



# Tools, Data and Research on Internet Topology and Geography

*Amogh Dhamdhere*



caida

# CAIDA's Mission Statement




The Center for Applied Internet Data Analysis (CAIDA) is an independent analysis and research group based at the University of California's San Diego Supercomputer Center. CAIDA investigates both practical and theoretical aspects of the Internet.




- CAIDA's active measurement infrastructure
- 102 monitors, growing by 1 or 2 per month
  - 37 IPv6 capable
  - 39 countries (88 cities)
  - 54 Raspberry Pis
- current projects
  - team-probing experiment to collect IPv4 and IPv6 topology
  - alias resolution measurements
  - research experiments, e.g., congestion measurement, spoofing measurement








## Raspberry Pi



- 700 MHz ARM CPU
- 512 MB RAM
- 100 Mbps Ethernet
- 2 x USB 2.0
- SD card slot
- HDMI display output
- **Cost only \$35**
- **Always looking for new vantage points: talk to me later if you can host one!**

# Archipelago monitors and data



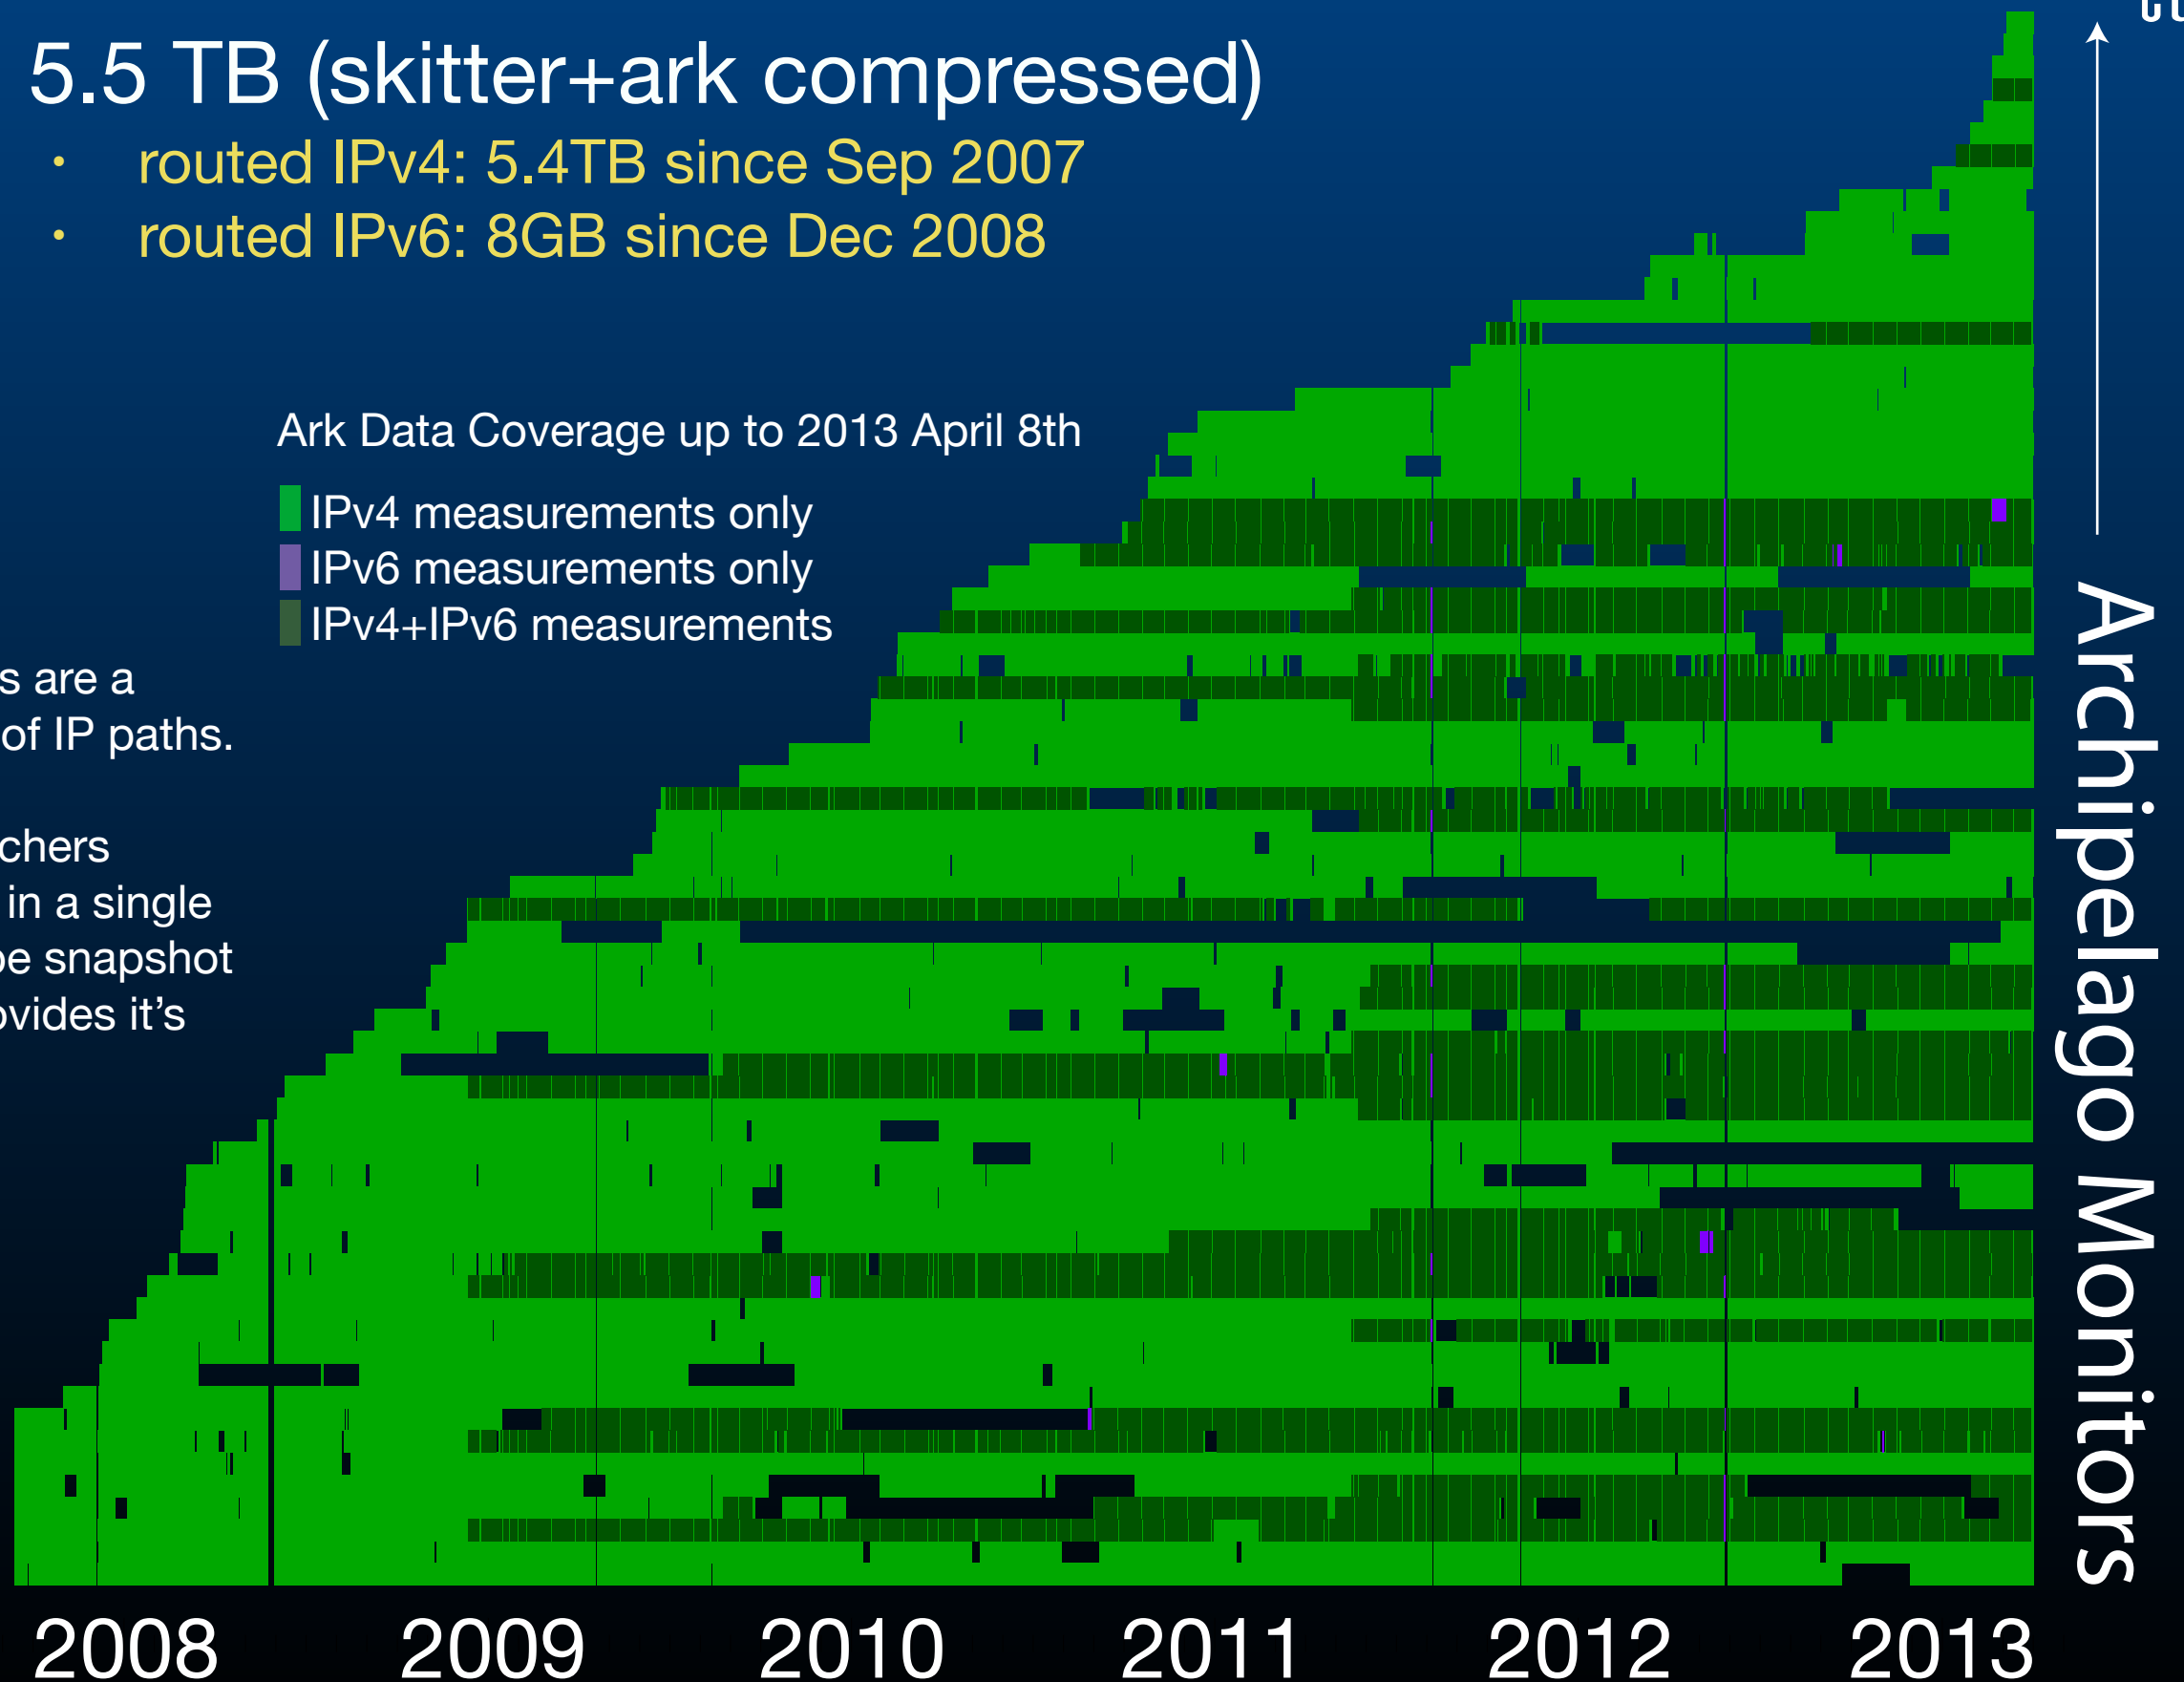
- 5.5 TB (skitter+ark compressed)
  - routed IPv4: 5.4TB since Sep 2007
  - routed IPv6: 8GB since Dec 2008

Ark Data Coverage up to 2013 April 8th

- IPv4 measurements only
- IPv6 measurements only
- IPv4+IPv6 measurements

Raw traces are a collection of IP paths.

For researchers interested in a single microscope snapshot CAIDA provides it's ITDK.



# Supporting rich queries on Ark data

- Goal: support rich queries on traceroute data + geolocation + annotated AS-level topology + router-level topology
- Example 1: Show all traces from a monitor in Canada to destinations in Canada that traverse at least N hops in the United States
- Example 2: Suppose we predict that a certain region will be affected by a natural disaster or political instability. Find all paths that currently traverse that region.
- Example 3: Show connectivity statistics from all monitors to all probed addresses in a given country
- Which types of queries would you like to see?



# Vela: Interactive topology-on-demand

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

- Vela: **interactive interface** to on-demand measurements from Ark monitors, currently ping and traceroute

## Create a Basic Measurement

Define a measurement to ping or traceroute a single target from a single source.

**Destination**  
Enter an address/prefix/hostname:

**Method**  
☒ ping  
☐ traceroute

**Protocol**  
☒ ICMP  
☐ UDP  
☐ TCP  
Note: ICMP is the only supported protocol for ping.

**Vantage Point**  
     
Monitors with IPv6 have an asterisk next to their name

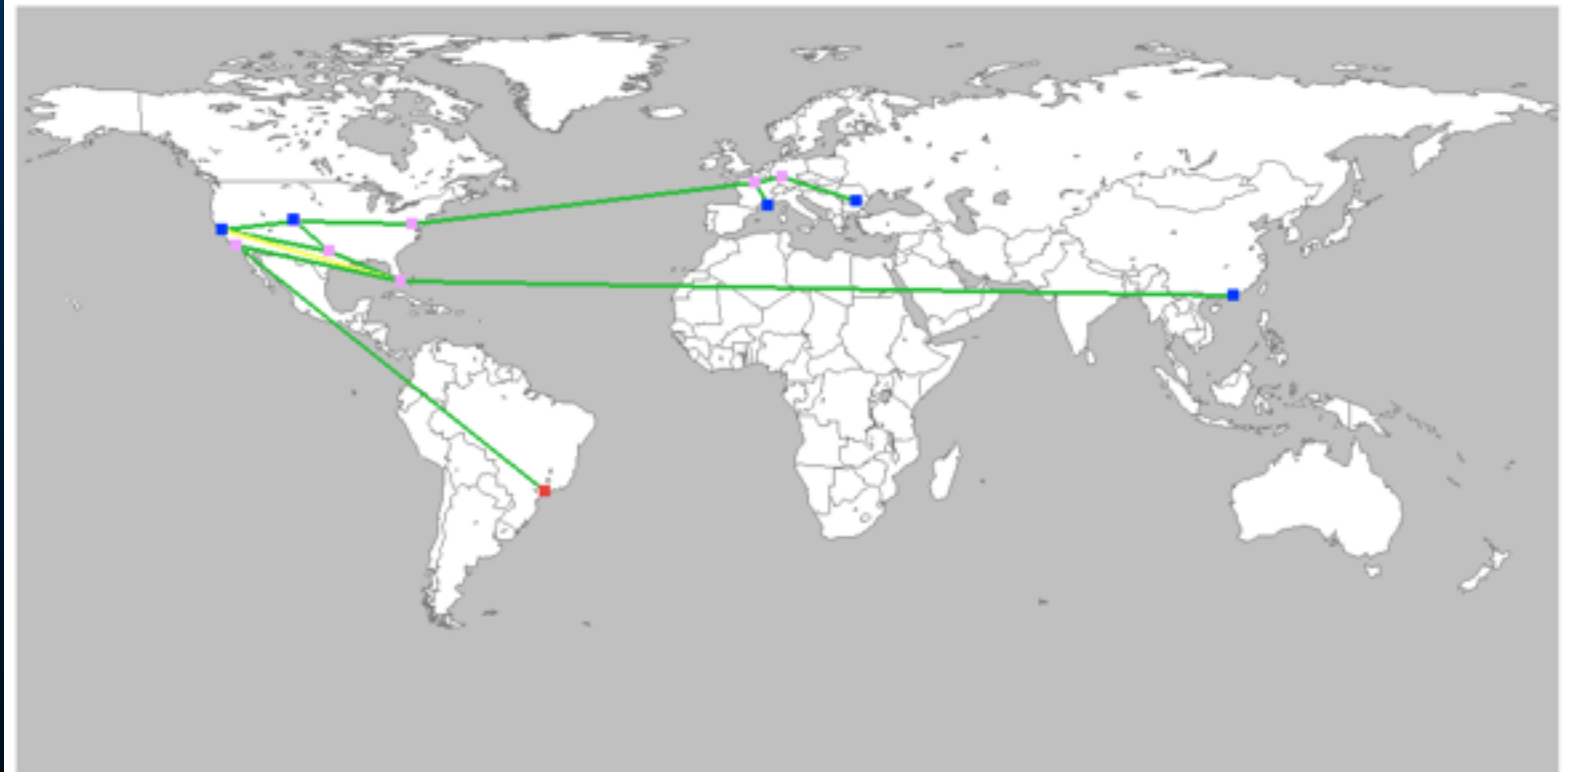
- business
  - bma-se \*
  - fmo-de
  - rek-is
  - sql-us \*
- commercial network
  - bjc-us
  - cdg-fr
  - hkg-cn \*
  - jfk-us \*
  - otp-ro \*
  - sjc2-us \*
- community network
  - vie-at \*
- educational network
  - sao-br
  - zrh2-ch \*

[Home](#)

## traceroute to [sao2-br.ark.caida.org](http://sao2-br.ark.caida.org) from *commercial network (6)* using ICMP

### Traceroute Geo Map

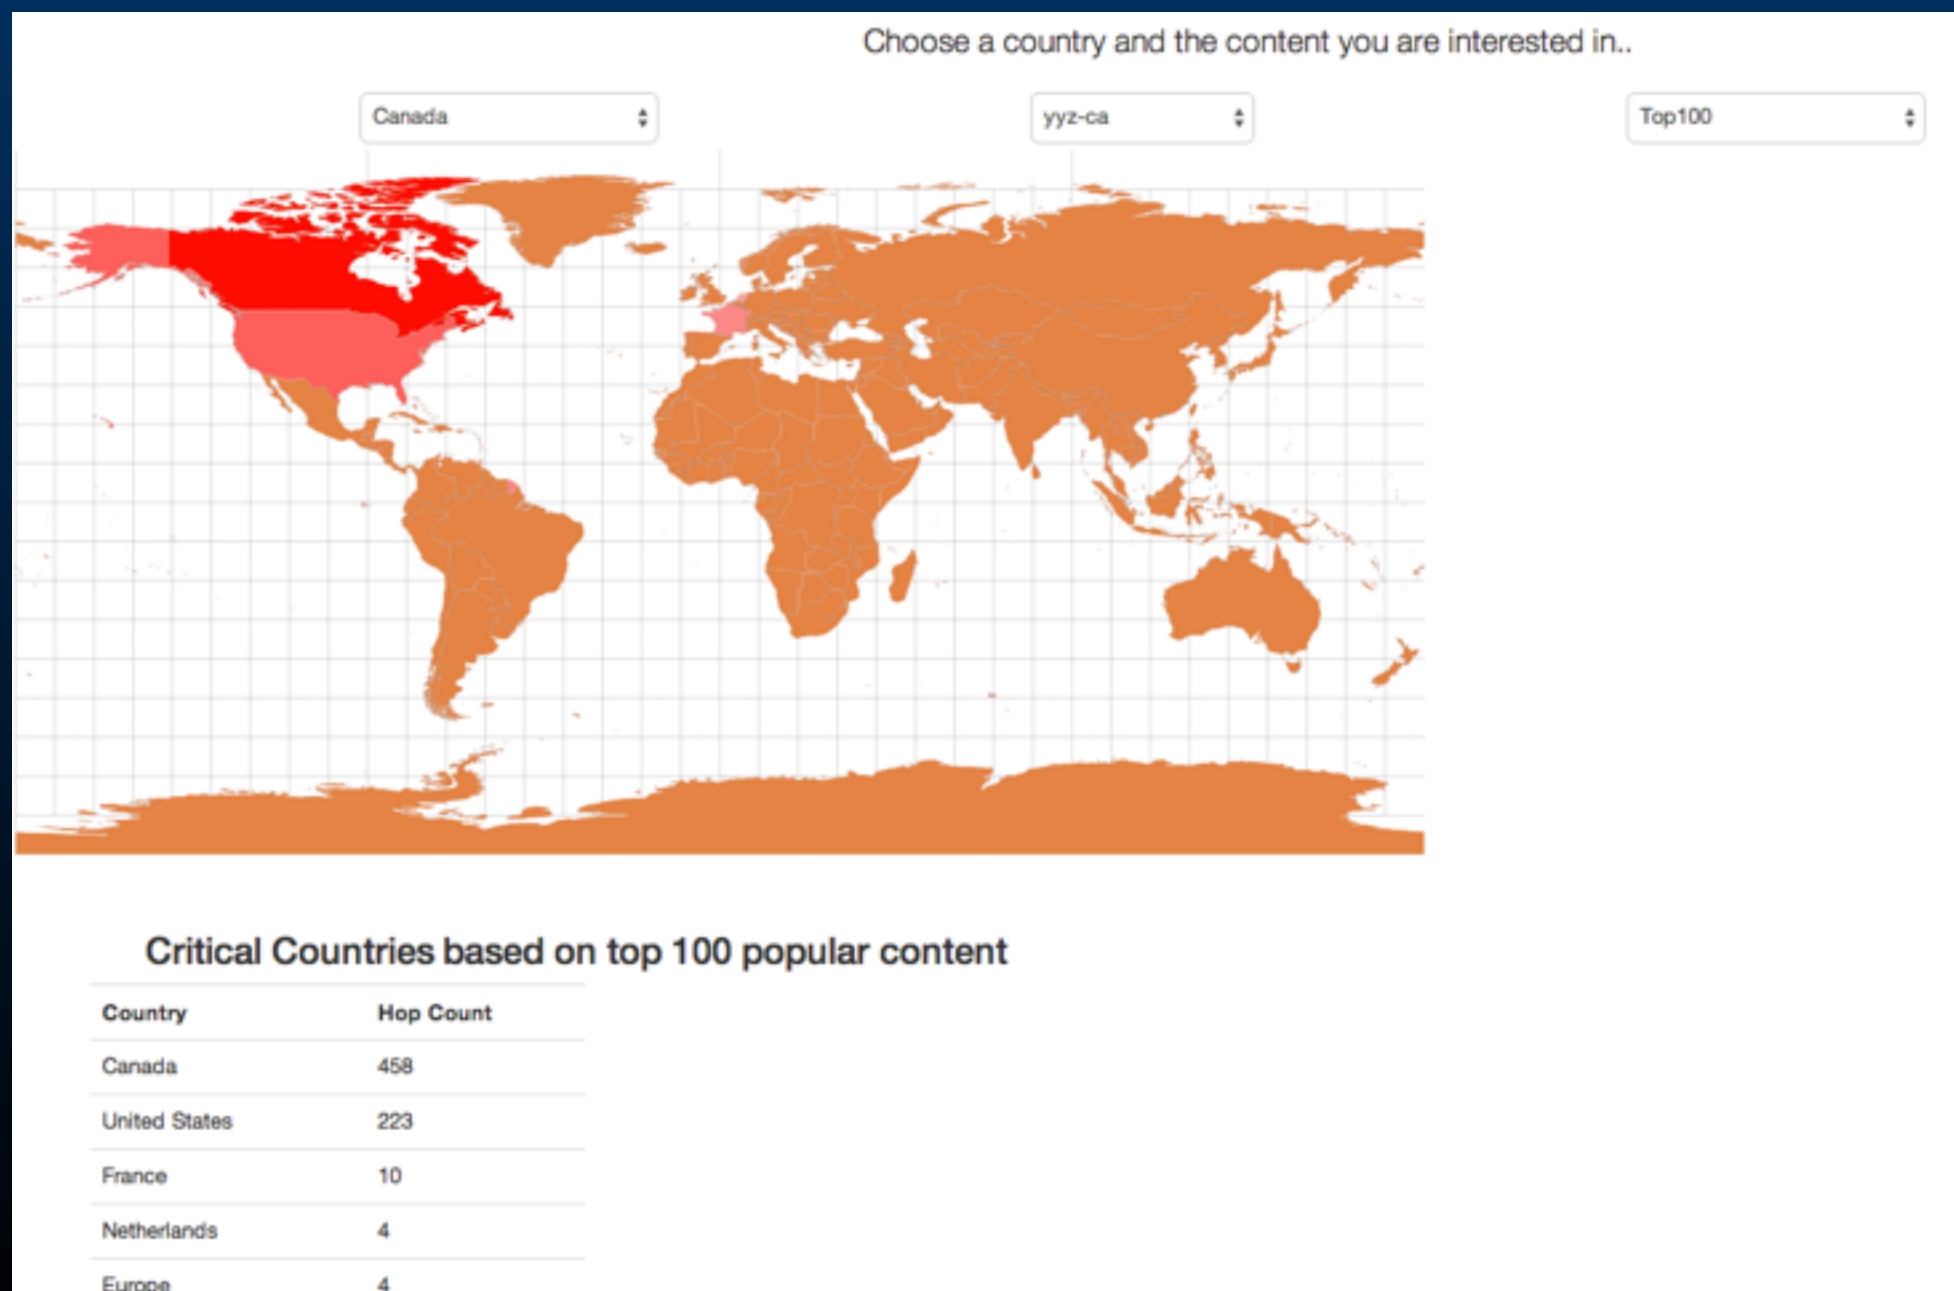
Node Color Key: ■ Source ■ Intermediate ■ Destination  
Link Color Key: — Direct — Indirect



Other Views: [USA](#) | [South America](#) | [Europe](#) | [China](#) | [Japan](#)

# Building specialized tools on top of Ark

- Student project: Visualize paths from Ark monitor in a country to top content (per Alexa list) in that country (work in progress)

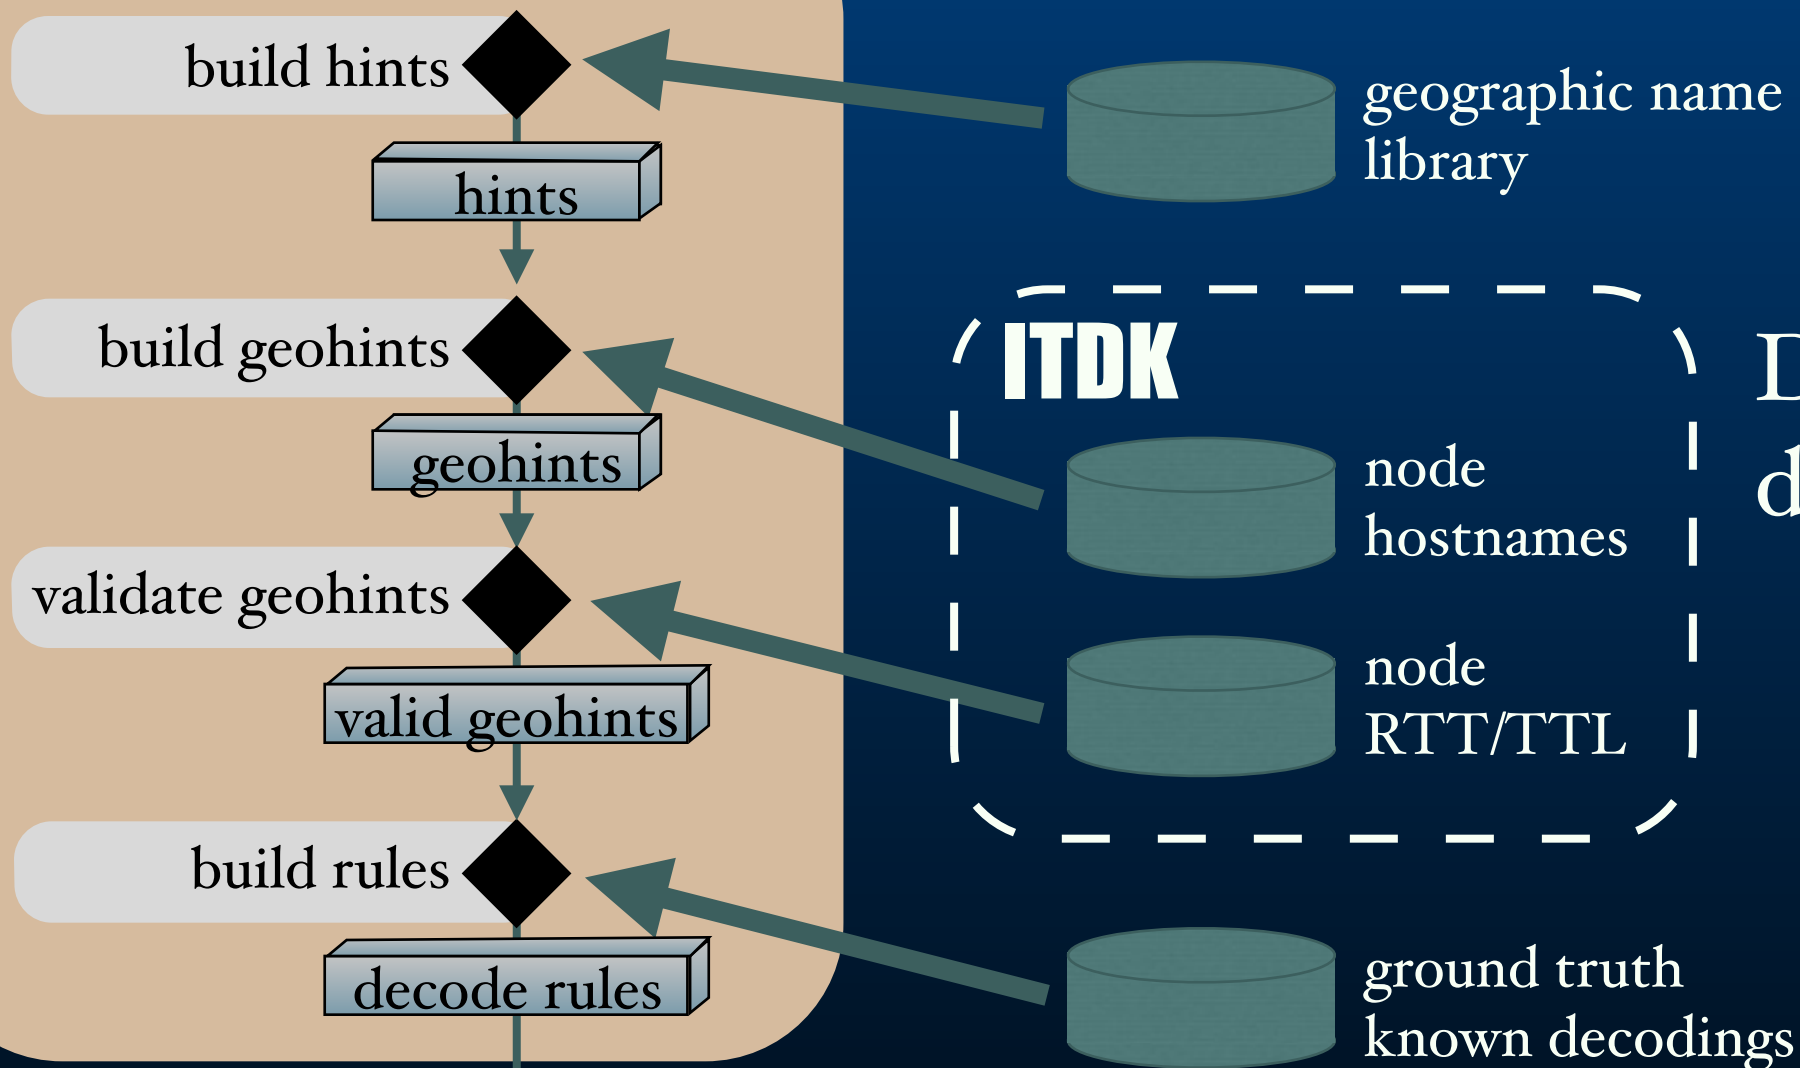





# DNS-based geolocation

[ddec.caida.org](http://ddec.caida.org)

## DRoP



DRoP: automated  
dns hint detection



DDec: public  
interface for lookups  
and corrections

## AS business relationships, customer cones, ranking

AS Ranking | Org Ranking | Information for a single AS | Information for a single Org | Background | Data Sources | Help

The top ASes ranked by customer cone size are displayed below.  
For information about a specific AS, enter its AS name, its AS number, or the name of the Org of which the AS is a member.

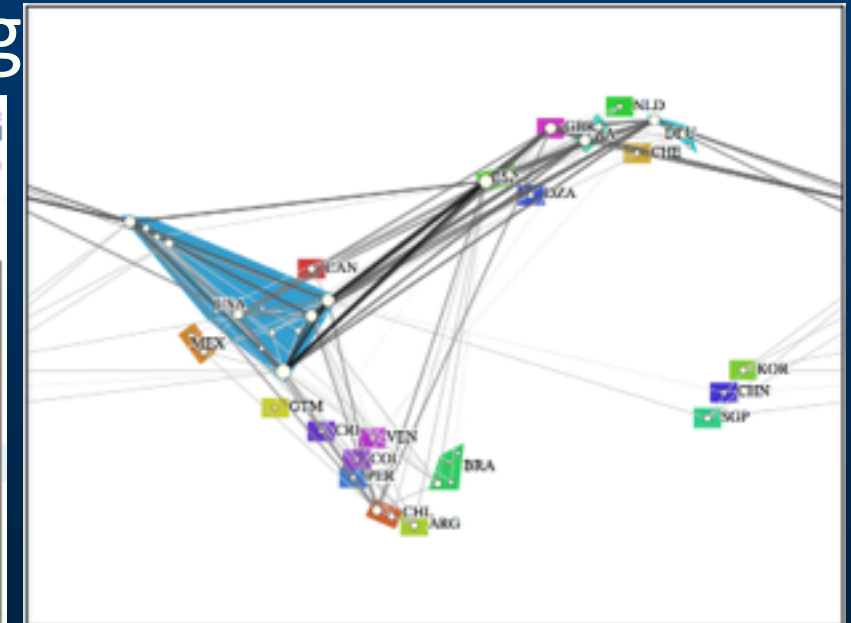
Look up an AS by number or name

Table shows 10 of 44086 ASes, sorted by number of ASes in customer cone

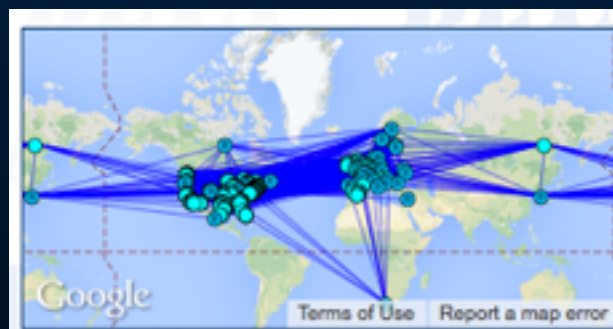
AS rank	AS number	AS name	Org name	Number of			Percentages of all			AS transit degree
				ASes	IPv4 Prefixes	IPv4 Addresses	ASes	IPv4 Prefixes	IPv4 Addresses	
1	3356	LEVEL3	Level 3 Communications	22,685	261,219	1,401,759,501	51%	57%	65%	3621
2	3549	LVL-3549	Level 3 Communications	15,103	200,586	696,222,855	34%	44%	52%	3264
3	3257	TINET-BACK...	Tinet SpA	14,873	188,737	709,433,321	33%	41%	53%	942
4	174	COGENT-174	Cogent/PSI	13,594	147,701	589,730,708	30%	32%	27%	3855
5	1299	TELIA.NET	TeliaNet Global Network	12,722	160,514	616,234,216	28%	35%	28%	764
6	2914	NTT-COMMUN...	NTT America, Inc.	11,159	169,846	711,971,065	25%	37%	33%	888
7	6453	AS6453	TATA Communications	7,062	120,037	459,993,873	16%	26%	21%	580
8	701	UUNET	McI Communications Services, Inc. d/b/a Verizon Business	5,402	96,864	738,082,126	12%	21%	34%	1693
9	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	4,808	61,319	190,002,775	10%	13%	8.8%	284
10	2828	XO-AS15	XO Communications	4,118	80,165	353,394,094	9.3%	17%	16%	1047

**data sources**

geo	database	2013.03.02	netacuity
organization	whois	0000.00.00	JPNIC, KRNIC, LACNIC
		2012.06.29	AFRINIC, APNIC, ARIN, LACNIC, RIPE
topology	BGP	2013.04.01, 2013.04.02, 2013.04.03, 2013.04.04, 2013.04.05	ripe, mc00, mc03, mc04, mc05, mc06, mc07, mc10, mc12, mc13, mc14, mc15
			routeviews, eqix, isc, jinx, kixp, linx, routeviews2, saopaulo, sydney, telxatl, wide
	ITDK	2012.07.23	MIDAR IF



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 1 65 3,789  
Provider Sibling Peer Customer

Router-level map

## Operator feedback

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

# Recent relevant research



- Inferring which networks peer at which IXPs  
“Inferring Multilateral Peering”, Giotsas, Zhou, Luckie, Claffy, ACM CoNEXT 2013
- Mining historical peeringDB data for colocation at IXPs, peering policies, geographical expansion  
“Using PeeringDB to understand the Peering ecosystem”, Lodhi, Larson, Dhamdhere, Dovrolis, Claffy, ACM SIGCOMM CCR 2014
- Investigating connectivity in the LACNIC region  
“Lacnic Connectivity”, Lutu, Bagnulo, Dainotti, Dhamdhere, Claffy, In progress
- <http://www.caida.org/publications/>



Thanks!  
[amogh@caida.org](mailto:amogh@caida.org)  
[www.caida.org](http://www.caida.org)