CYBER SECURITY DIVISION – Program Manager – Ann Cox, PhD.

Reverse site visit, Washington, DC

Cartographic Capabilities for Critical Cyberinfrastructure ("C4"): Internet topology and performance analytics for mapping critical network infrastructure

(TTA#7: Network Mapping and Measurement)

CAIDA/UCSD

PI k claffy

13 November 2015

Team Profile

The Center for Applied Internet Data Analysis (CAIDA)

- Founded by PI and Director k claffy
- Independent analysis and research group
- 15+ years experience in data collection, curation, and research
- Renowned world-wide for data collection tools, analysis, and data sharing
- located at the University of California's San Diego Supercomputer Center

Key personnel: Bradley Huffaker, Young Hyun, Marina Fomenkov, Josh Polterock, Ken Keys, Matthew Luckie (now at Waikato), Amogh Dhamdhere, Vasilieos Giotsas

Outline

Team Profile
Project Motivation and Description
Ark Infrastructure and Measurements
Outcomes of Data Curation and Sharing
Topology Analysis Methods and Workflow
    Increasing Completeness
    Increasing Accuracy
    Increasing Annotations
Much supporting software (mostly open source), algorithms,
    interactive interfaces to data, visualizations
Future Work
Quad Chart

http://www.caida.org/funding/c4/

Project Description

The project integrates strategic Internet measurement and data analysis capabilities to provide comprehensive annotated Internet topology maps that improve our ability to identify, monitor, and model critical cyberinfrastructure.

Motivation: there is no map of the net

- The best available data about the global interconnection system that carries most of the world's communications traffic is incomplete and of unknown accuracy.
- There is no map of physical link locations, capacity, utilization, or interconnection arrangements.
- This opacity hinders R&D efforts to: model network behavior and topology; design protocols and new architectures; and assessment of real-world security and stability properties such as hygiene, robustness, resilience, and economic sustainability.

We designed, implemented, deployed, and operated a secure infrastructure named Archipelago (Ark) that supports large-scale active measurement studies of the global Internet.

Motivation (from DHS BAA)

"The protection of cyber infrastructure depends on the ability to identify critical Internet resources, incorporating an understanding of geographic and topological mapping of Internet hosts and routers. A better understanding of connectivity richness among ISPs will help to identify critical infrastructure. Associated data analysis will allow better understanding of peering relationships, and will help identify infrastructure components in greatest need of protection. Improved router level maps (both logical and physical) will enhance Internet monitoring and modeling capabilities to identify threats and predict the cascading impacts of various damage scenarios."

Motivation: mapping capabilities are weak

- Active Measurement Approach:
  - Send tailored probes, elicit specific behaviors, make inferences
  - Vantage points may also serve as destinations
  - Integrate with BGP, traffic, other data for deeper and broader view
- Examples of Internet-wide active measurements:
  - Vulnerability scanning (e.g., heartbleed, blind TCP attacks)
  - Topology mapping (e.g., interconnection of service providers)
  - Routing stability (e.g., comparing IPv4 and IPv6)
  - Network hygiene (e.g., ingress filtering BCP compliance)
  - Infrastructure robustness (e.g., DNS readiness for root key roll)

Research Challenges: why this is hard

(Rob Beverly, NPS)

- Internet (and TCP/IP protocol suite) not designed to be measured
  - Must be reverse-engineered
  - Many tools and techniques are “Tricks and Hacks”
    - Network may “lie” in responses
  - Service providers don’t want to be measured (competitive, economic reasons)
  - Best common security practices often prevent measurement
- Millions or billions of measurements often required
- Dependence on location and quantity of vantage points
- Lots of large data (packets, flows, routing messages, topology, etc)
- → Needle in haystack (data mining)

Monitor Deployment

Legend: - Raspberry Pi - FreeBSD

- 135 monitors in 44 countries
  - 87 Raspberry Pi's
  - 56 have IPv6
  - 35 have RADclock

Continent

39	North America
6	South America
39	Europe
5	Africa
14	Asia
4	Oceania

Organizations

48	academic
24	residential
23	commercial/business
10	network infrastructure
2	other

Raspberry Pi

1st gen

- 700MHz ARMv6
- 512MB RAM

2nd gen

- 900MHz quad-core ARMv7
- 1GB RAM

both

- 100 Mbps Ethernet
- 8GB SD card
- $35 for bare board

~$68 complete system

Outcomes of Ark: Datasets

- **Publicly Available Datasets**
  - –The Ark IPv4 Routed /24 Topology Dataset (data older than 2 years) PREDICT
  - –The Ark IPv4 Routed /24 DNS Names Dataset (data older than 2yrs) PREDICT
  - –Ark Internet Topology Data Kits (ITDK) (data older than two years) PREDICT
  - –The Ark IPv6 Topology Dataset PREDICT
  - –The Ark IPv6 DNS Names Dataset PREDICT
  - –IPv4 Routed /24 AS Links (September 2007 - ongoing)
  - –IPv6 AS Links (December 2008 - ongoing)
  - –AS Relationships
  - –AS Classification

Outcomes of Ark: Datasets

- **Restricted Access Datasets**
  - The Ark IPv4 Routed /24 Topology Dataset (most recent two years) PREDICT
  - The Ark IPv4 Routed /24 DNS Names Dataset (most recent two years) PREDICT
  - Ark Internet Topology Data Kits (ITDK) (most recent two years) PREDICT
  - PAM 2010 "Improving AS Annotations" Supplement

- **Web-based Datasets**
  - Ark statistics http://www.caida.org/projects/ark/statistics/
  - AS Rank http://as-rank.caida.org/ Interactive
  - DNS Decoding Database (DDec) http://ddec.caida.org/ Interactive

Outcomes of Ark: Supported Projects

- **Hosted Measurements**
  - The Spoofer Project (w/ Robert Beverly @ NPS)
  - TCP Behavior Inference (w/ Matthew @ Waikato)
  - IPv4 and IPv6 stability (w/ Ioana @ Simula)
  - DNS Health (Casey Deccio at Verisign)
  - TCP HICCUPS (w/Robert @ NPS)
  - others at http://www.caida.org/projects/ark/
- **Ongoing Measurements**
  - IPv4 and IPv6 topology discovery
  - Congestion
    - Time-Sequence Ping (TSP)
    - ISP border mapping

Outcomes of Ark: Supported Projects

- **On-demand Measurements**
  - tod-client (topology on-demand)
  - Vela.caida.org: Web Interface
- **Developing Experimental Capabilities**
  - Outage detection (e.g. cable cuts, natural disasters)
  - BGP hijacks
  - <*your experiment here*>

Approach

- **Active measurement using Archipelago measurement infrastructure**
  - Continuous random probing of IPv4 address space at /24 granularity
  - 135 monitors and growing (53 IPv6, 86 Raspberry PIs, 35 RadClock) — *ask us if you want one*
- **Alias resolution measurements**
  - Improved tools and techniques
  - IPv4 and IPv6 (different methods)
- **Validate and annotate topology graphs**
  - BGP, traceroute, IXP route servers
  - WHOIS, geolocation
  - Economic data: relationships, type, legal ownership
  - Performance

MIDAR: IPv4-scale Alias Resolution

MIDAR: Monotonic ID-Based Alias Resolution

- **Monotonic Bounds Test**: for two addresses to be aliases, their combined IP-ID time series must be monotonic
- 4 probing methods: TCP, UDP, ICMP, "indirect" (traceroute-like TTL expired)
- sliding-window probe scheduling for scalability
- multiple sources

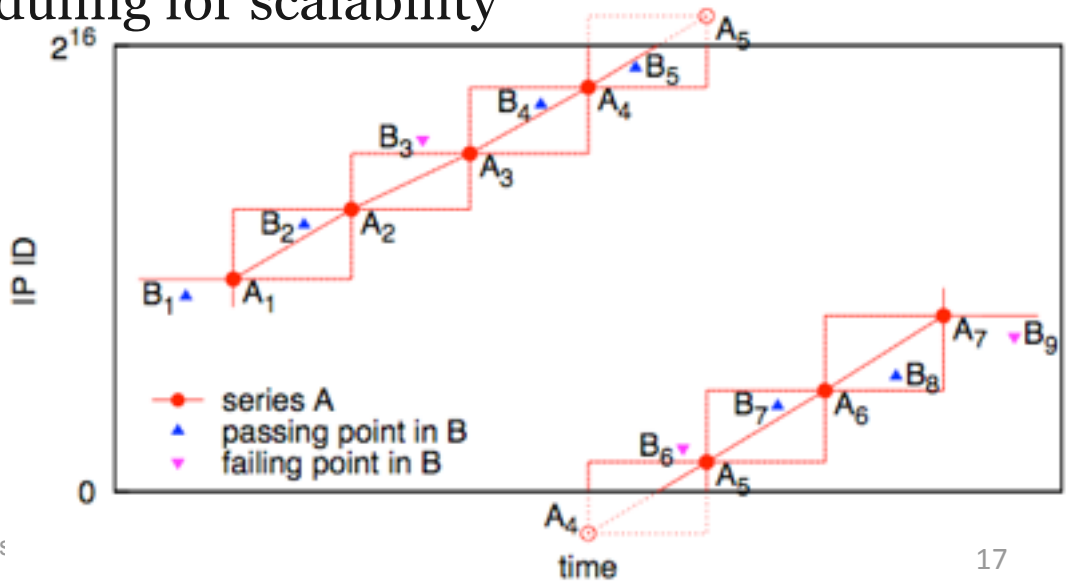

Speedtrap: IPv6 Alias Resolution

- **Speedtrap provides Internet-Scale IPv6 alias resolution**
  - uses IP-ID (slow) velocity and and Monotonic Bounds Test (MBT)
  - IPv6 header does not include ID field, but fragment header does..
  - so, trigger fragmentation (which occurs only at source)
  - challenges (besides slow velocity and large packets)
    - only 32% of interfaces send frags w incrementing IP-IDs
    - 18% send random IPIDs
    - 30% do not respond to ping
  - validates extremely well (>99%) for 2% of routers we could validate

"Speedtrap: Internet-Scale IPv6 Alias Resolution" in IMC Oct 2013
http://www.caida.org/publications/papers/2013/speedtrap/

Approach

- **Active measurement using Archipelago measurement infrastructure**
  - Continuous random probing of IPv4 address space at /24 granularity
  - 135 monitors and growing (53 IPv6, 86 Raspberry PIs, 35 RadClock) — *ask us if you want one*
- **Alias resolution measurements**
  - Improved tools and techniques
  - IPv4 and IPv6 (different methods)
- **Validate and Annotate topology graphs**
  - BGP, traceroute, IXP route servers
  - WHOIS, geolocation
  - Economic data: relationships, type, legal ownership
  - Performance

Dolphin bulk DNS resolution system

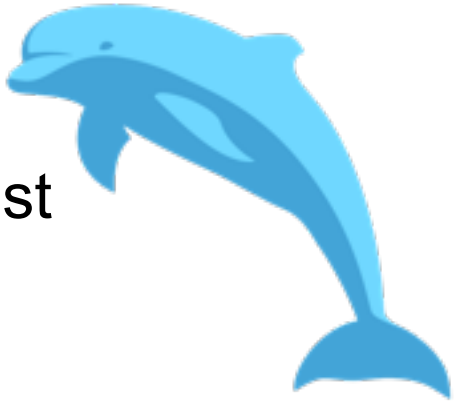

- **conducts parallel DNS lookups**
  - millions of lookups per day from a single host
  - PTR records for IPv4 and IPv6 addresses
- **retries failed lookups automatically**
  - retries once per day for up to 3 days
- **ensures targets only looked up once in any 7 days**
  - remembers all targets queried in most recent 7 days
  - reduces load on authoritative DNS servers, independent of TTL

DRoP (DNS-based Router Positioning)

automated hostname geohints (geolocation) inference

- **automated detection of geographic hint in router hostnames**
  - previous efforts have been manual
- **methodology**
  - Find possible geographic hints
  - Create validation vector from RTT and TTL measurements
  - Train classifier with known domains' hint
  - Build domain specific rules from "likely" hints

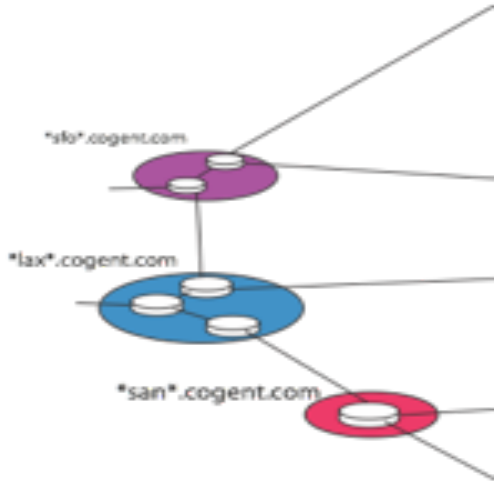

<IATA>\d+.cogent.com
<city name>-gw.routers.es.aau.dk

<CLLI>.([a-z]+\.){2}.ntt.net
<IATA>.above.net

DDec (DNS Decoded)

public database of hostname heuristics

- **central repository of hostname heuristics**
  –DRoP, undns, sarang
- **hostname simplified RegEx**
  <router=D+>.<iata3><pop=D+>.<country2>.opencarrier.eu
- **web interface for operator feedback**
  http://ddec.caida.org
- **community resource for validated hostname decoding**

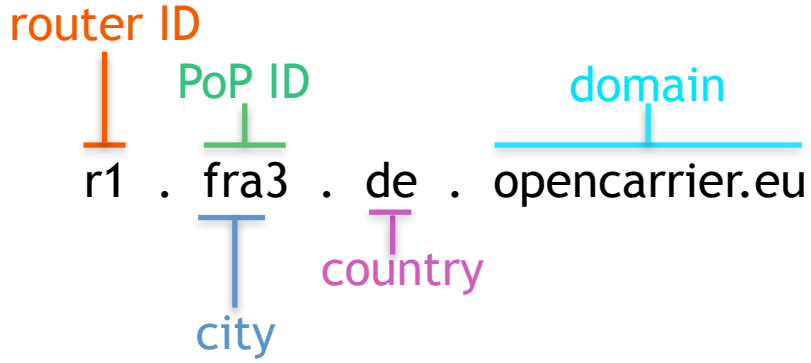

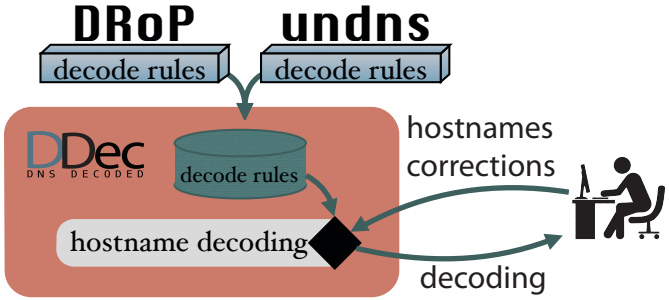

Result: ITDK Datasets

- **Prior releases**
  - **Newly released ITDK-2015-08**
  - ITDK-2014-04 ITDK-2014-12
  - ITDK-2013-04 ITDK-2013-07
  - ITDK-2012-07
  - ITDK-2011-04 ITDK-2011-10
  - ITDK-2010-01 ITDK-2010-04 ITDK-2010-07
  - historical ITDK releases 0204 and 0304 from April 2002 and 2003 collected with skitter (use with caution)
- **All in PREDICT**

AS Relationship Inference Algorithm

- Built new AS relationship inference algorithm and new customer cone inference algorithm
- Performed unprecedented validation
  - 99.6% p2c, 98.7% p2p, 34.7% of 126,082 inferences
- Released code and 97% of validation data to promote reproducibility

""AS Relationships, Customer Cones, and Validation" in IMC Oct 2013
http://www.caida.org/publications/papers/2013/asrank/

AS annotations: companies, types

- "AS2org": mapping each AS to its holding legal entity
  - Infer which ASes belong to the same organization
  - Processed as part of WHOIS data collection machine learning used to infer the business model for individual ASes
- AS type
  - trained against self-identified ASes on PeerDB
  - dataset used for inference:

    AS relationships, customer cone sizes,

    size of advertised address space,

    /24s seen in darknet, websites in Alexa

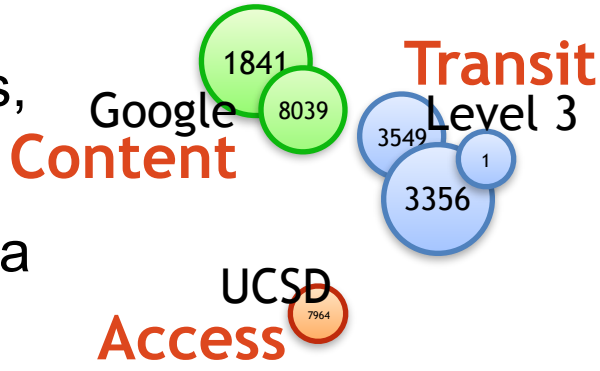

Increased Completeness, Accuracy and Richness of Annotations

Interface to allow operators to submit corrections to our relationship inferences for neighboring ASes.

neighbor				inferred relationship type	actual relationship type
AS rank	AS	AS name	Org name		
5	1299	TELIANET	TeliaNet Global Network	↑ provider	
46	11164	INTERNET2-TRANSITRAIL-CPS	National LambdaRail, LLC	↑ provider	
9	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	↔ peer	
13	6939	HURRICANE	Hurricane Electric, Inc.	↔ peer	
15	3491	BTN-ASN	Beyond The Network America, Inc.	↔ peer	

(correct)
↓ customer
↑ provider
↔ peer
⇔ sibling
(remove entry)

Increased Completeness, Accuracy and Richness of Annotations

Information about an individual AS

AS number:	7922
AS name:	COMCAST-7922
Org name:	Comcast Cable Communications, Inc.
AS rank:	27
Country:	US
Customer cone size:	1,607
AS transit degree:	136 (2 Provider, 82 Peer, 106 Customer, 23 Sibling)
Type:	Tr/Ac

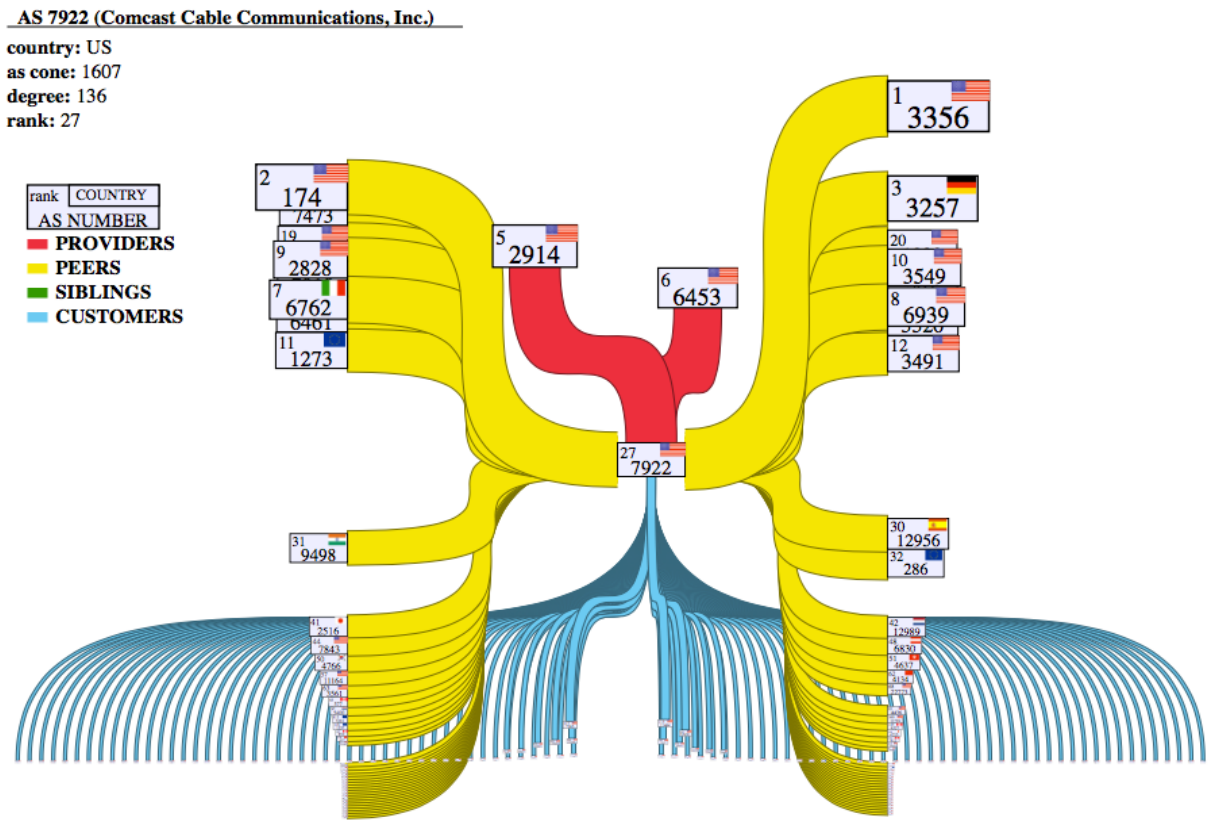

AS relationship graph for Comcast AS 7922

Increased Completeness, Accuracy and Richness of Annotations

PoP-level, router-level and geographic maps for each AS

AS number:	174
AS name:	COGENT-174
Org name:	Cogent/PSI
AS rank:	4
Country:	US
Customer cone size:	13,594
AS transit degree:	3,855

0	1	65	3,789
Provider	Sibling	Peer	Customer

Router-level map

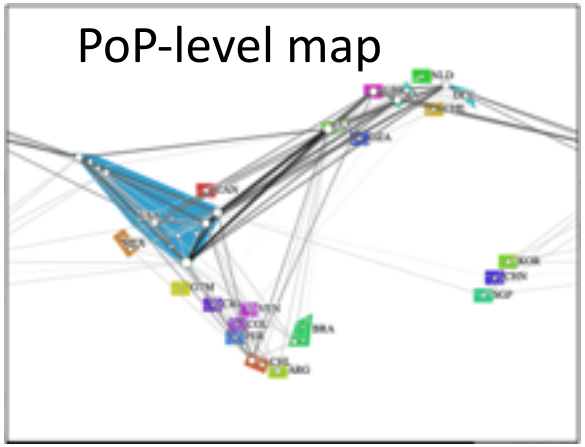

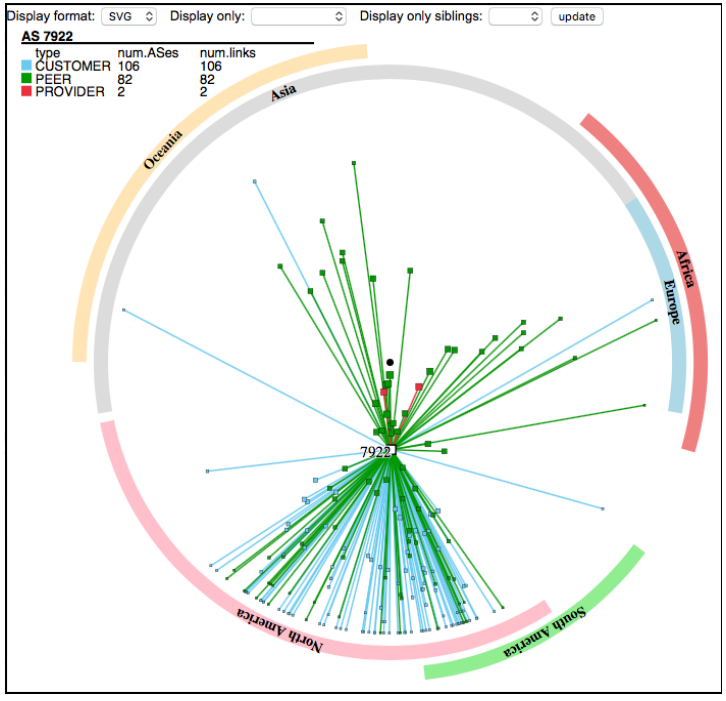

Geographic location and customer cone

Router-to-AS assignment

- Determine the Autonomous System (AS) that owns each router
- Based on router interfaces (known and inferred) assign router to AS via heuristics:
  - Election, Customer, Degree, Neighbor
  - Election+Degree (best combination)

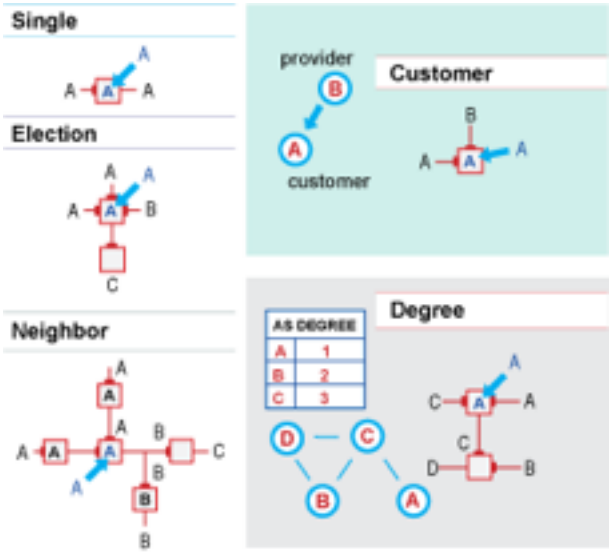

"Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers" in PAM Apr 2010
http://www.caida.org/publications/papers/2010/as_assignment/

Internet Topology Datasets Process

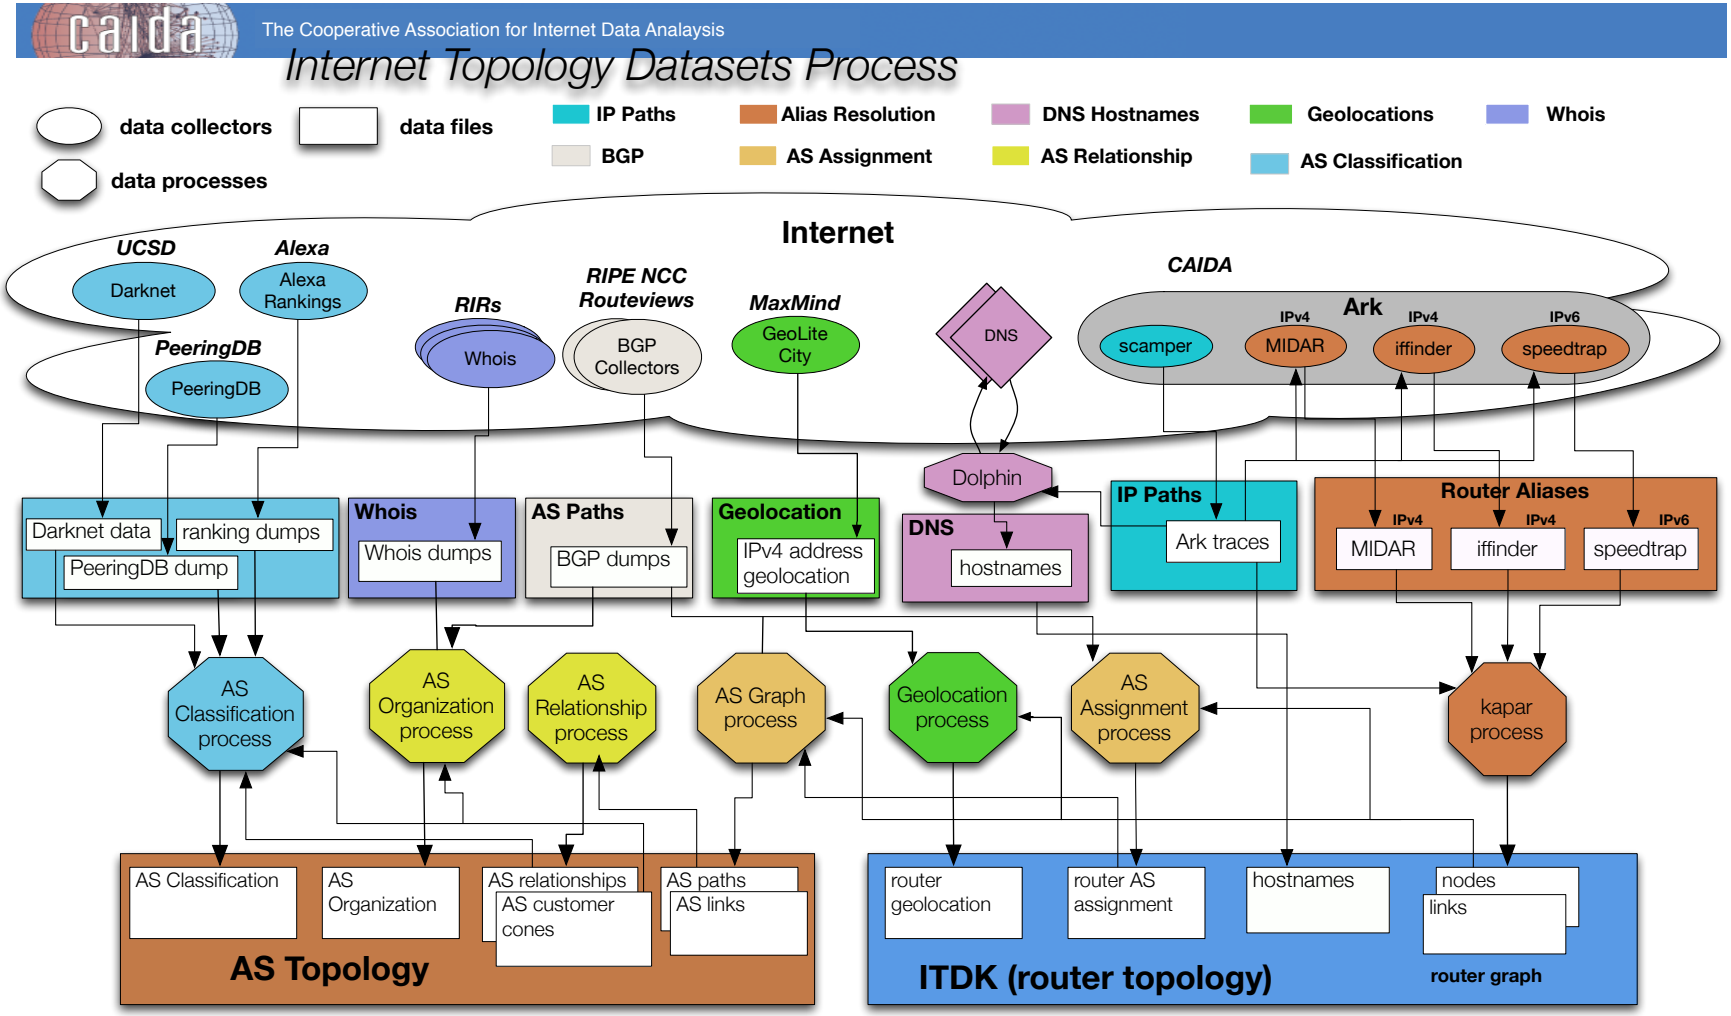

AS Relationships: Process

caida The Cooperative Association for Internet Data Analysis

AS Topology Datasets Process

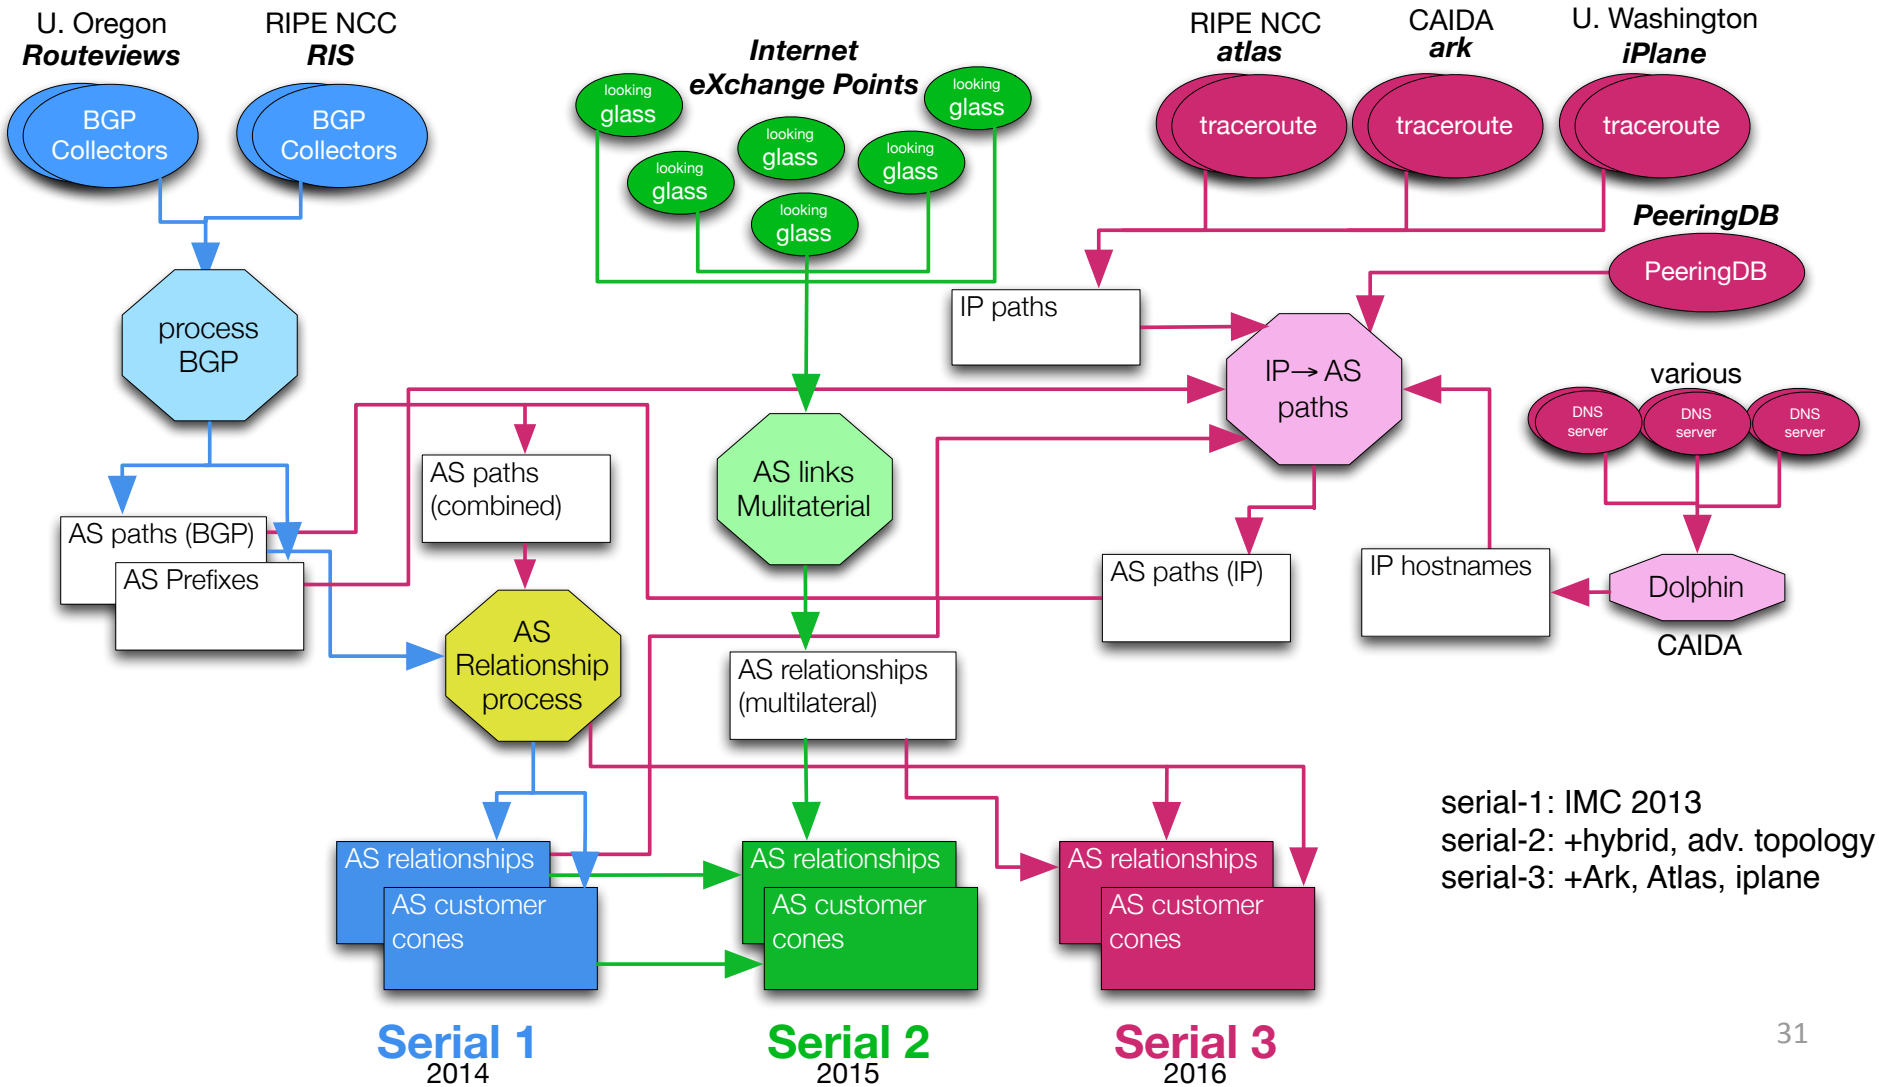

serial-1: IMC 2013
serial-2: +hybrid, adv. topology
serial-3: +Ark, Atlas, iplane

AS Relationships: Serial-1 (2014)

- **AS links inferred from BGP collectors**

  AS Relationships, Customer Cones, and Validation, IMC 2013

  1. Extract all AS links from RouteViews snapshots.
  2. Infer customer-provider relationships, and annotate AS links.
  3. Infer peer-to-peer relationships, and annotate AS links, possibly overriding customer-provider relationships inferred in step 2.
  4. Heuristically fix suspicious looking inferred relationships (e.g., a low-degree AS acting as provider to a high-degree AS).

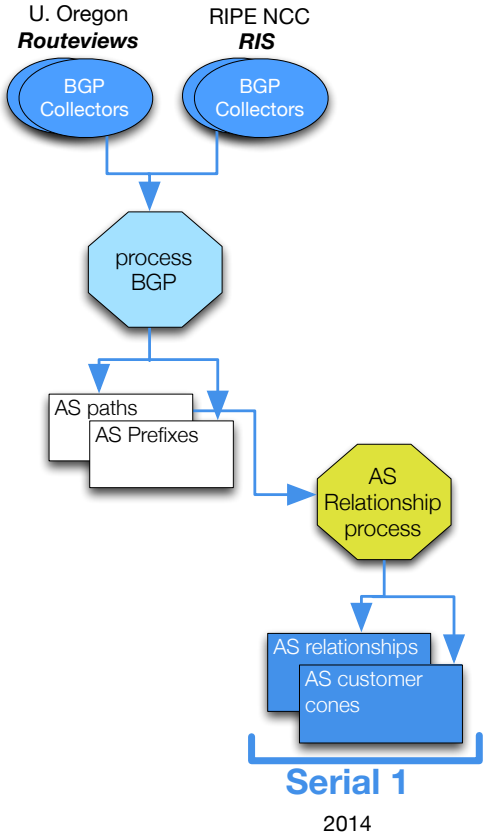

AS Relationships: Serial-2 (2015)

- **Serial-1, plus**
- **AS links inferred from BGP communities collected from IX looking glass servers**

  [Inferring Multilateral Peering, CoNEXT 2013]
- Inferring Multilateral Peering
  1. Collect BGP communities from IX looking glass servers.
  2. Infer peering links between pairs of AS which accept routes from each other at.
  3. Infer peering links at points in the observed AS paths that cross an known IX.
  4. Merge all newly inferred links to the serial-1 graph as peering links

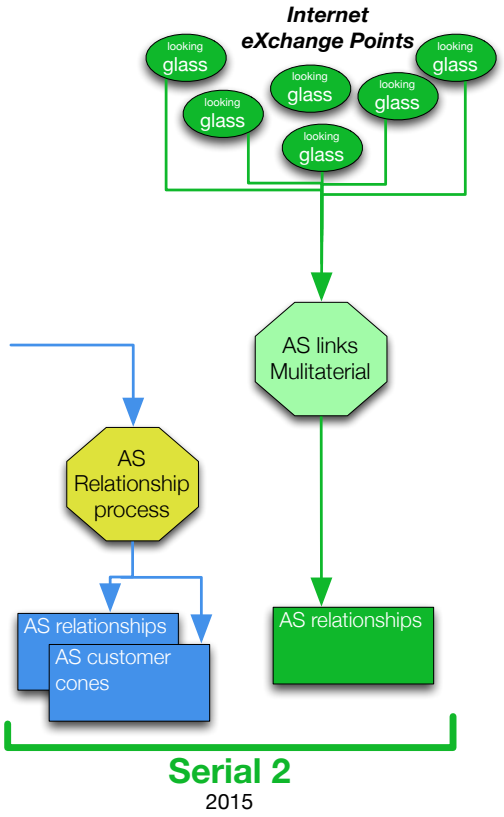

AS Relationships: Serial-3 (2016)

- **AS links inferred from BGP communities collected from IX looking glass servers**
  1. Complete Serial-1 process to generate AS relationships
  2. Collection traceroutes from RIPE NCC's **atlas**, CAIDA's **ark**, U. Washington's **iPlane**
  3. Covert IP paths to AS paths
  4. Merge AS paths from BGP and traceroute
  5. Run Serial-1 algorithm on combined AS paths

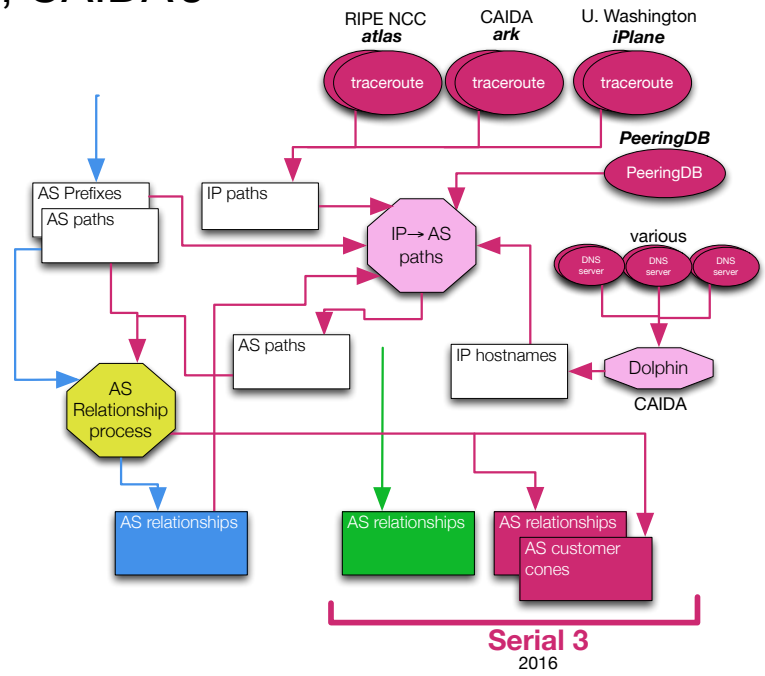

Current Status

- **Remaining Deliverables**
  - Monthly data collection (ongoing)
  - Evaluate traceroute-based Internet topology (Jan 2016)
  - Analyze alias resolution data, derive topology graphs at various levels of granularity and make data available (Jan 2016)
  - Enable queries regarding observable performance changes and trends across specific regions of the world (Jan 2016)
  - Create AS-traceroute measurement tool (Jan 2016)
  - Final Report (Jan 2016)

- **Schedule**
  - Phase 1: Applied Research (18 mo, ended Jan 2015)
  - Phase II: Development (12 mo, Feb 2015 - Jan 2016)
  - Optional Phase III: Deployment (6 mo, Feb 2016 - July 2016)

Current Status: Resulting Data Products

Synthesized, annotated, traceroute-based Internet topology from all available sources; Ark traceroutes, BGP (Route Views + RIPE-NCC RIS), and IXPs.

Partly integrated to CAIDA's ASRank, in processing of curating/documenting data set for sharing. (See earlier diagram)

Current Status: Resulting Science

- **Papers**
  - "Internet-Scale IPv4 Alias Resolution with MIDAR" (ToN13)
  - "AS Relationships, Customer Cones, and Validation" (IMC13)
  - "Inferring Multilateral Peering" (CoNEXT13)
  - "A Second Look at Detecting Third-Party Addresses in Traceroute Traces with the IP Timestamp Option" (PAM14)
  - "DRoP: DNS-based Router Positioning" (CCR14)
  - "Spurious routes in Public BGP Data" (CCR14)
  - "Challenges in Inferring Internet Interdomain Congestion" (IMC14)
  - "Inferring Complex AS Relationships" (IMC14)
  - "Measuring and Characterizing IPv6 Router Availability" (PAM15)
  - "IPv6 AS Relationships, Clique, and Congruence" (PAM15)
  - "Resilience of Deployed TCP to Blind Attacks" (IMC13)
  - "Mapping Peering Interconnections at Facility Level" (CoNEXT 15)
- **Community** support: hosted **AIMS 2013, 2014, 2015** (http://www.caida.org/workshops/aims) all reports published in CCR

Benefits

- Improved situational awareness of the Internet through:
  - **Increased completeness**
    - Increased measurement infrastructure
    - Expanded and more efficient probing
    - New methods to synthesize disparate Internet topology data
  - **Increased accuracy**
    - Filter out (some) false link inferences, assess impact
    - Improve AS business relationship inference
  - **Improved richness of topology maps**
    - Better geolocation accuracy
    - Router level: aliases resolved w/2 methods (min FP or max coverage)
    - Increased connectivity at router-level
    - Physical facility awareness
    - IP, router, PoP, and AS-level
    - AS-level annotations: org, type, relationship, performance

~~Competition~~ – Related Work

- RIPE Atlas (http://atlas.ripe.net/)
- Internet Atlas (http://internetatlas.org/)
- iPlane datasets (http://iplane.cs.washington.edu/data/data.html)
- DIMES (http://www.netdimes.org/)
- zMap (https://zmap.io/)
- ISI Census (http://isi.edu/ant/address)

- Renesys (http://www.renesys.com/) recently acquired by Dyn

Next Steps (Data)

Integrate additional links & annotations into AS Relationships

- serial-1 (IMC 2013)
- serial-2 (hybrid, advanced topology)
- serial-3 (Ark, Atlas, iPlane)

Next Steps (Data Accessibility)

Create an interface for **browsing**, **querying**, and **visualizing** the data gathered by the infrastructure.

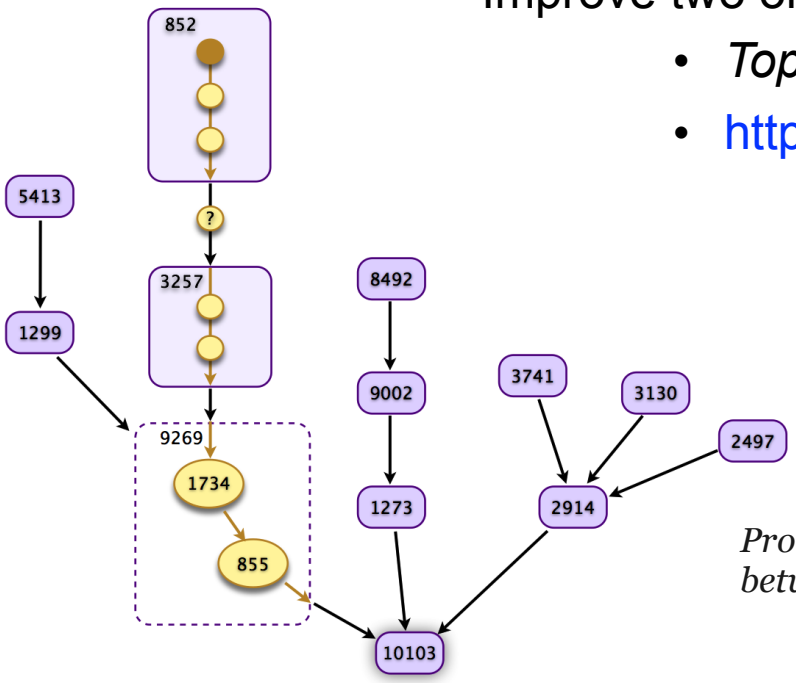

Improve two on-demand topology measurement tools

- *Topo-on-demand* – CLI to Ark platform
- https://vela.caida.org/ web GUI to Ark platform

Prototype visualization showing differences between a traceroute path and BGP AS paths

Next Steps (Data Accessibility)

- *browsing* interface
  - view broad properties and summary statistics over multiple time scales and aggregation levels
    - example: trace counts and response rates; path-length and RTT distributions; inferred AS links

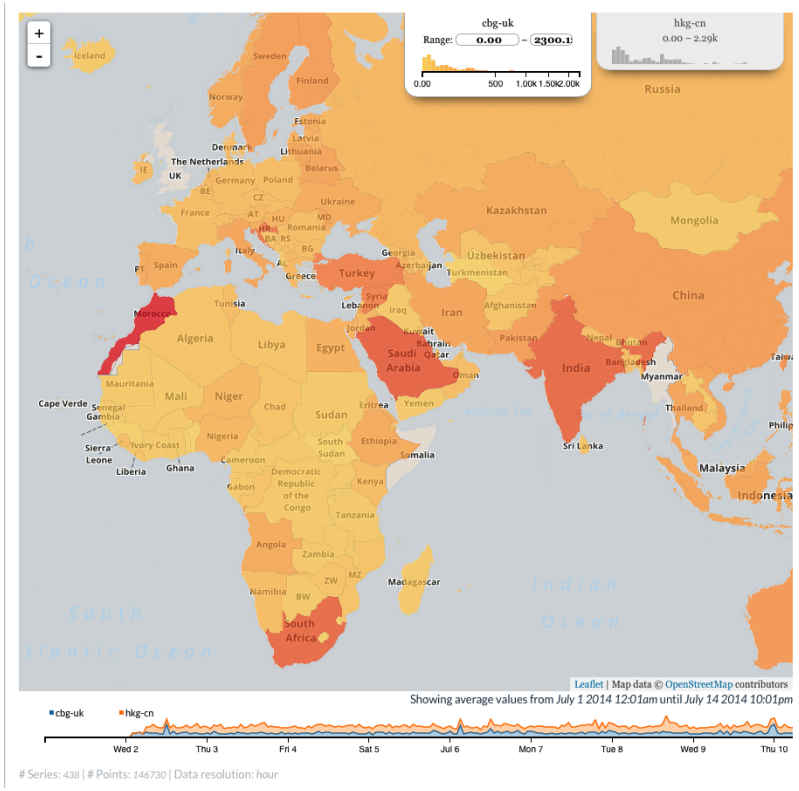

Prototype view of traceroute RTTs implemented with CAIDA's Charthouse

Next Steps (Data Accessibility)

- *query* interface
  - find the most relevant historical data for one's research
    - either directly answers a question, or identifies data to download for further study

  examples:
    - all traceroutes through a given region and time period toward/across a particular prefix/AS/[*country?]
    - router address aliases for a given IP address
    - all inferred links to a router identified by a given IP address
    - all routers in a given city*

*[*blocked on improved geolocation of routers]*

Project Quad Chart

Proposal Title: Cartographic Capabilities for Critical Cyberinfrastructure **15 June 2015**

Photograph or Artist concept / Technical Approach:

Operational Capability/Benefits:

- Expanded measurement platform; provide ongoing Internet topology probing to regularly updated set of targets; contribute results to DHS S&T PREDICT
- Deploy 60+ additional monitors; collect, curate, analyze, and aggregate 1.5-2.5 TB of topology data per year.
- Estimated system maintenance cost: $300K/year.
-

Deliverables and Progress to date:

- Operations and maintenance of globally distributed active measurement infrastructure: 60 new nodes, 118 total
- Developing rich cybersecurity-relevant annotated maps of critical resources at physical and logical levels.
- Release of periodic Internet Topology Data Kit (ITDK)
  - Additional data sources to expand completeness
  - AS to organization, type mapping and validation
  - GUI for interactive validation and correction of AS meta-data and PoP/city-level map

Schedule, Cost and Deliverables:

- Applied Research Phase: Oct 2012-Jan2015
- Development Phase: Feb 2015-Jan 2016
- Deployment Phase: Feb 2015-Jul 2016

Cost:

- $3M

Deliverables:

- Monthly reports; datasets; topo-on-demand monitoring; GUI-based topology maps; final report.

Contract End Date:

- Jan 28 2016 (optional period: Jul 27 2016)

Highlights: Derived and evaluating experimental traceroute-based Internet topology

- Cost
- Schedule
- Technical

POC Information:
Jennifer Ford, UCSD Contracts and Grants
Address: 9500 Gilman Drive, MC 0934, La Jolla, CA 92093-0934
Email/Phone jjford@ucsd.edu / (858)657-5107