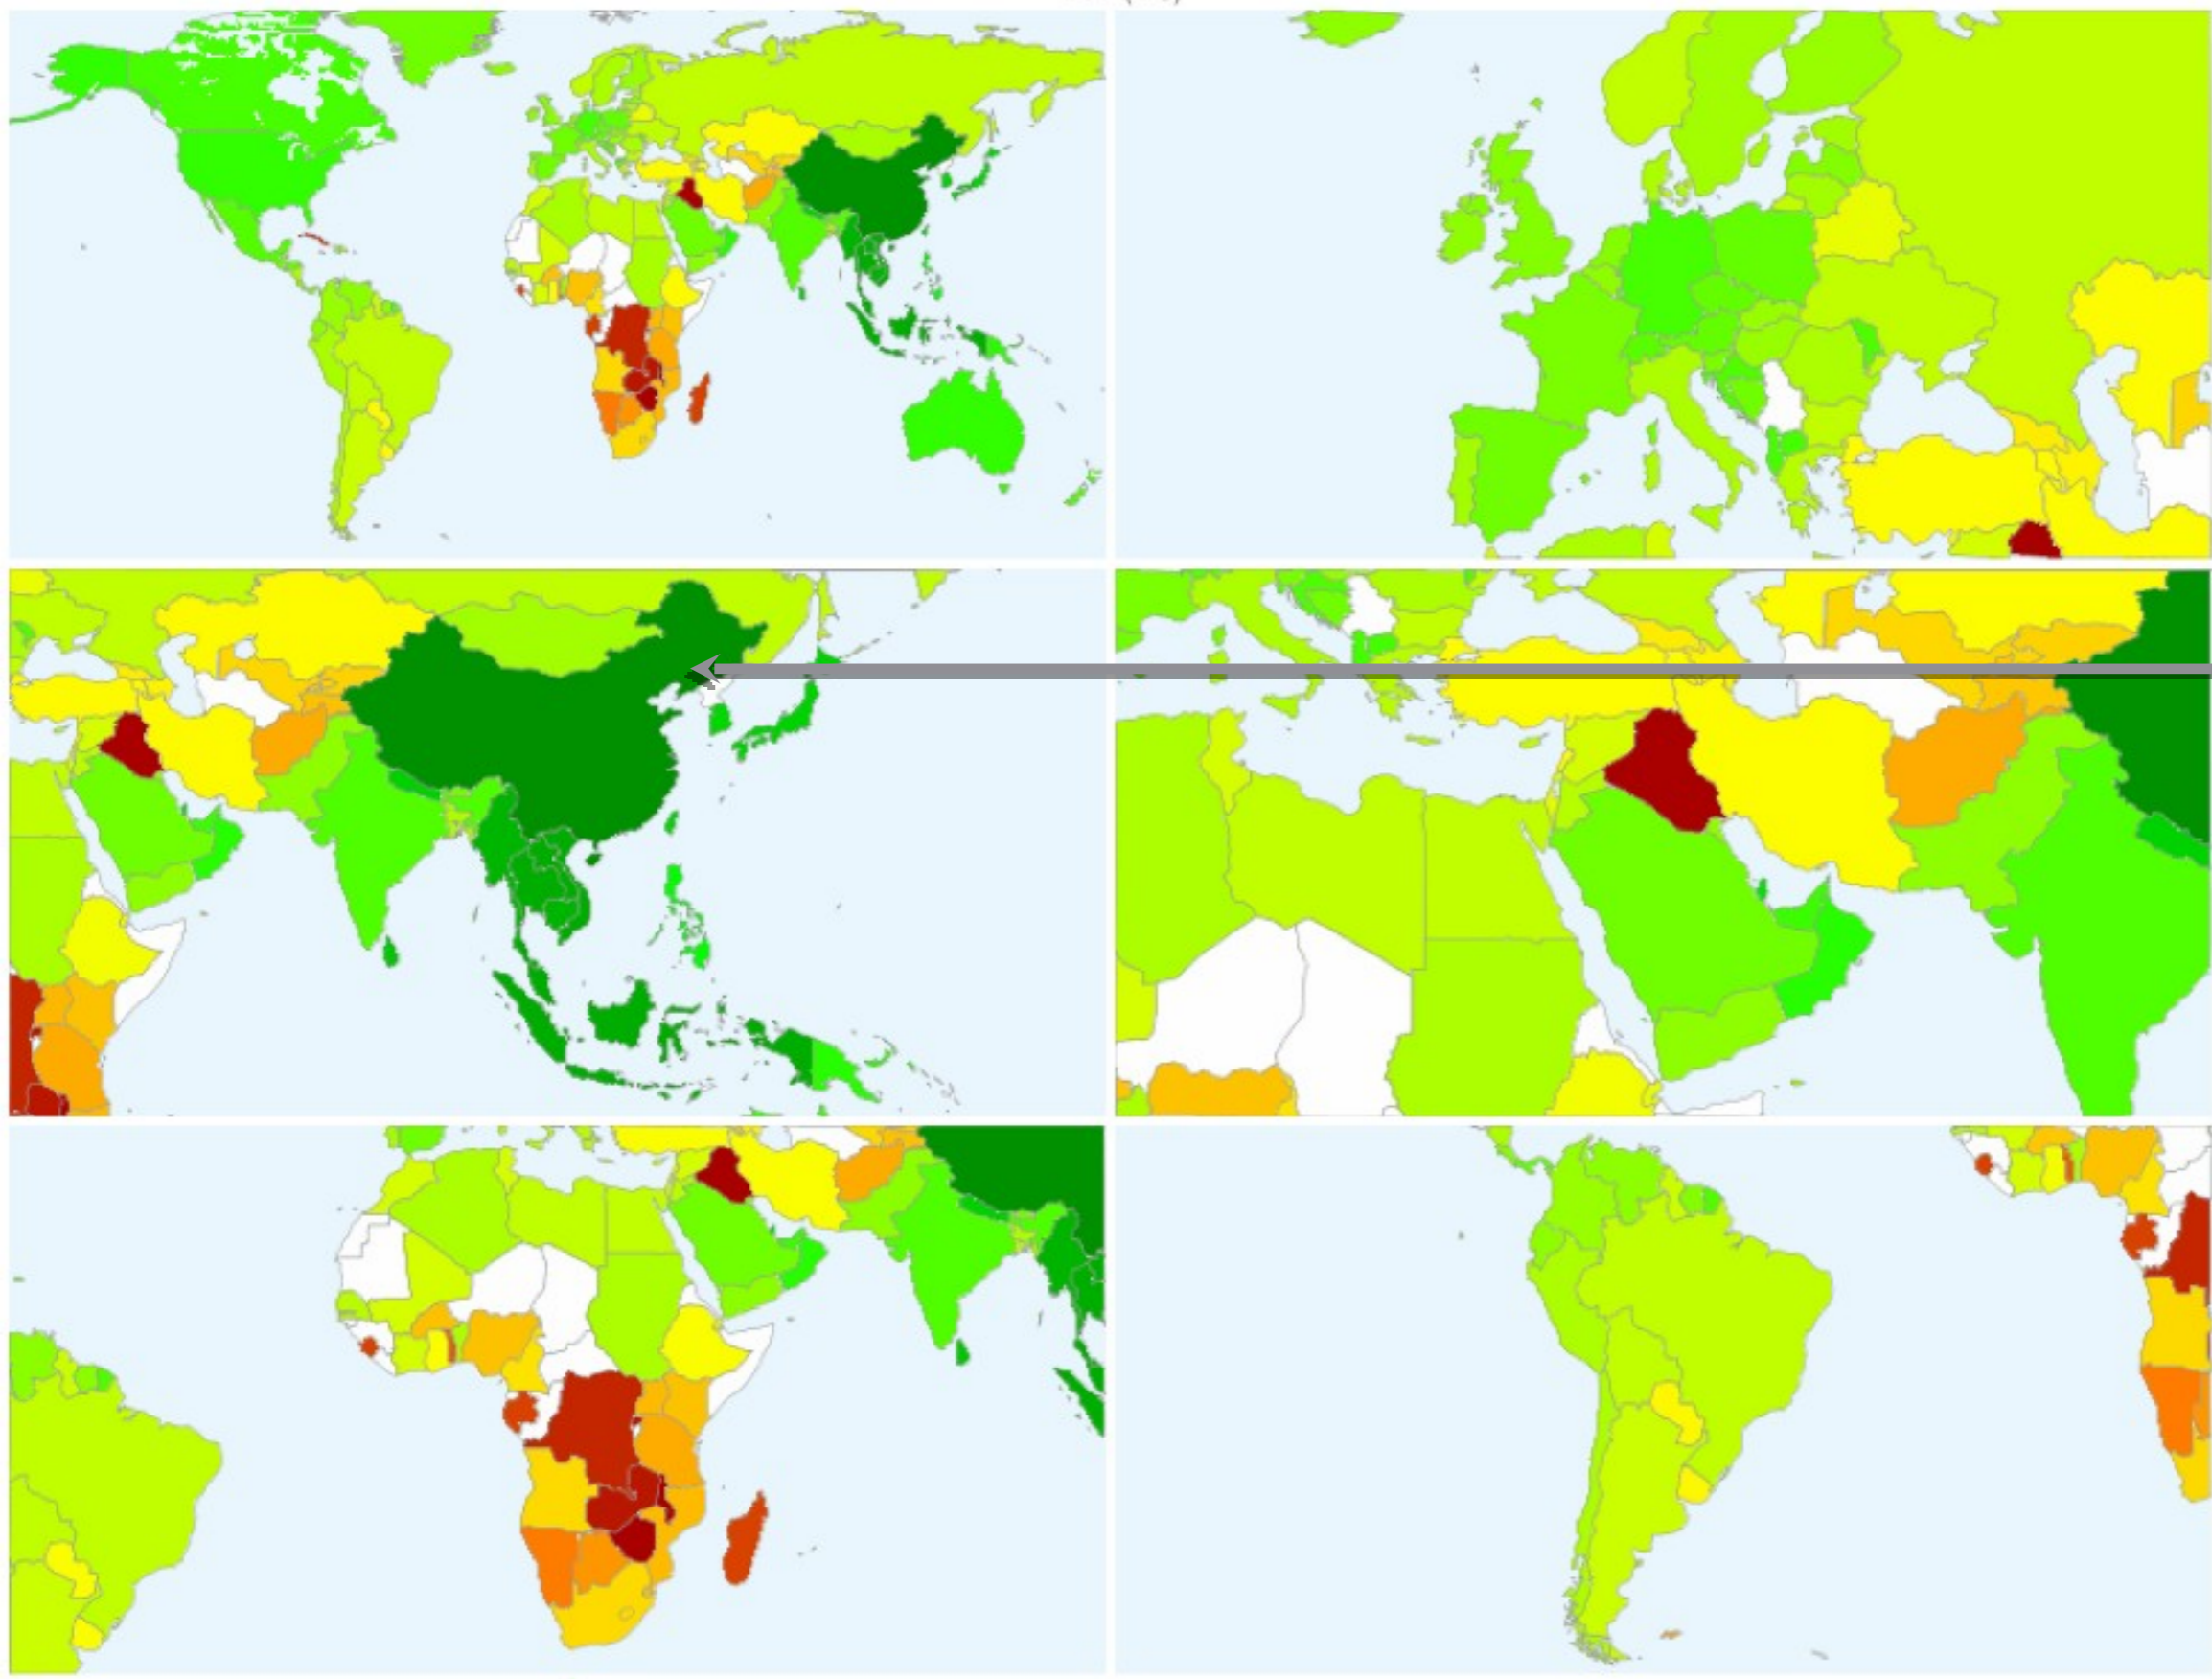



- work in progress: RTT plotted by country
 - geolocate destinations with NetAcuity
 - color each country by median RTT of destinations







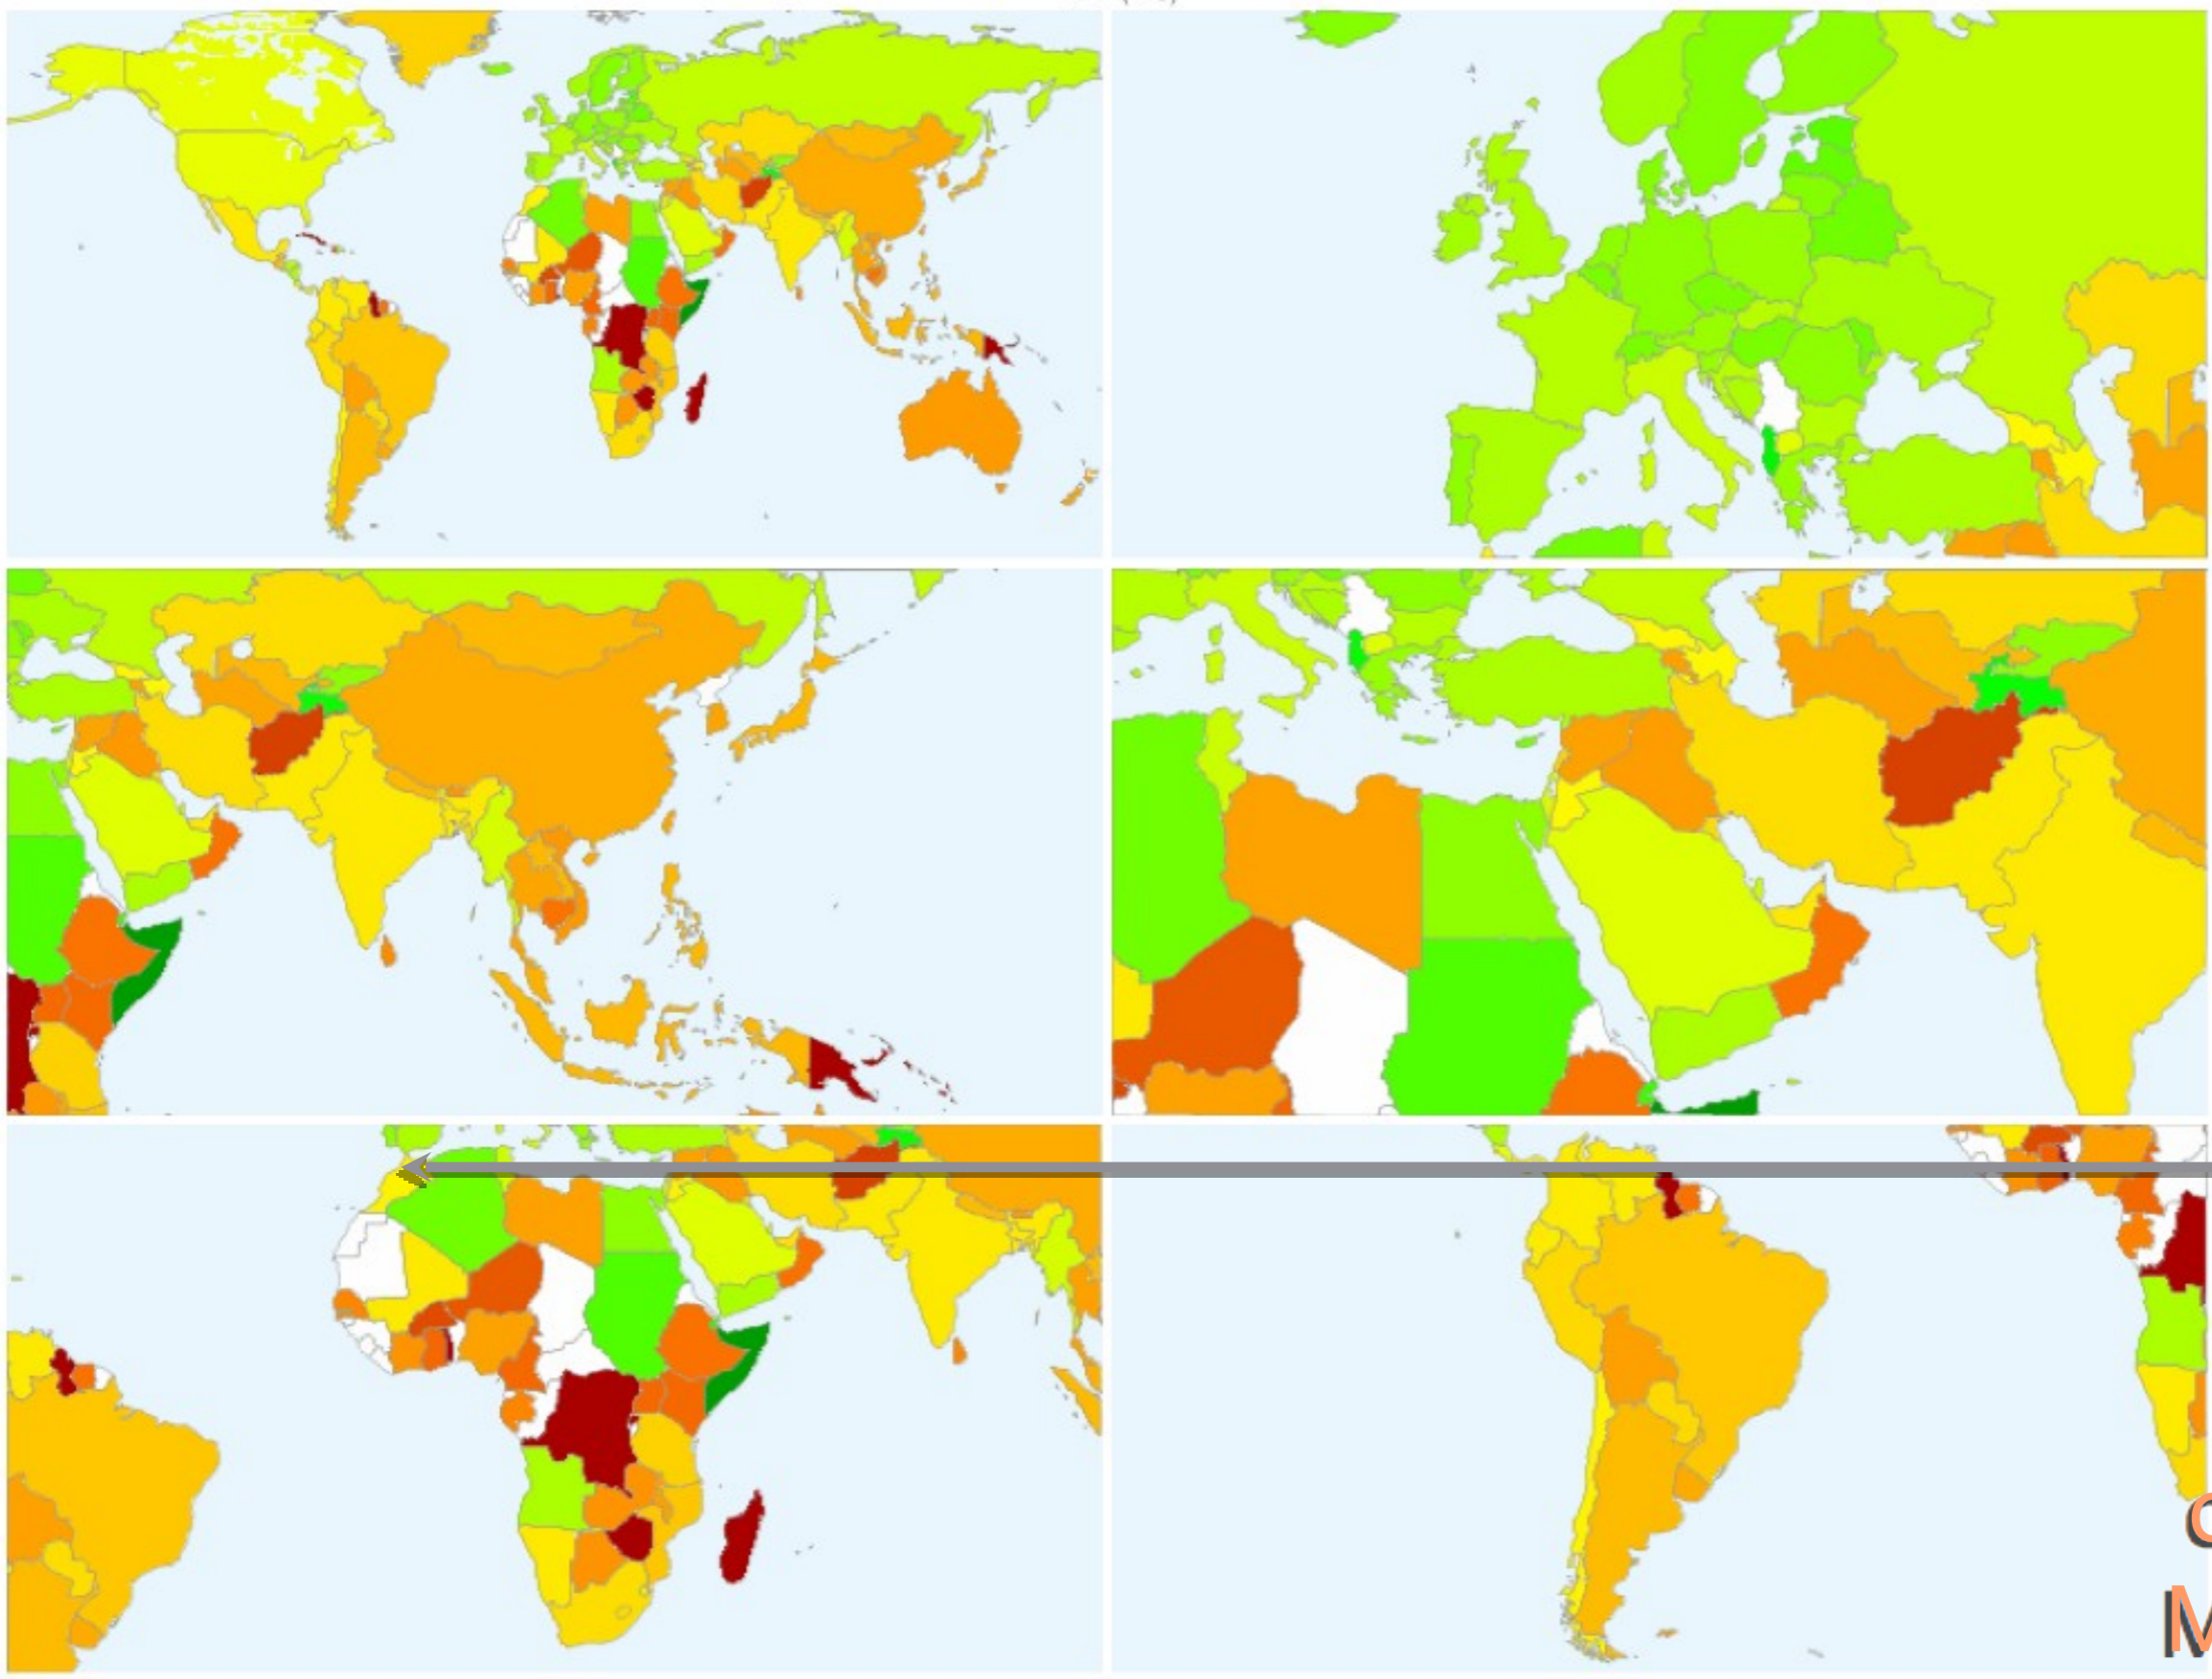
view
from
ams-nl
Netherlands



view
from
she-cn
China




caida
www.caida.org



view
from
cmn-ma
Morocco

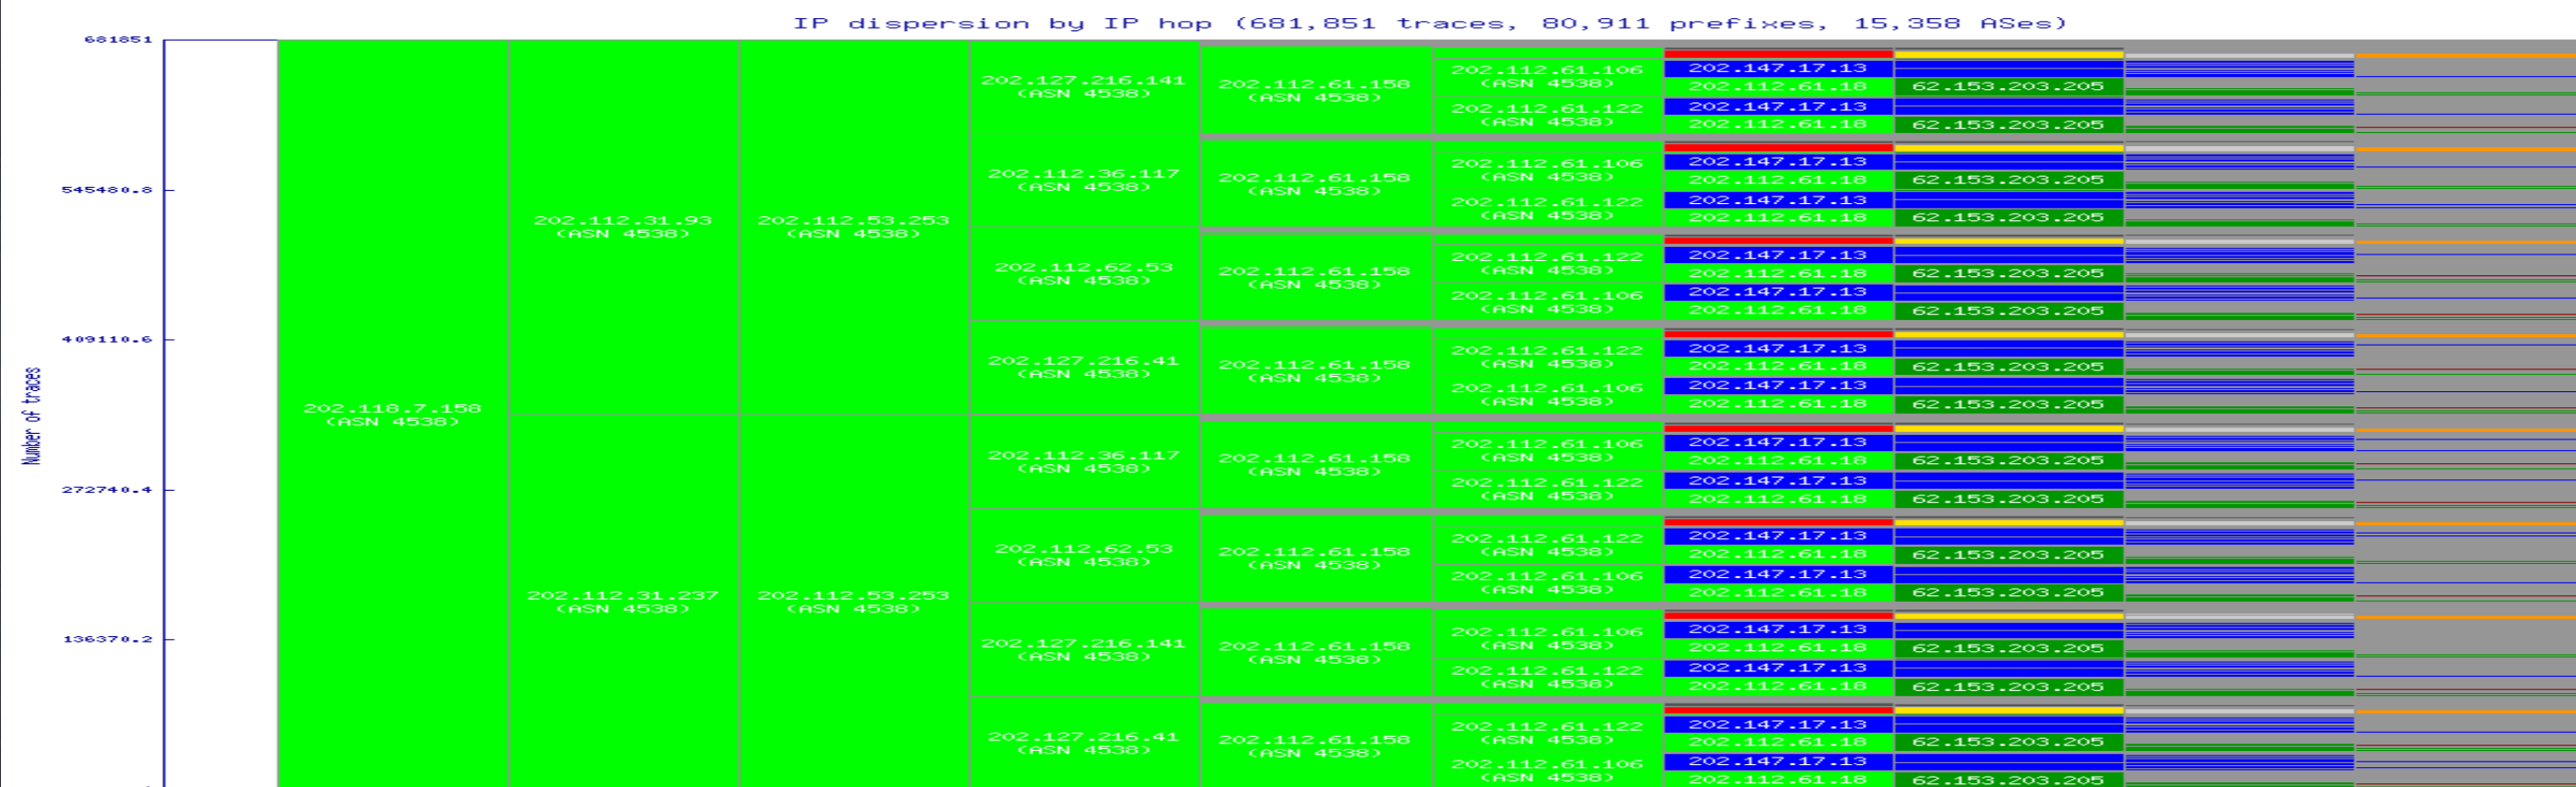
technical accomplishments: views from monitors



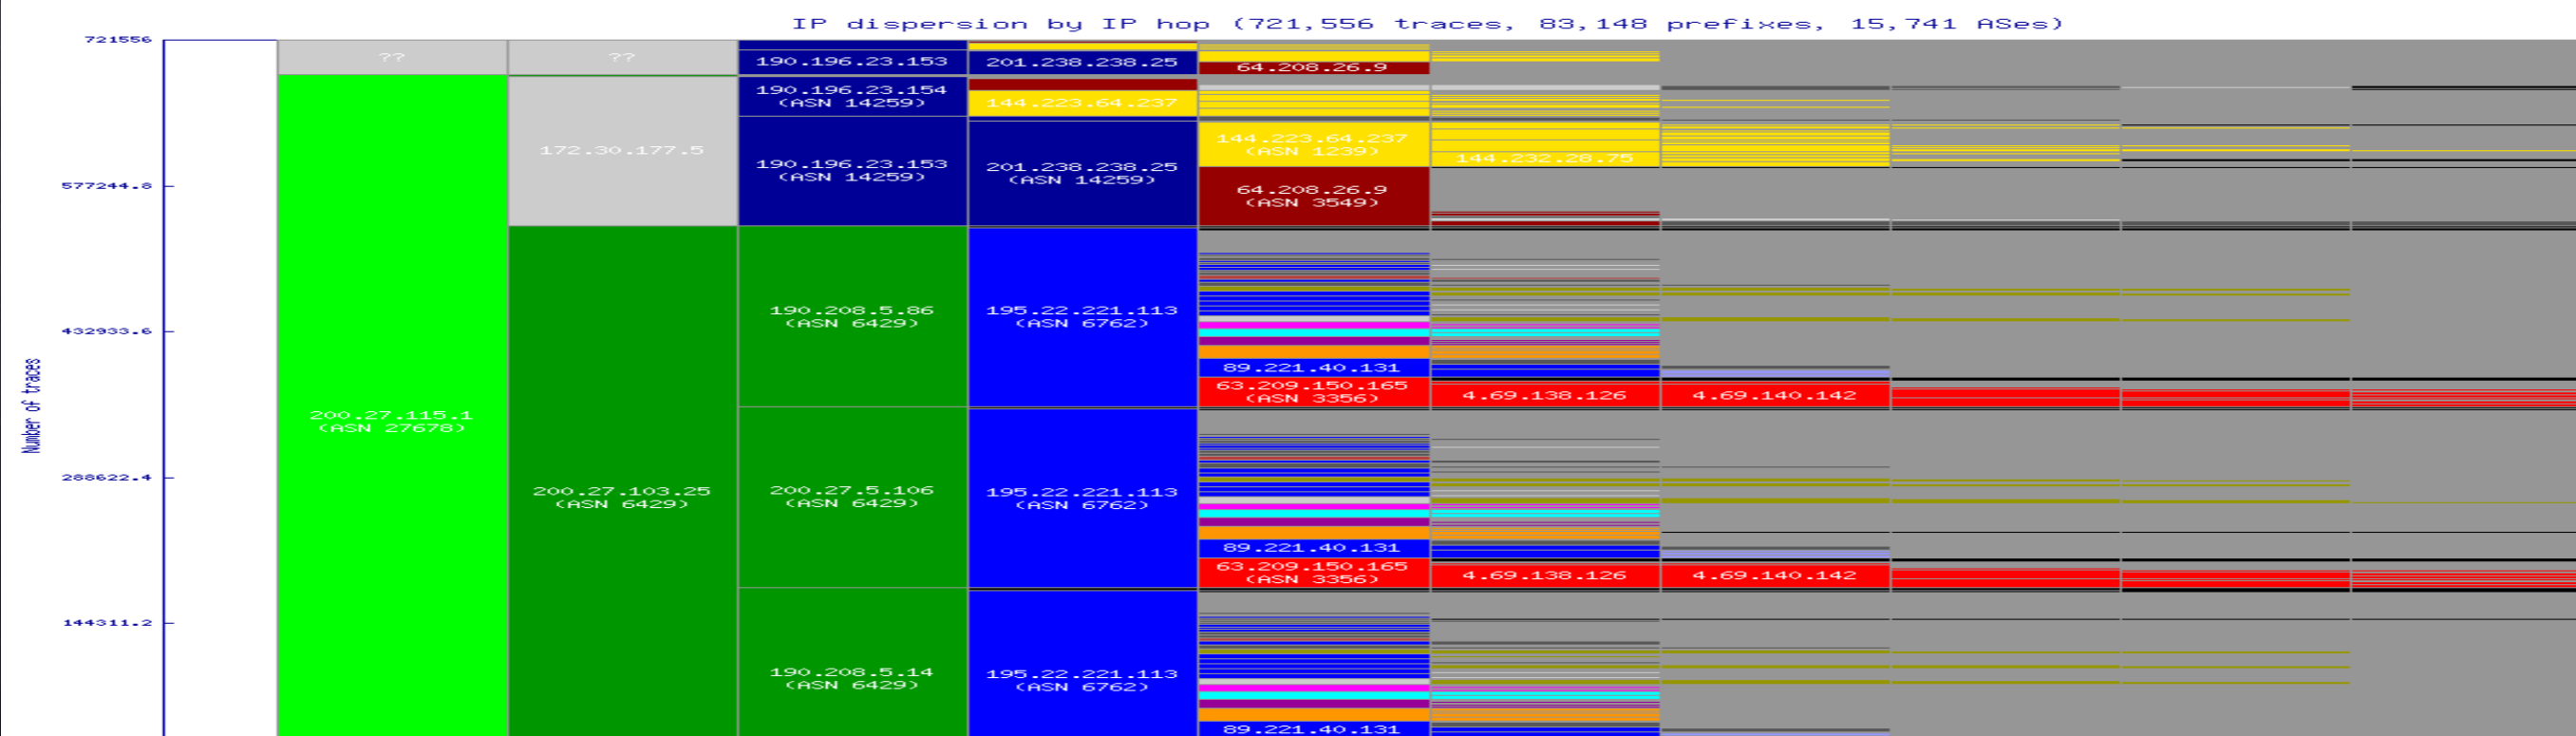
caida
www.caida.org

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



Other Links



- IRNC-SP: Sustainable data-handling and analysis methodologies for the IRNC networks

<http://www.caida.org/funding/irnc/>

- Archipelago (Ark) network measurement platform

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

- Archipelago Monitor Statistics

<http://www.caida.org/projects/ark/statistics/>

- Coralreef

<http://www.caida.org/tools/measurement/coralreef/>