CYBER SECURITY DIVISION
2014 PRINCIPAL INVESTIGATORS'

Cartographic Capabilities for Critical Cyberinfrastructure (C4)

CAIDA/UCSD
PI k claffy

13-14 March 2014

Homeland Security
Science and Technology

caida

Team Profile

The Cooperative Association for 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

Customer Need

Global Cybersecurity Challenges

President Obama has declared that the "cyber threat is one of the most serious economic and national security challenges we face as a nation" and that "America's economic prosperity in the 21st century will depend on cybersecurity."

To help address these threats, DHS needs:

- New measurement and data collection technologies
- Infrastructure to improve situational awareness
- Better understanding of the structure, dynamics and vulnerabilities of the global Internet

Approach

- Active measurement using Archipelago measurement infrastructure
  - Ongoing measurements
  - Randomly probe entire IPv4 address space at /24 granularity
  - 83 monitors and growing (35 IPv6, 35 Pi's, 36 RadClock)
- Alias resolution measurements
  - [Every six months]
  - Improved tools and techniques
- Collect and analyze additional data on Autonomous Systems
  - Enriched annotations
  - BGP, WHOIS, performance data
  - [Financial data]

Approach

- Collection and synthesis of data required to publish the Internet Topology Data Kit (ITDK)
  - Data sources: active measurement of multiple topological levels, BGP, DNS, geolocation data
  - Derived data: IP paths, AS paths, DNS lookups, router aliases, device locations
  - Results: AS relationships, AS paths/links, router locations, router to AS assignments, hostnames, router graphs including nodes and links

Increased coverage of Internet

Task 1: Improve completeness of macroscopic Internet maps

Archipelago Measurement Infrastructure

Increased Completeness, Accuracy and Richness of Annotations

Task 2: Increase accuracy of macroscopic Internet maps

AS Ranking of Autonomous Systems

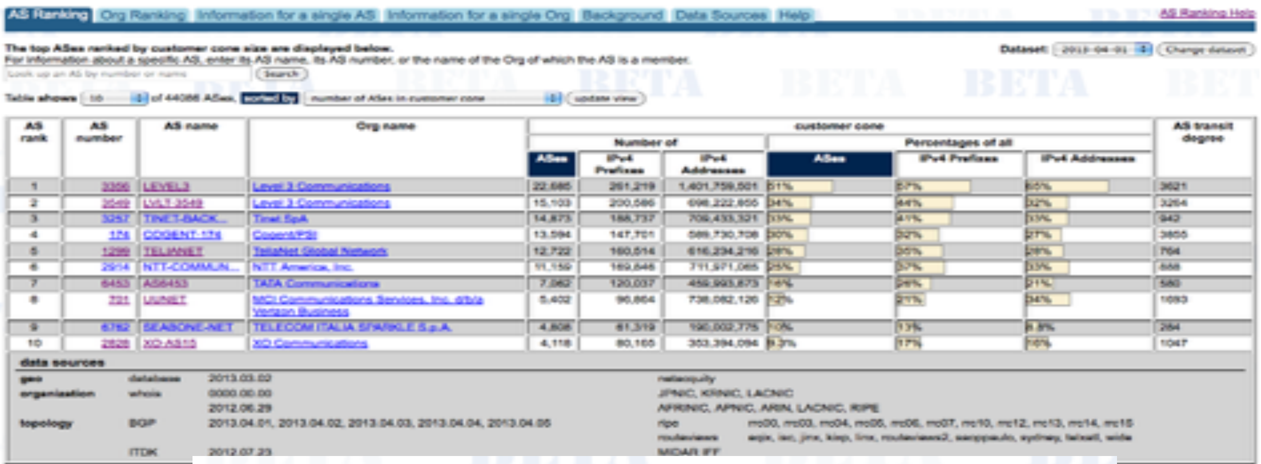

PoP-level map

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 Provider, 1 Sibling, 65 Peer, 3,789 Customer)

Router-level map

Operator feedback

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	

Improved Topology Maps

Task 3: Increase the richness of macroscopic Internet maps

AS Core network visualizations

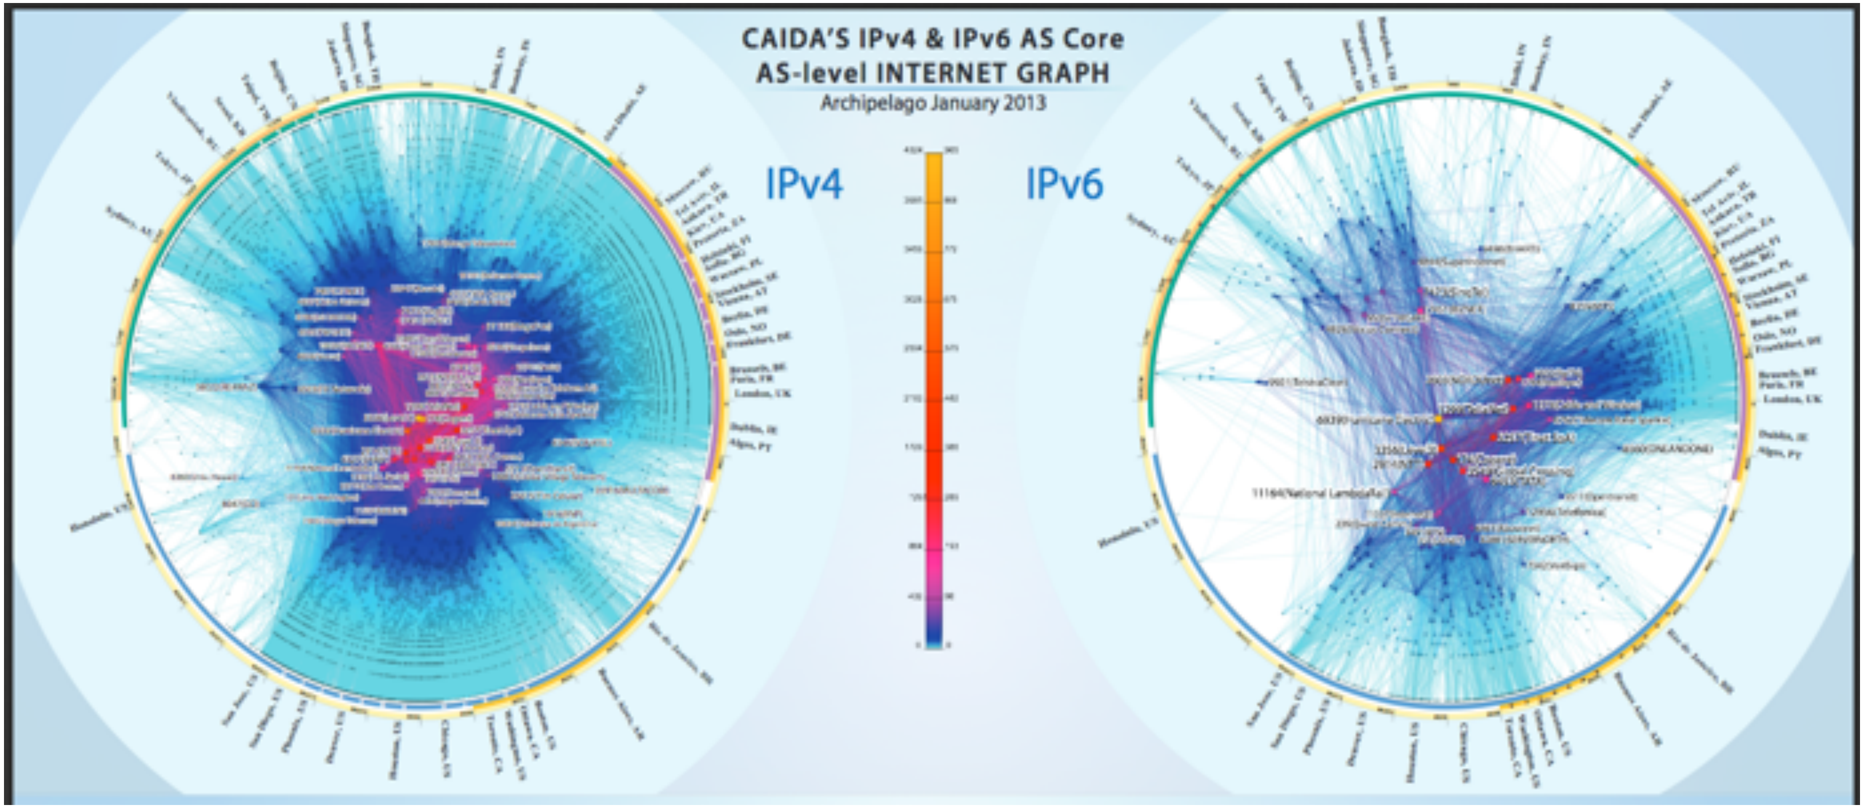

Benefits

- 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
    - Dual maps, aliases resolved with :
      1. MIDAR+iffinder – highest confidence aliases, minimize false positives
      2. MIDAR+iffinder+kapar - increased coverage at cost of false positives
    - Increased connectivity at router-level
    - IP, router, PoP, and AS-level

~~Competition~~ – Related Work

- In academics, we view as related work rather than competition and try to reduce unnecessary redundancy.

- RIPE Atlas (http://atlas.ripe.net/)
- iPlane datasets (http://iplane.cs.washington.edu/data/data.html)
- DIMES (http://www.netdimes.org/new/)
- Renesys (http://www.renesys.com/)
- zMap (https://zmap.io/)

Current Status

- Deliverables (will be late due to funding delay)
  - Monthly data collection (ongoing)
  - Evaluate traceroute-based Internet topology (Aug 2014)
- Milestones
  - Activated 10 new Ark nodes (last 6 months)
  - Evaluated scalable probing algorithms
  - Increased pool of IP addresses for alias resolution
  - Investigated the impact of false link inferences on the router-level, PoP-level, and AS-level graphs (latter a CCR paper)
- Schedule – near term
  - Beta-version of interactive intermediate (PoP/city-level) map validation functionality for testing and feedback (April 2014)
  - Applied Research Phase through March (now July) 2014

Next Steps

- Based on the success of our tech transfer approach on a previous BAA (07-09), we plan to transfer an array of academic research related to homeland security challenges into a production resource of practical utility to DHS needs. We plan to:
  1) release two Internet Topology Data Kits per year;
  2) develop a user-friendly interactive visual interface to topology data and meta-data; and
  3) implement two on-demand topology measurement tools
     1) *Topo-on-demand* – CLI to Ark platform
     2) https://vela.caida.org/ web-based GUI to Ark platform

Recent work co-funded by c4

- Speedtrap: IPv6 alias resolution technique (IMC 2013)
- AS Rank: Improved AS rankings, annotations (IMC 2013)
  – AS-to-Organization mapping: siblings
- IPv4 Transfer Market (CoNEXT 2013)
- Inferring multilateral peering (CoNEXT 2013)
- A Second Look at Detecting Third-Party Addresses in Traceroute Traces with the IP Timestamp Option (PAM 2014)

Speedtrap: IMC 2013

- Speedtrap offers a step toward IPv6 alias resolution at Internet scale
  - uses IP-ID (in fragment headers) to fingerprint IPv6 routers
  - induce velocity in a counter that usually has little
- IPv4 methods (Mercator, DisCarte, RadarGun, MIDAR) do not apply
  - Source routing deprecated in IPv6; no ID field in header
- **Too-Big-Trick**:
  - Send 1300B ICMP echo request.
  - If echo reply > 1280B, send Packet Too Big (PTB)
  - Host should respond to further echo requests with fragmented echo replies with IP-ID until Path-MTU cache entry expires (typically >= 2hrs)
- Developed and validated the technique, code available at: http://www.caida.org/tools/measurement/scamper/

AS Rank: IMC 2013

- 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

http://www.caida.org/publications/papers/2013/asrank/

A Second Look at Traceroute Flaws

- "A Second Look at Detecting Third-Party Addresses in Traceroute Traces with the IP Timestamp Option" (PAM'14)
- Revisit PAM2013 result that would have invalidated decades of traceroute research (but had no validation)
- Underlying (false) assumption: traceroute-reported IP address was off-path if subsequent probe toward same dest could not trigger pre-specified timestamp.
- We inferred an inbound IP interface by attempting to infer if its /30 or /31 subnet mate is an alias of previous hop.

  http://www.caida.org/publications/papers/2014

Improved router geolocation [ongoing]

- Commercial geolocation providers focus on edge hosts
  - home users, commercial servers

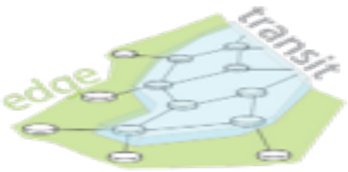

- ITDK needs better transit
  - routers, PoPs

- Solutions
  - DRoP (DNS-based Router Positioning)
  - DDec (DNS Decoded)

DRoP (DNS-based Router Positioning)

automated hostname geohints 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

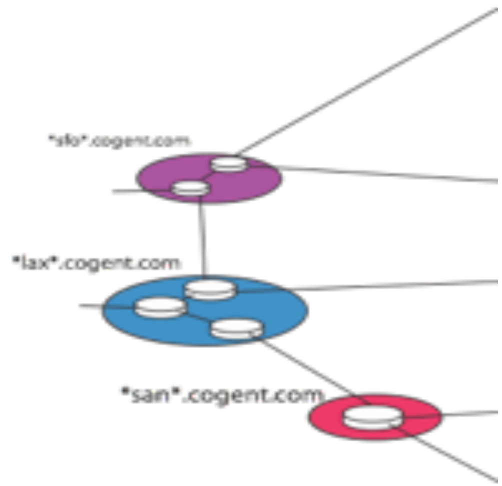

<IATA>\d+.cogent.com
<city name>-gw.routers.es.aau.dk
<CLLI>.([a-z]+\.){2}.ntt.net
<IATA>.above.net

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
- community resource for validated hostname decoding

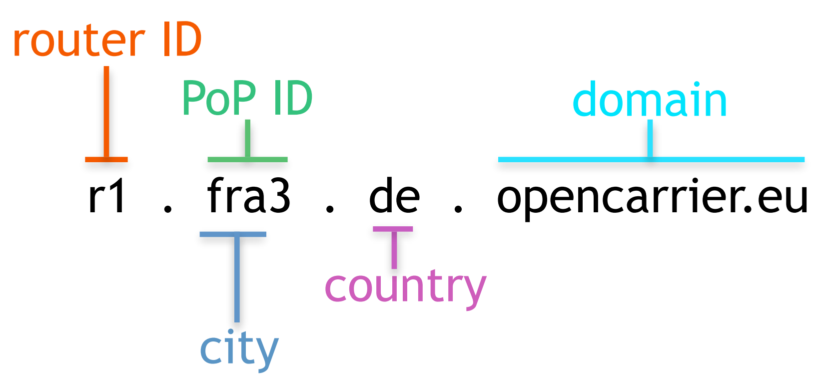

A First Look at IPv4 Transfer Markets

- Much debate on impact of transfer markets on IPv6 adoption
- Part I: empirical study of IPv4 transfer market using lists of transfers published by three RIRs; ARIN, APNIC, RIPE-NCC.
- From the lists of published transfers we found that:
  - 70% of transferred address blocks are legacy blocks
  - Transferred blocks are in lightly utilized before transfer
  - Transferred blocks generally appear in BGP within 3-6 months
  - Observable transfer market thus far seems to be facilitating a healthy redistribution of address space

A First Look at IPv4 Transfer Markets

- Part 2: attempt to detect transfers using public BGP routing table snapshots
- Changes in origin AS for a prefix may be transfers
- Designed filters to rule out transfers due to routing transients, traffic engineering etc.
- Conclusion: even filtered BGP data still too noisy, and produces many apparent transfers.
- Currently investigating the use of DNS data and IP-level (traceroute) paths to detect transfers.

AS2Organization [ongoing]

mapping ASes to the same organization

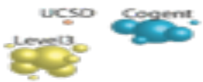

- add accounting code to record sources of all inferences
- delay all inferences until final stage
- flag sources of conflicts with operator feedback

- Evaluate new RIR organization ids

  - Organizations own multiple ASes
  - With improved AS relationship inference, operators have begun to provide more sibling links in feedback
  - Accounting

Inferring Multilateral Peering (CONEXT'13)

- Motivation: topology sources capture only a small fraction of Autonomous System (AS) p2p links
    - 50K peer links in single IXP (Ager2012), 142K peer links (PCH2011)
- We found 206K peering links
    - 88% missing from public BGP data
- Collected and published data
- low measurement cost -> repeatability

Inferring Multilateral Peering (CONEXT'13)

- Two peering paradigms

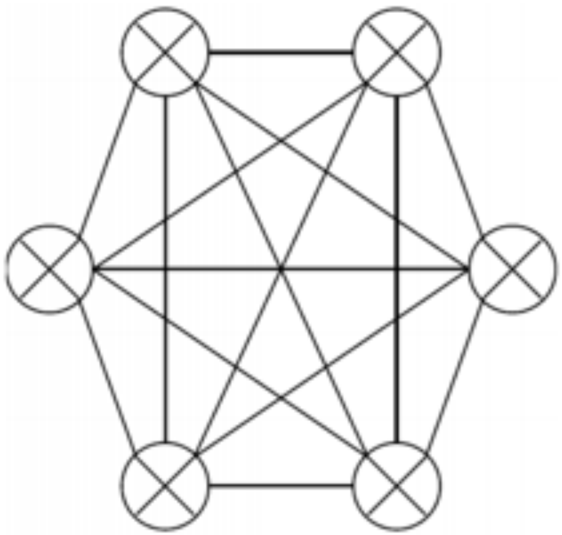

- Bilateral peering
  - Separate BGP session per peering
  - Tight control of peering
  - Poor scalability

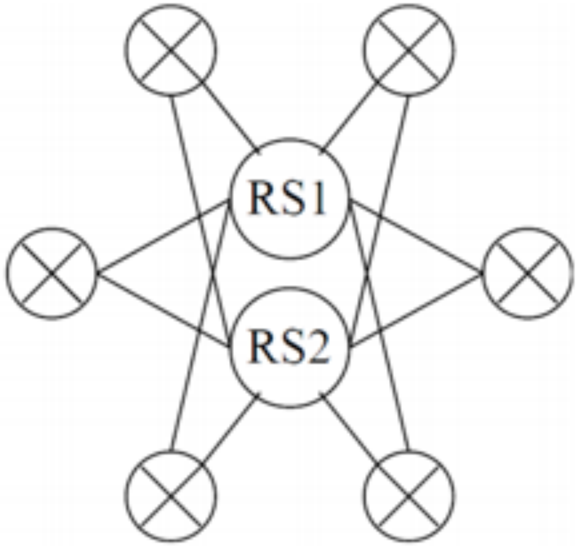

- Multilateral peering (**MLP**)
  - BGP session only with Route Servers (RS) for all links
  - Loose control of peering
  - Great scalability/flexibility

Inferring Multilateral Peering (CONEXT'13)

- Comparison against observable peer-to-peer links

- 12% overlap with passive BGP measurements (Routeviews+RIPE RIS+PCH)
- 2% overlap with active traceroute (Ark+Dimes)

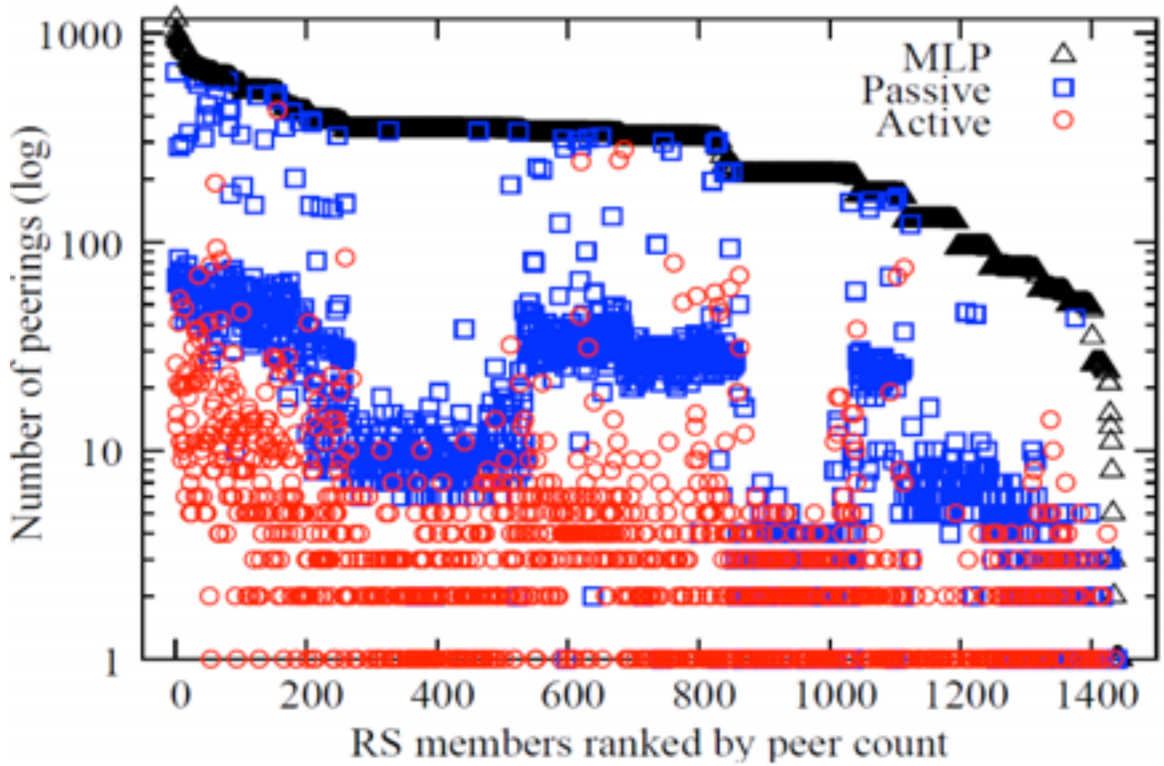

Other recent work of interest

- Inferring Interdomain Congestion (presentation at Yahoo)
- Passive measurements for an Internet Census (CCR)
- Revisiting BGP churn growth (CCR)
- Open Peering by Internet Transit Providers: Peer Preference or Peer Pressure? (Infocom 2014)
- Blog entry: "*CAIDA delivers more data to the public*"

Internet Interdomain Congestion [ongoing]

- Modern peering disputes among access, content, and transit providers manifest as congested links, which affect everybody
- Data on location of congested links is sparse & anecdotal
- Goal: characterize extent of interdomain congestion
  - Methods to detect and localize congestion
  - Map of interdomain links and their congestion state
  - Data to help transparency, empirical grounding of debate
- Trying to infer which network actors are responsible, or the incentives for their behavior is not our focus
- Early work: developing method, seeking feedback/validation

Internet Interdomain Congestion

- This project aims to characterize the extent of interdomain congestion
- **Our goals (1) Methods to detect and localize congestion, (2) Map of interdomain links and their congestion state, (3) Data to improve transparency, empirical grounding of debate**
- Trying to infer which network actors are responsible, or the incentives for their behavior is not our focus
- This is early work: we are still developing the method, and seeking feedback/validation

Congestion Trends

- Three interconnection links of an access network over time

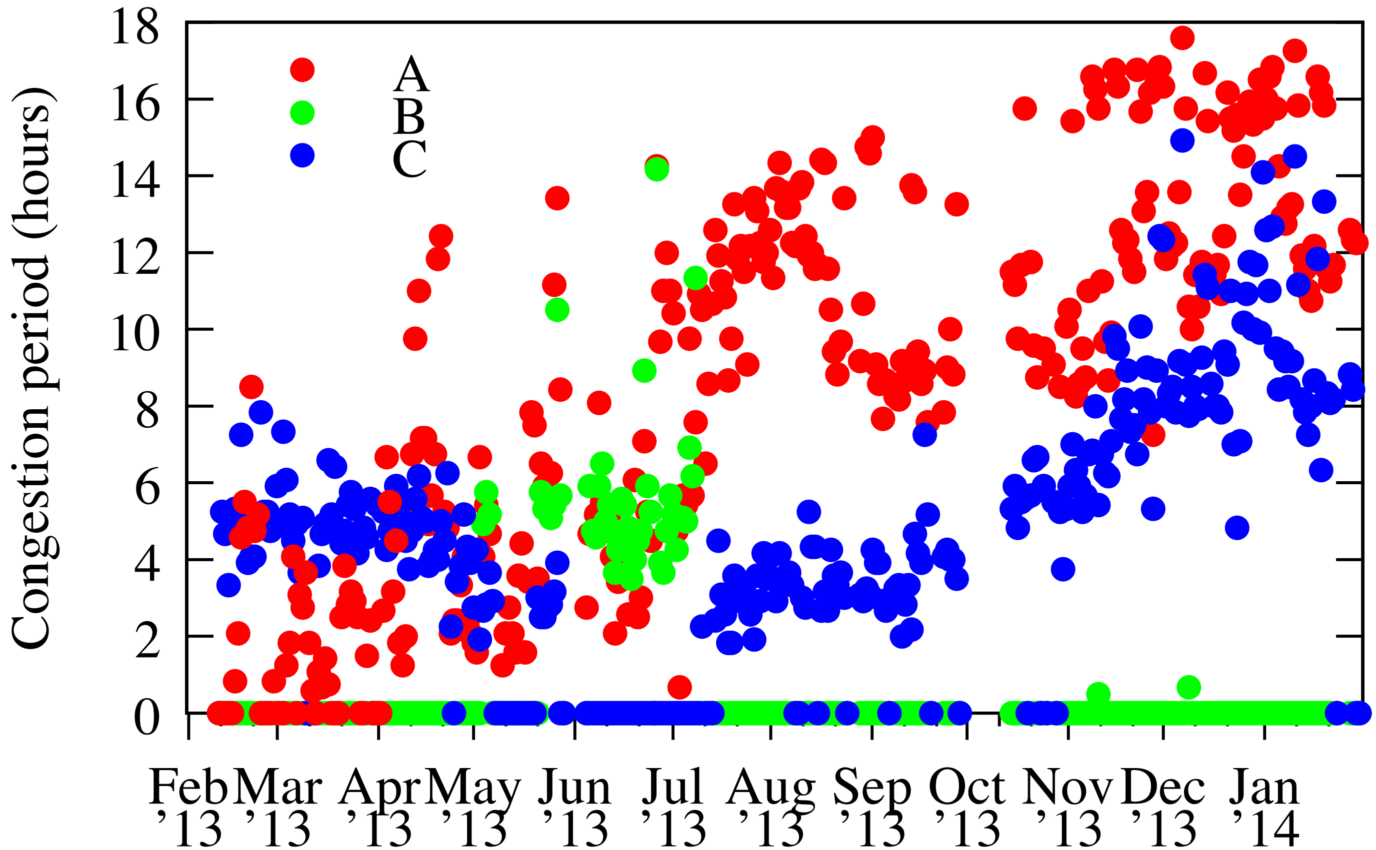

A "passive" Internet Census?

- Passive measurements can discover regions of the IPv4 space not seen by active approaches (e.g., ISI census)

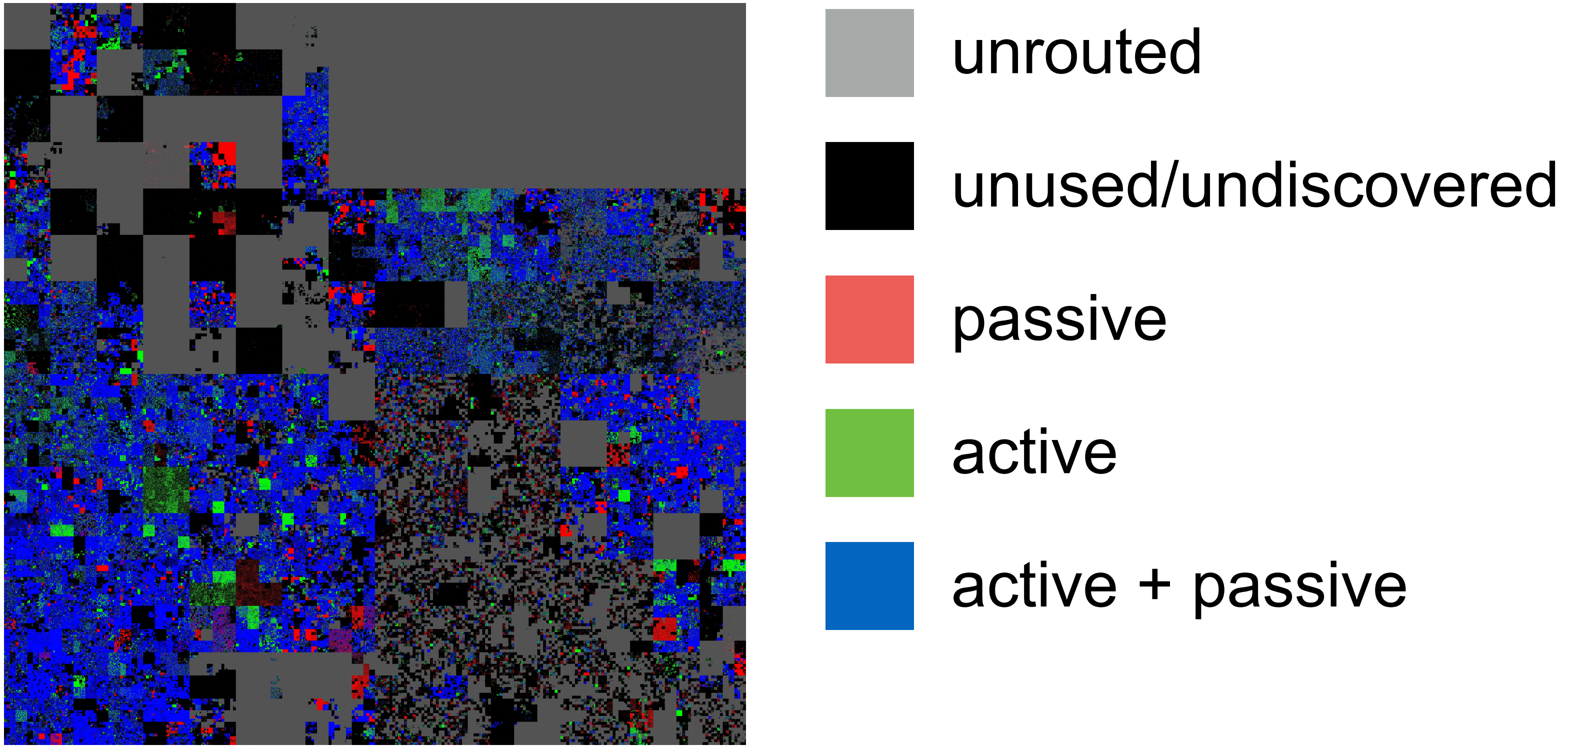

** /24 granularity*

A “passive” Internet Census?

- “Estimating Internet Address Space Usage through Passive Measurements” (CCR January 2014)
  - passive approaches yield significant contribution
  - source-spoofed traffic makes it challenging
  - best approach is to combine active + passive
- work in progress (IMC 2014):
  - more (and diverse) vantage points
  - larger coverage
  - deeper analysis of results

Revisiting BGP Churn Growth (CCRJan’14)

- We found that update churn grows linearly in IPv4 and exponentially in IPv6
- Developed a model of update churn, accounting for topological properties: path length, #updates observed after a routing change, prefix activity in a given interval
- Explains observed linear growth of IPv4 and exponential growth of IPv6 in terms of few measurable parameters
- Result: aggregate IPv6 churn normalized by the size of the topology is constant, similar to IPv4
- For individual prefixes, IPv6 shows more instability -- number of times that prefixes are active every day is higher in IPv6 than in IPv4.

Open Peering by Internet Transit Providers: Peer Preference or Peer Pressure? (Infocom '14)

- Self-reporting in PeeringDB shows that most transit providers advertise open peering policy
- Goal: game-theoretic model for decision process by transit providers to analyze the dynamics leading to open peering.
- Some providers may see incentive in peering with customers of their peers, thereby "stealing" transit traffic from their peers
- Peers then forced to do the same, i.e., peering with customers of peers in order to recover some transit traffic.
- Providers are drawn into a sub-optimal equilibrium due to: 1) Myopic decisions 2) Lack of co-ordination among transit providers
- All providers do not show same attraction: small transit providers more likely to gain from and therefore adopt open peering.

Delivering More Data to Public

- As of February 1, we **converted several popular restricted CAIDA datasets into public datasets**, including skitter and older Ark data.
- We have now made **all IPv4 measurements older than two years (which includes all skitter data) publicly available**.
- Includes derived datasets such as Internet Topology Data Kits (ITDKs).
- To encourage research on IPv6 deployment, we made our **IPv6 Ark topology and performance measurements publicly available.**

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

Contact Information

k claffy

kc@caida.org

http://www.caida.org/