AS Core: Visualizing the Internet

CAIDA

SDSC/UCSD

CSE 91
4 March 2011

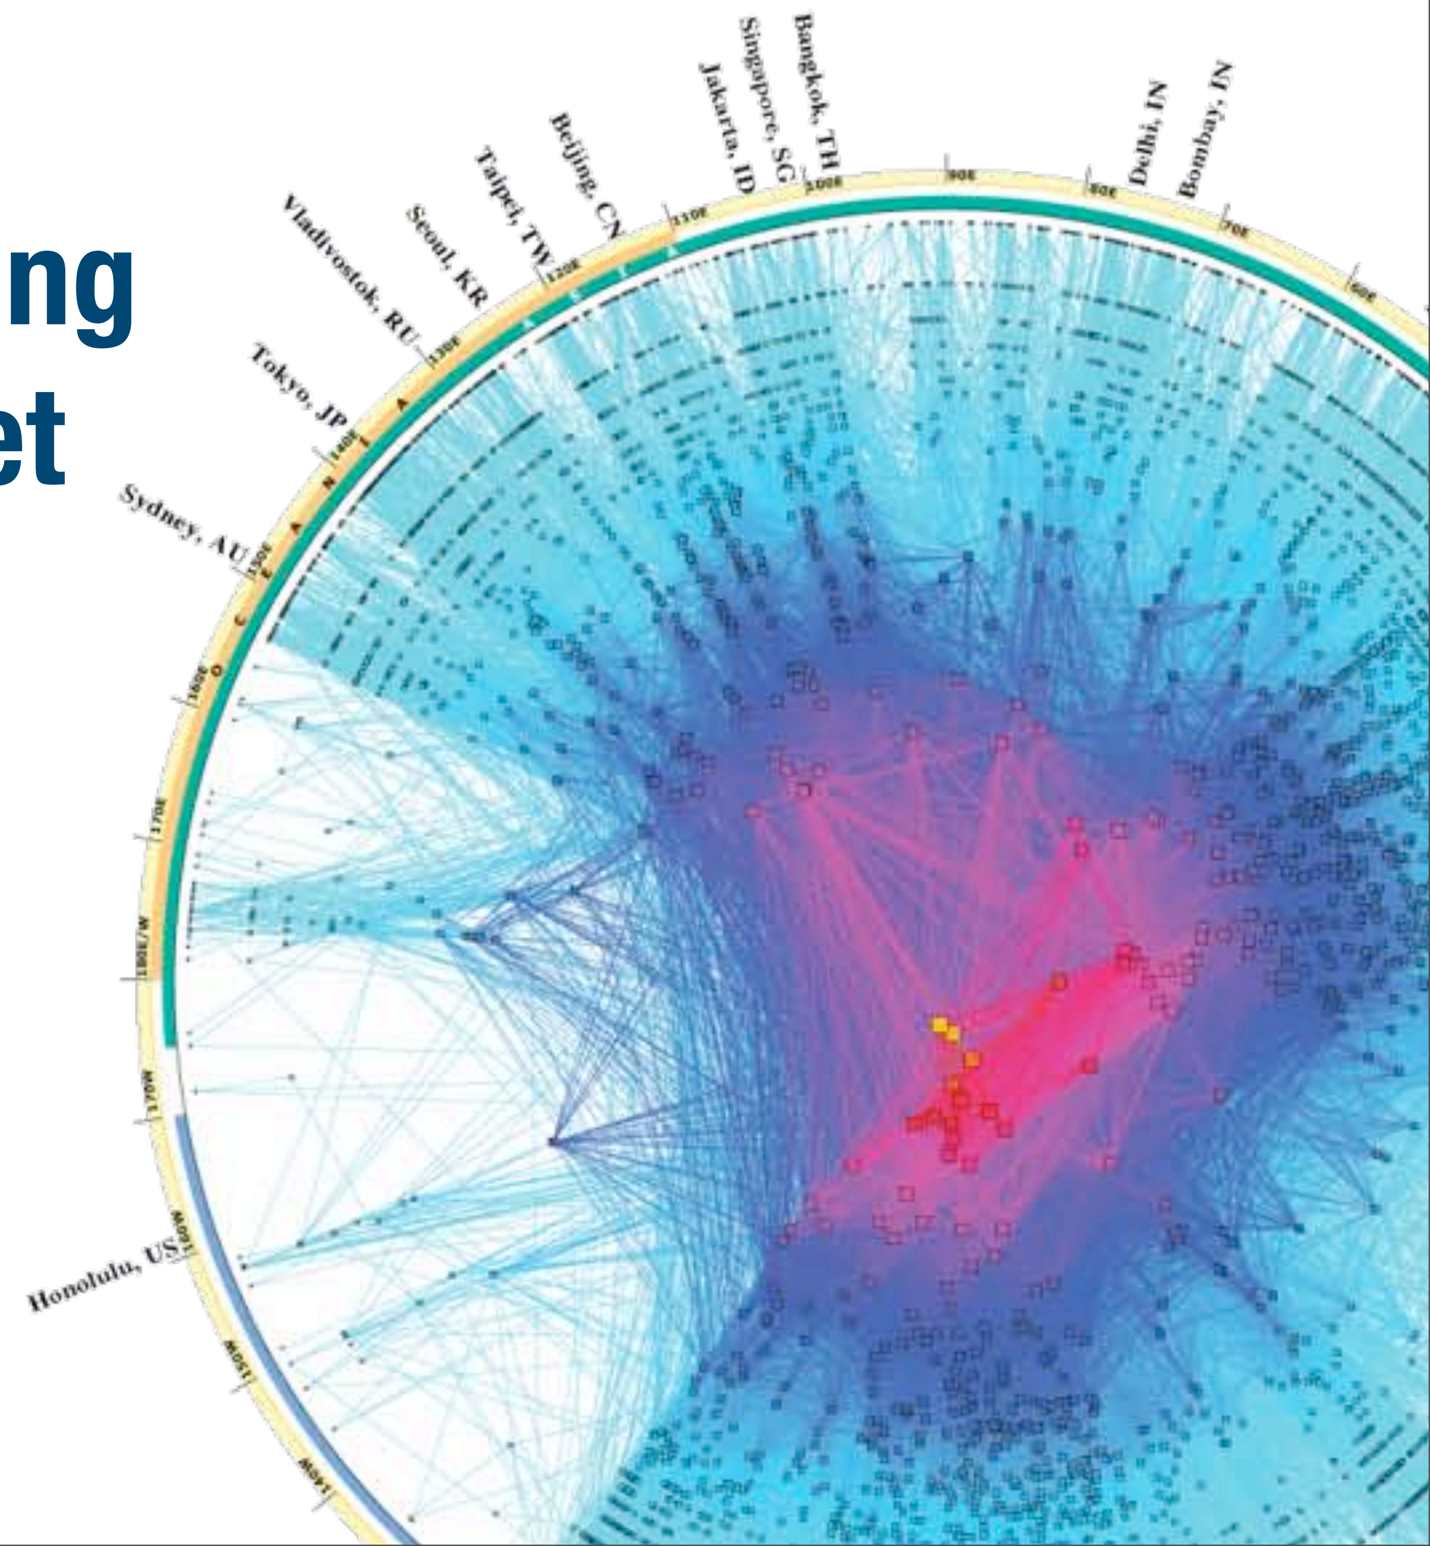

overview

- overview
- data sources
- data processing
- visualization breakdown
- IPv4 vs IPv6
- summary

what we want

Provide a visual representation of the AS level Internet.

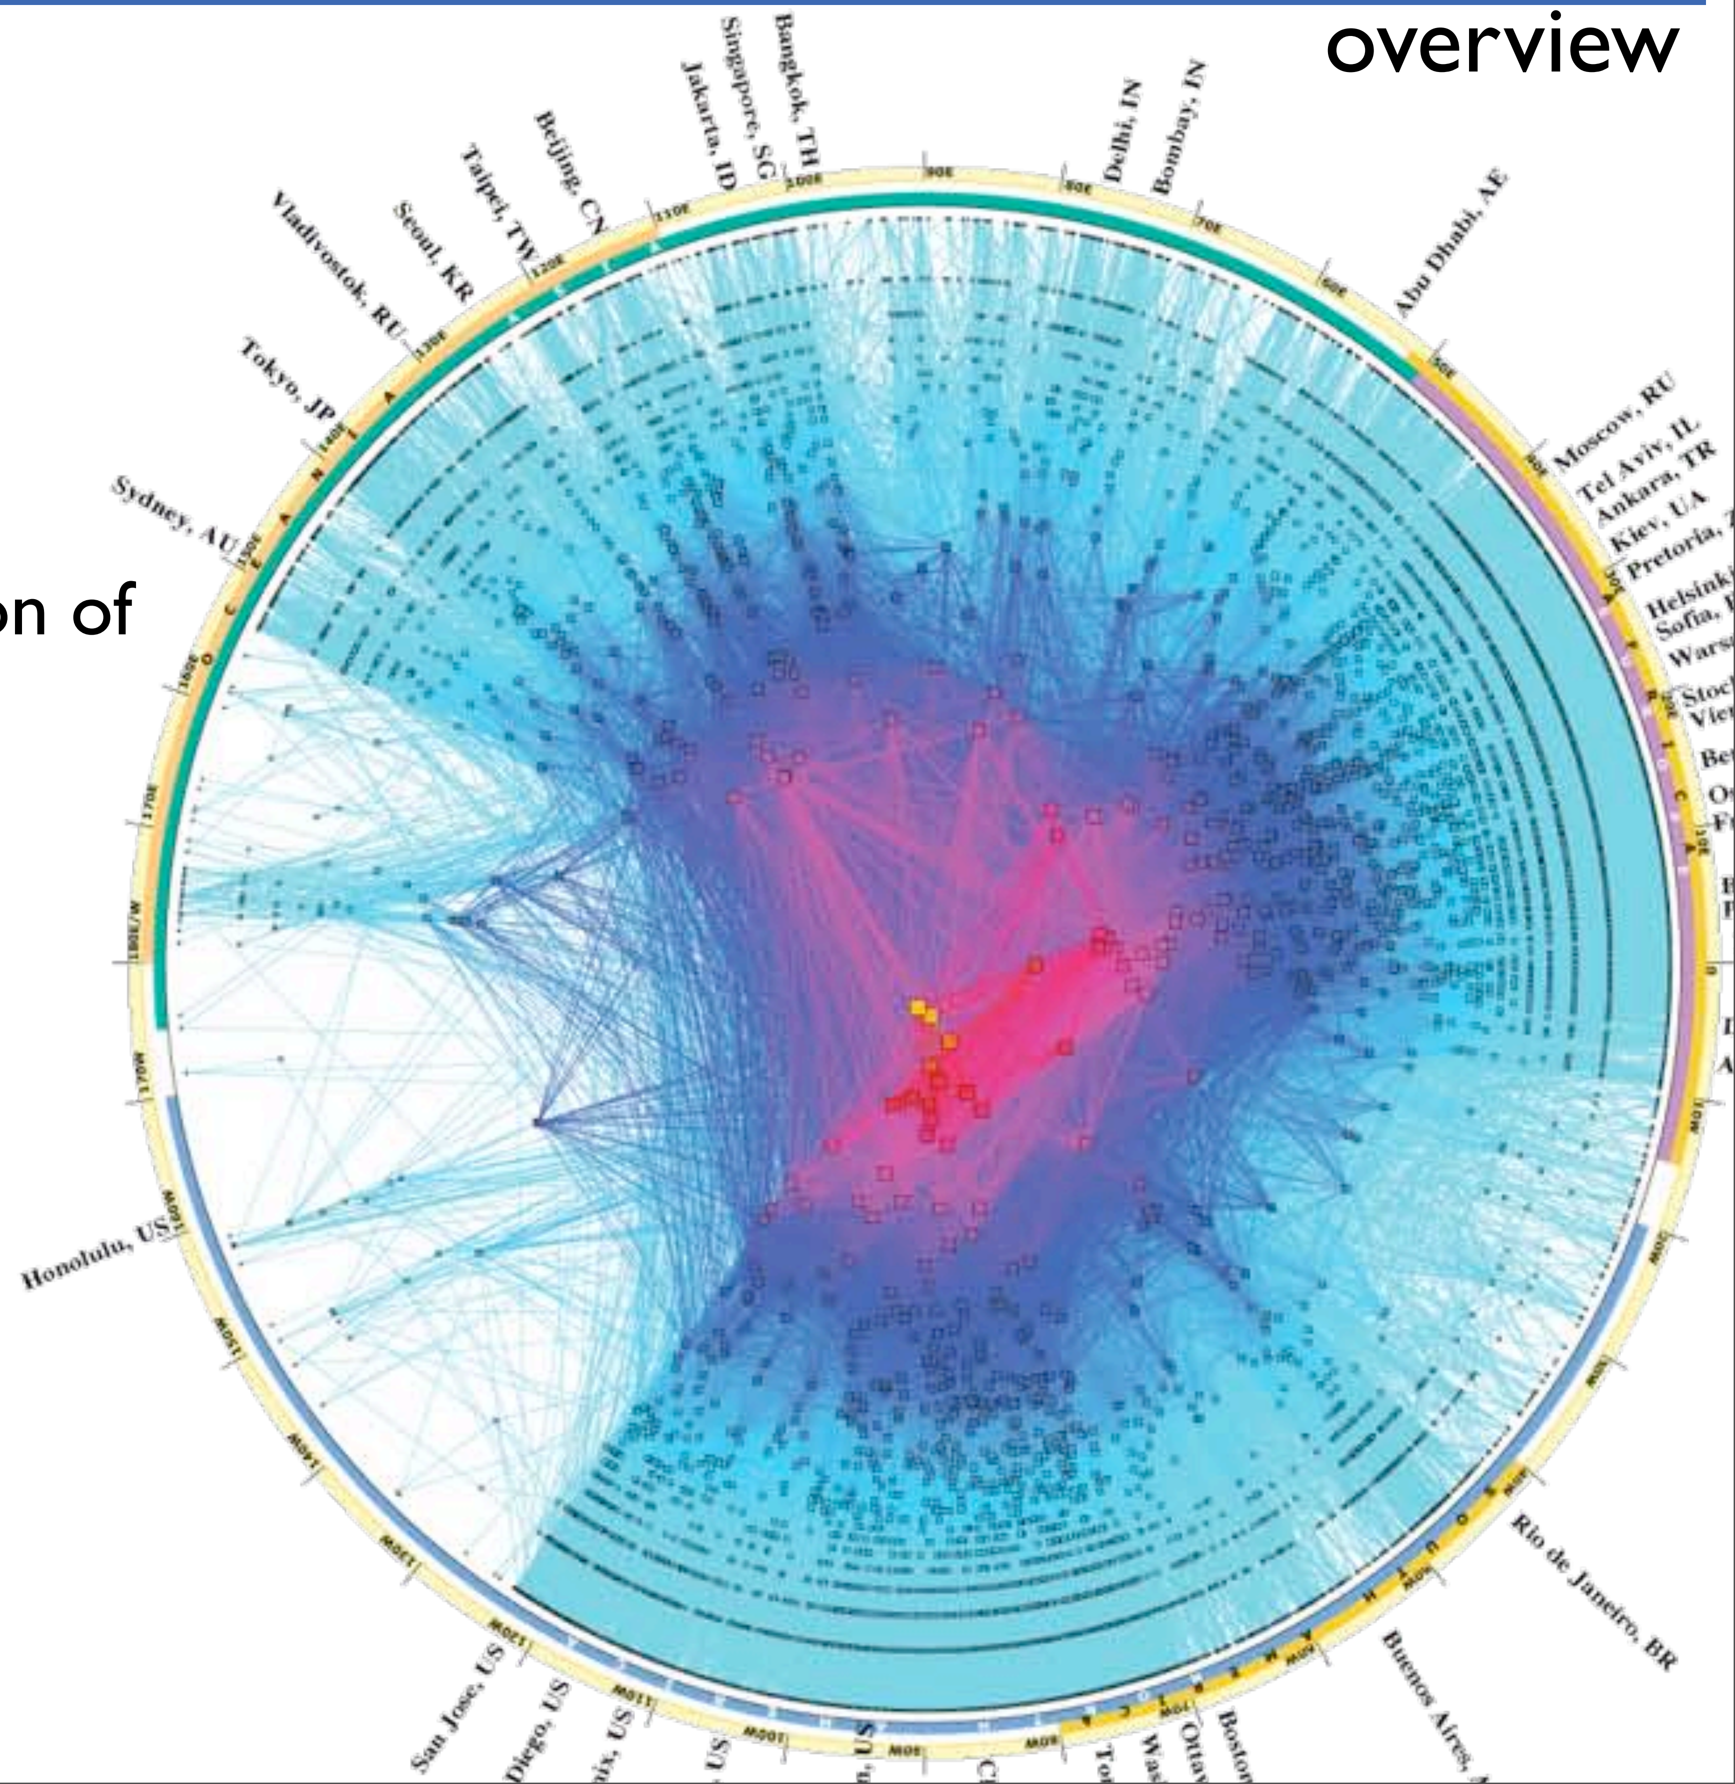

what are the nodes?

Autonomous System (AS)
an entity in the routing system
that announces and provides connectivity to
networks through a global routing protocol.

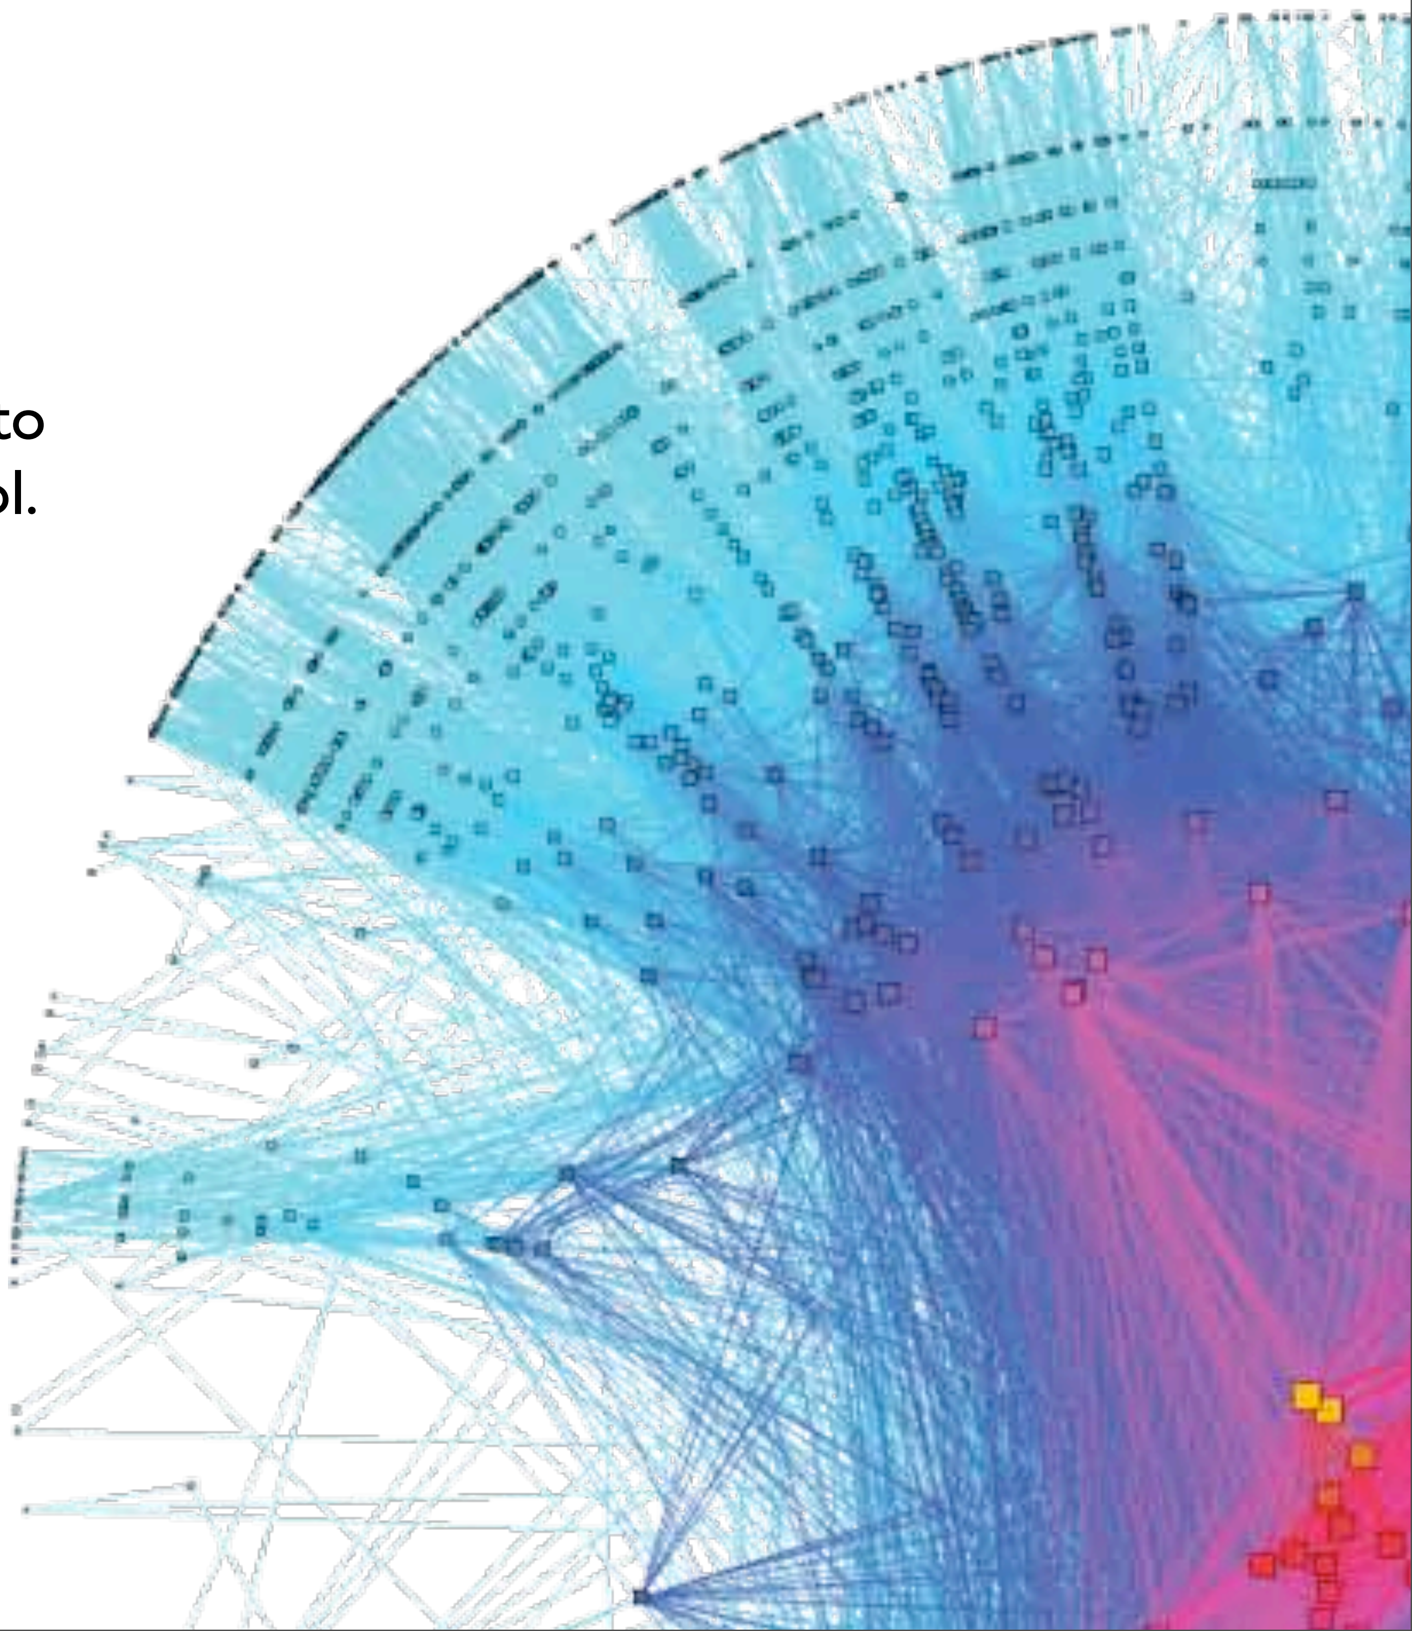

what are the nodes? (2)

Each AS is roughly a company or network operator.
UCSD has several.
Some companies use multiple ASes (M&As), so not one-to-one.

On the graph: A single node is a single AS, although nodes with the same coordinate values will overlap.

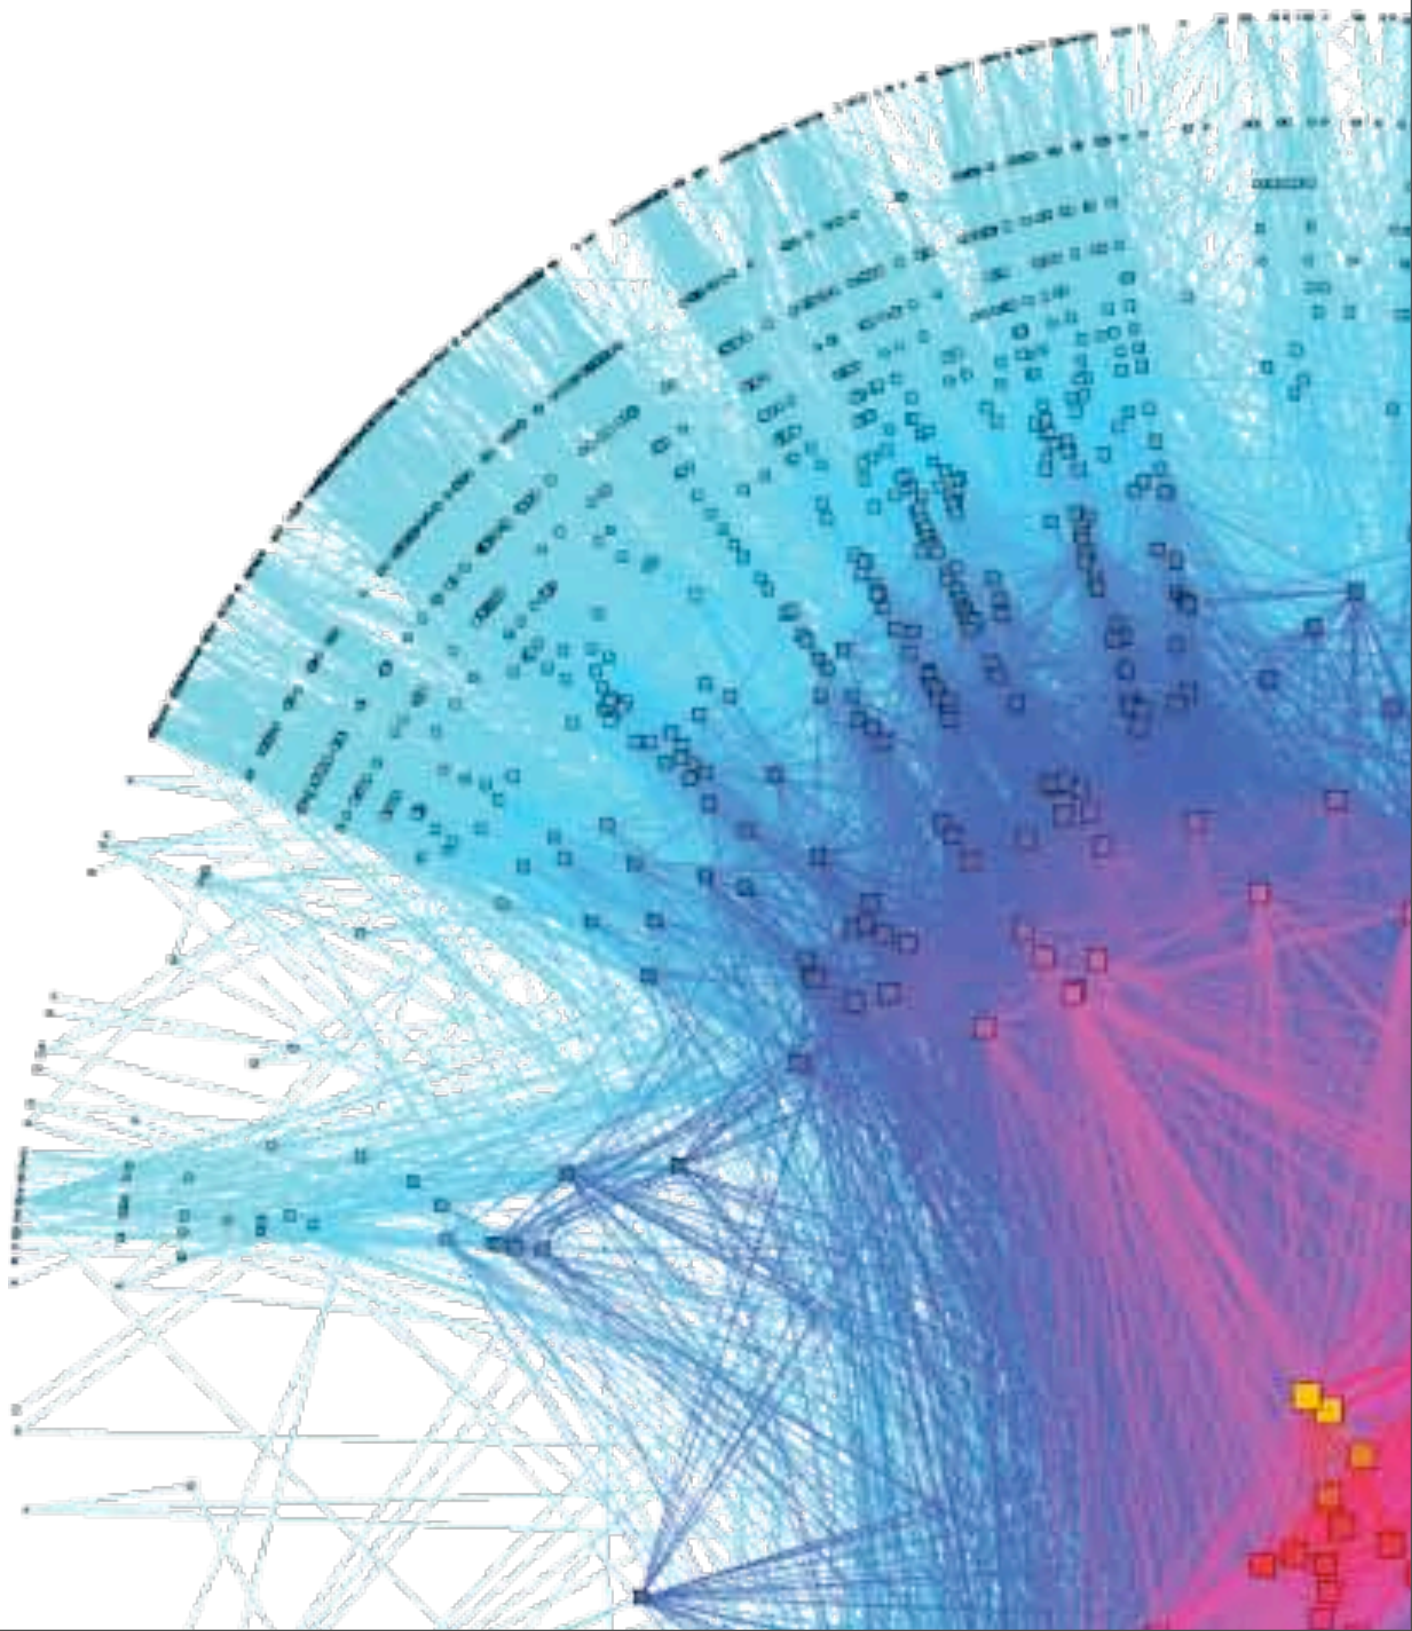

what we need to draw a node

- AS's name
- AS's longitude
- AS's neighbors
- AS's degree (# neighbors)

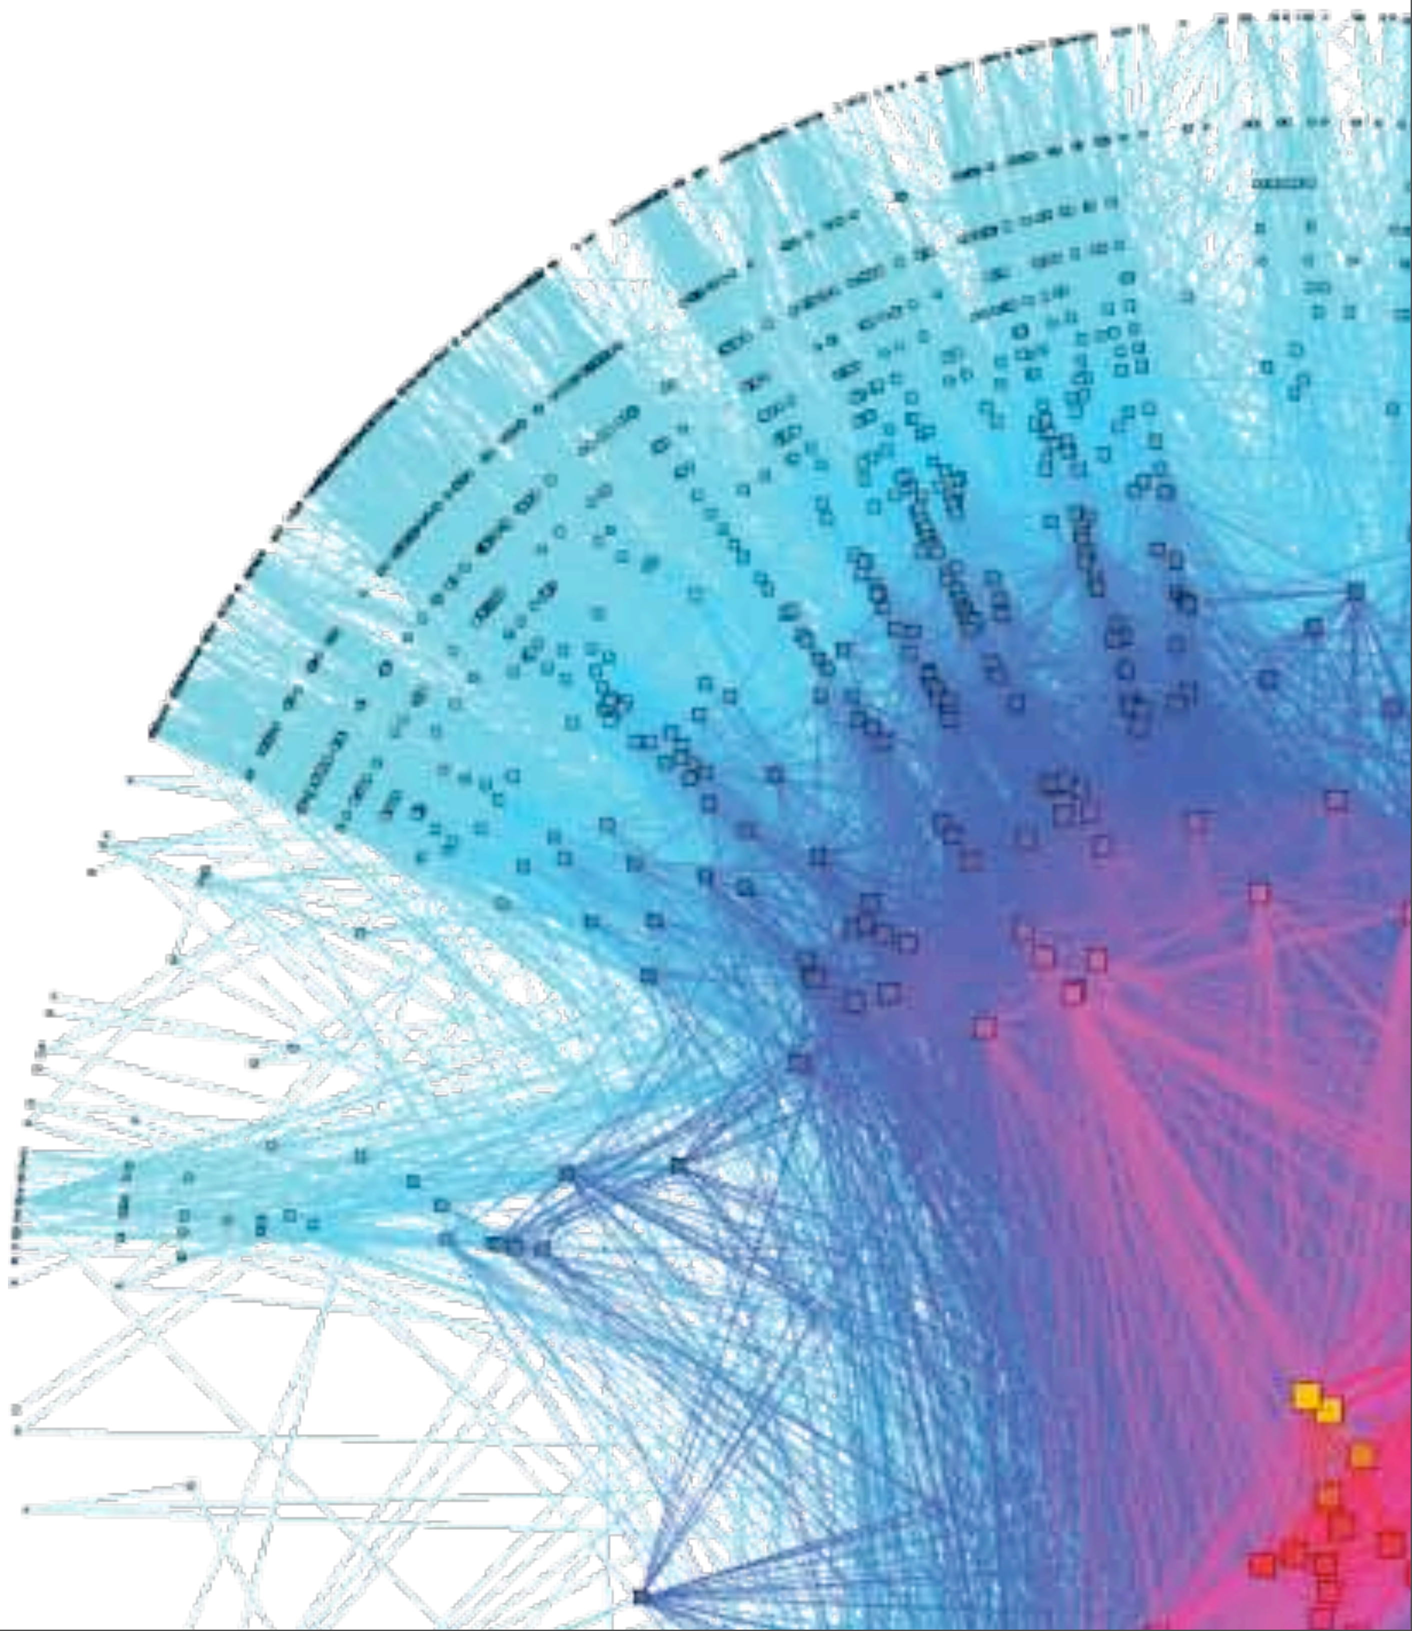

how we get the data

overview

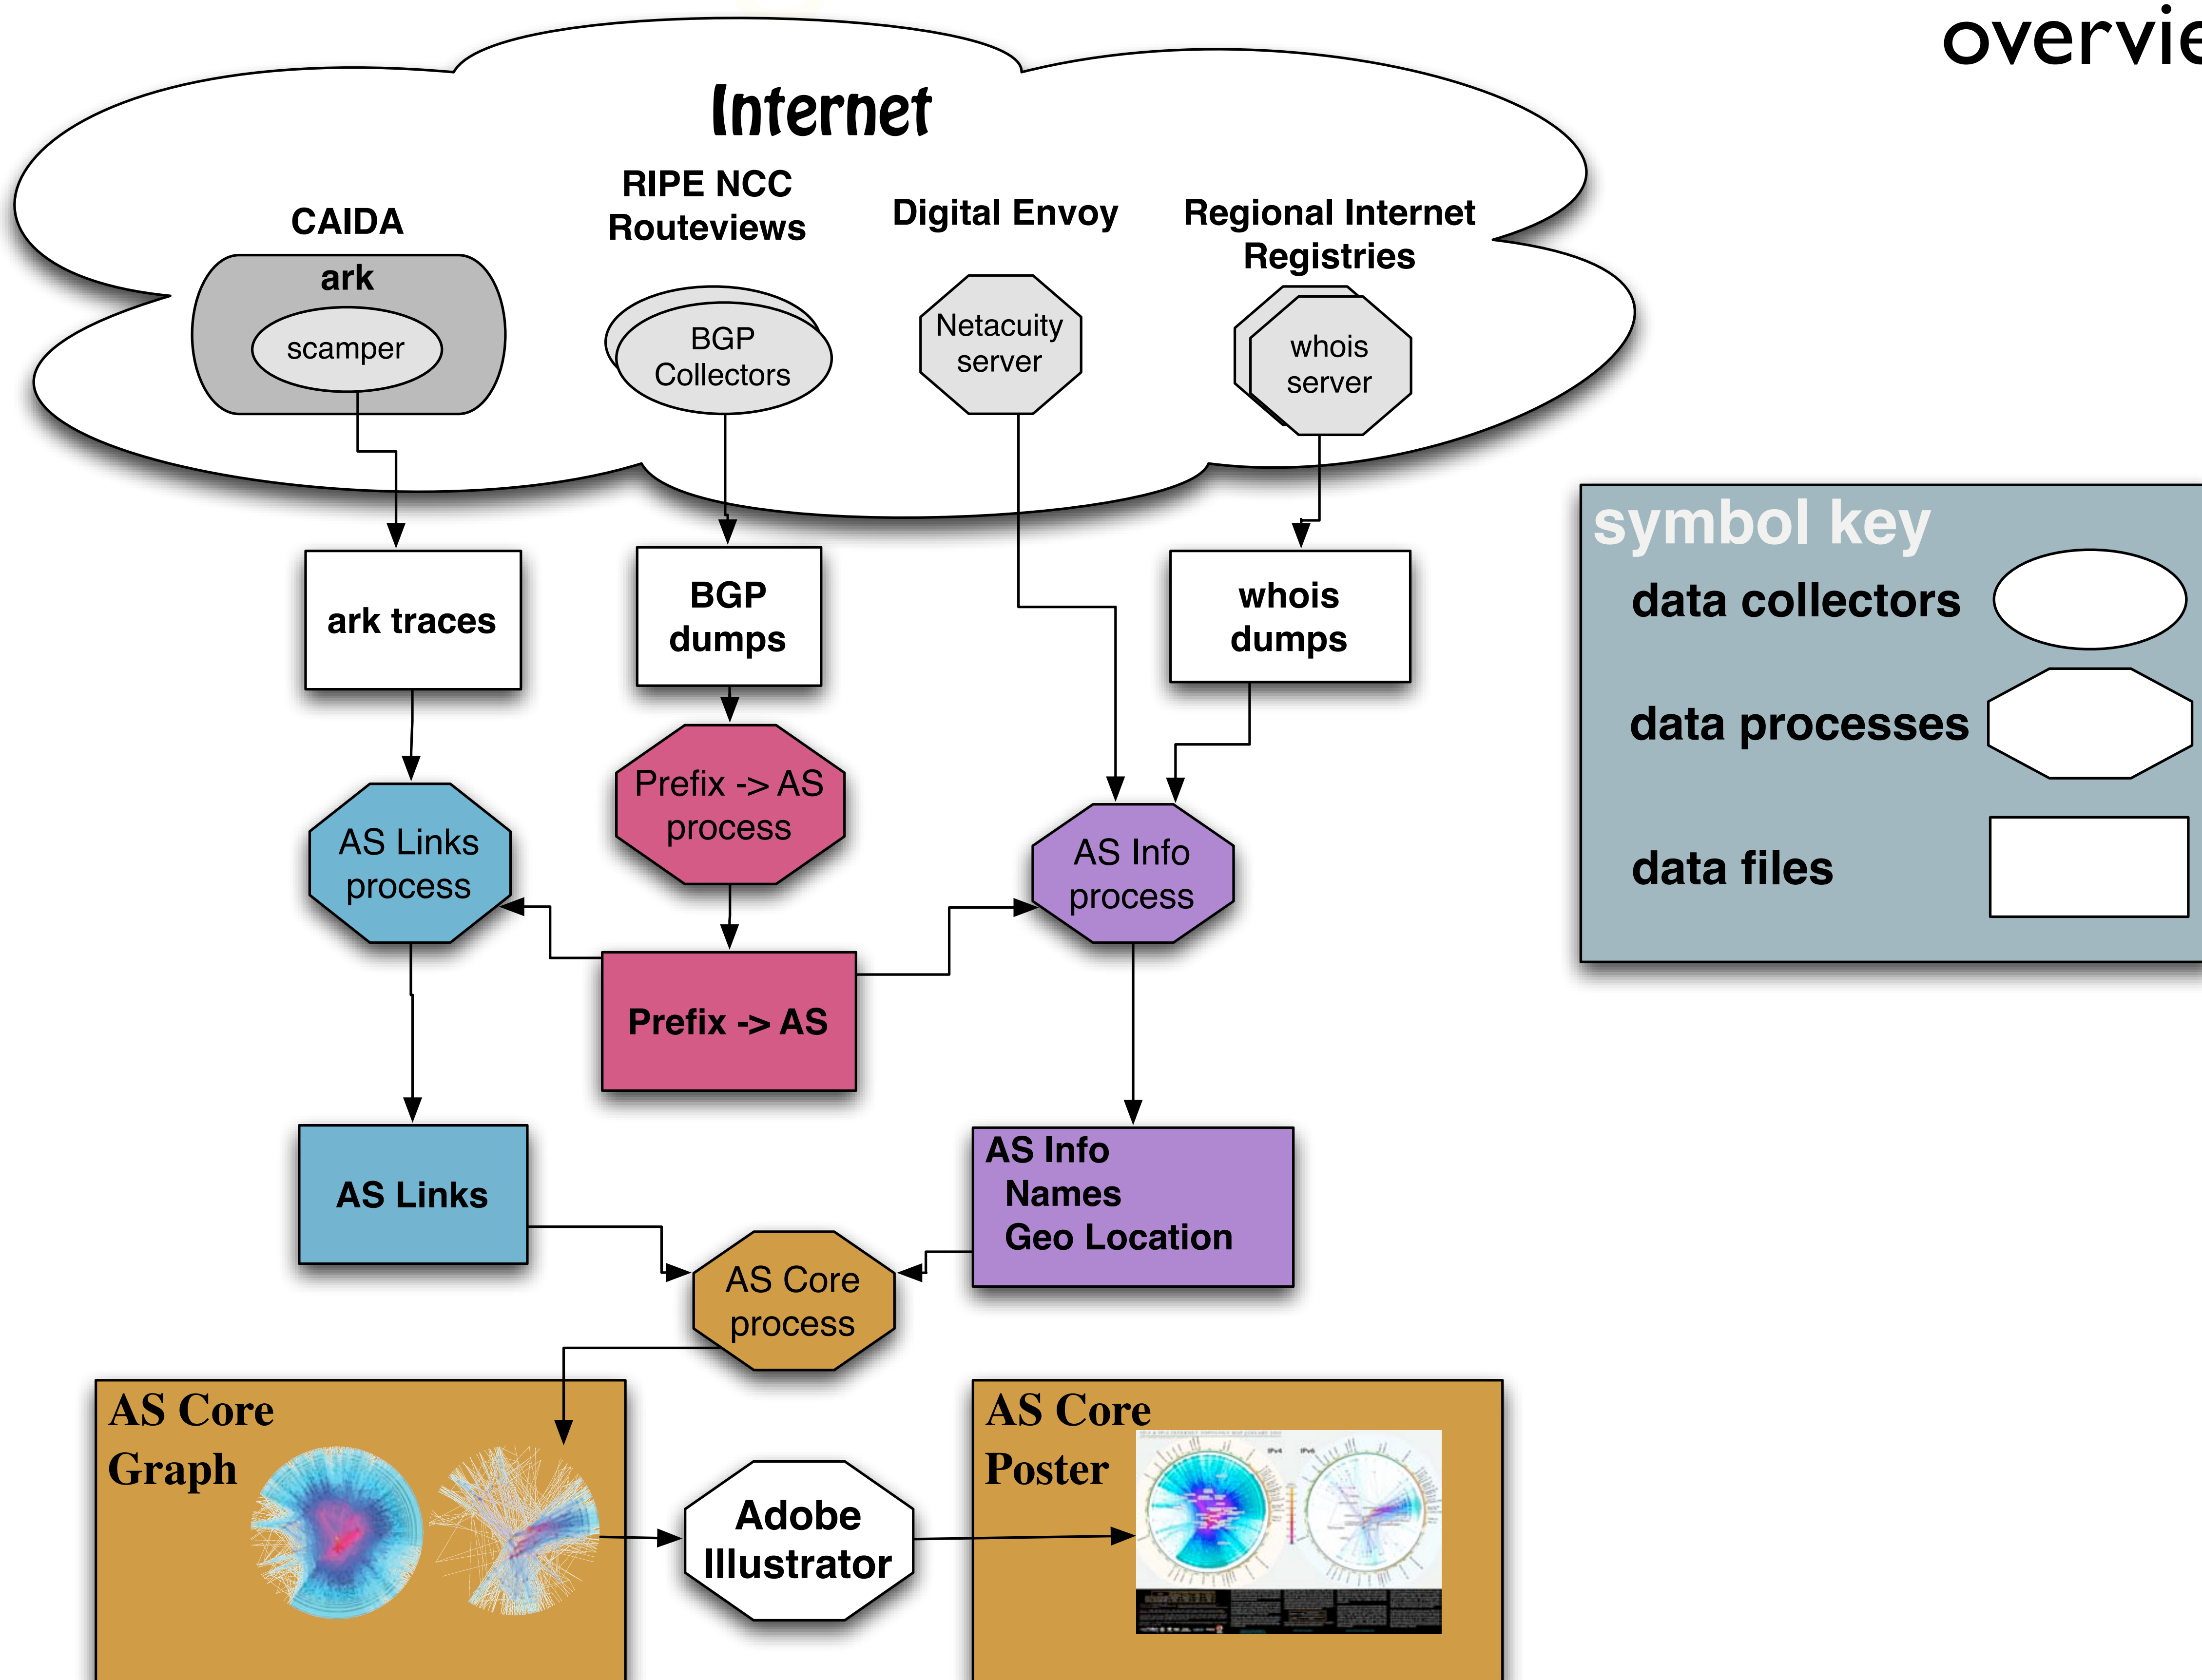

data sources

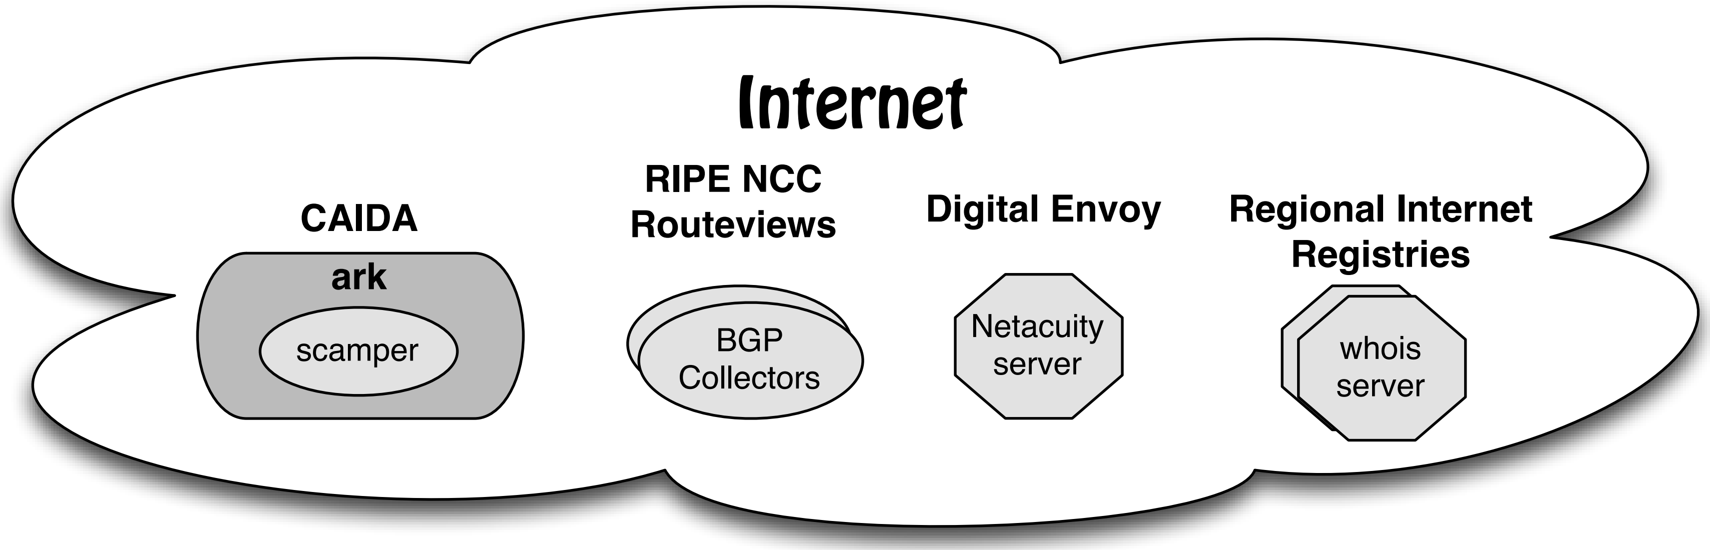

- **Archipelago** (ark)
  - platform that continually collects traceroute (topology) measurements
- **BGP collectors**
  - collects inter-domain (Border Gateway Protocol) routing tables and updates
- **Netacuity**
  - database of IP address geographic locations
- **WHOIS**
  - database(S) of registered users or assignees of Internet resources

Archipelago (ark)

data sources

- CAIDA's active measurement infrastructure
- 43 monitors - growing 1 or 2 per month
- 11 w/IPv6 connectivity
- Team-probing collecting IPv4 and IPv6 topology
- http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml

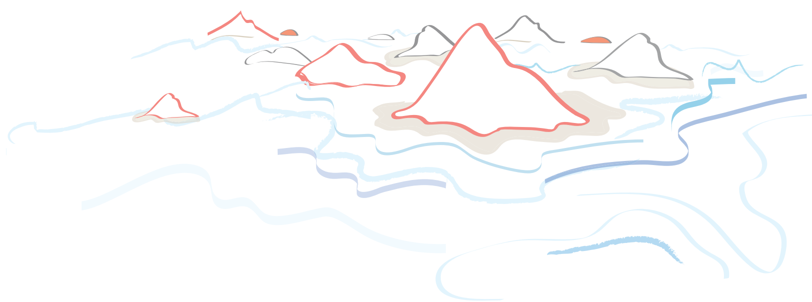

traceroute/topology data (not what is collected, but similar)

```
1  pinot-g1-0-0 (192.172.226.1)  0.856 ms  0.334 ms  0.374 ms
2  dolphin.sdsc.edu (198.17.46.17)  0.888 ms  0.461 ms  0.452 ms
3  dc-sdg-agg1--sdsc-1.cenic.net (137.164.23.129)  0.495 ms  0.486 ms  0.463 ms
4  dc-riv-core1--sdg-agg1-10ge.cenic.net (137.164.47.111)  3.462 ms  3.364 ms  3.215 ms
5  dc-lax-core1--riv-core1-10ge-2.cenic.net (137.164.46.57)  4.774 ms  4.815 ms  5.515 ms
6  dc-lax-peer1--lax-core1-ge.cenic.net (137.164.46.116)  12.970 ms  4.619 ms  4.560 ms
7  gi1-1--46.tr01-lsanca01.transitrail.net (137.164.131.245)  4.664 ms  4.655 ms  4.849 ms
```

hop | hostname | IP address | Round Trip Time (RTT)

from IP to AS Graph

summary

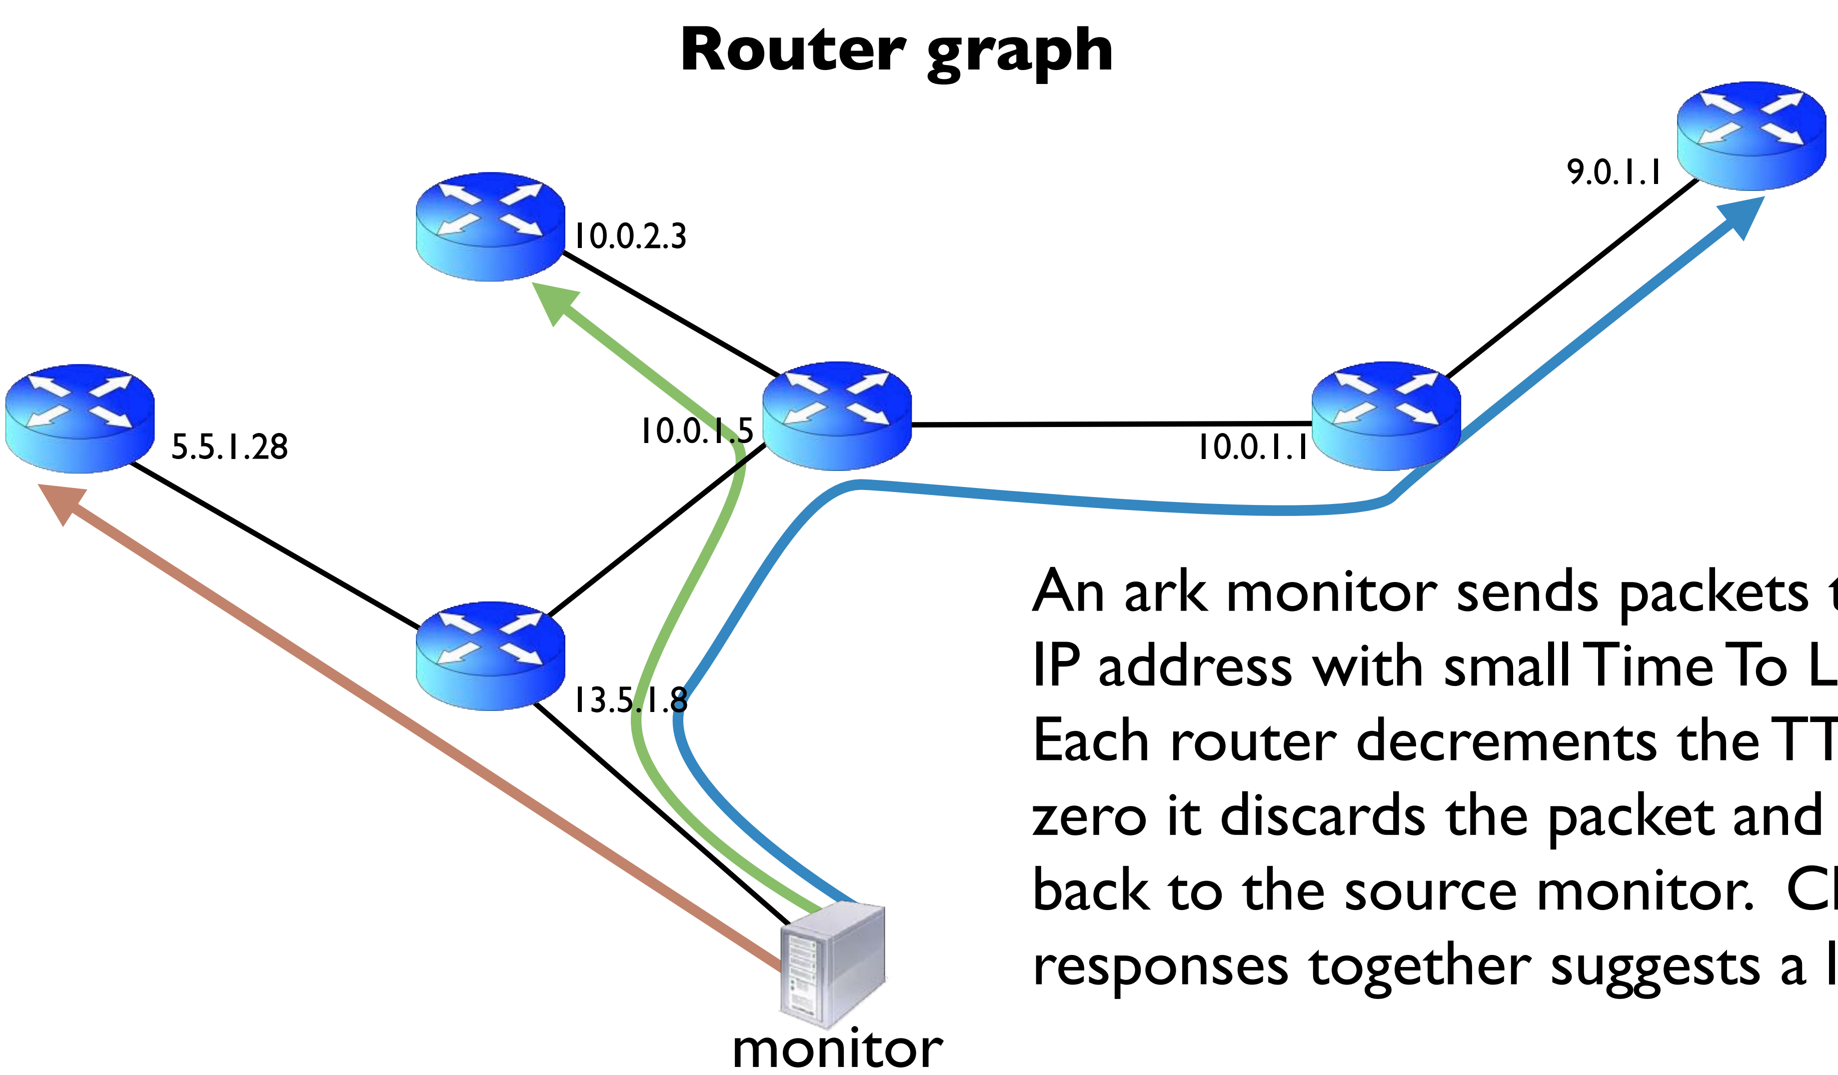

An ark monitor sends packets toward a destination IP address with small Time To Live (TTL) values. Each router decrements the TTL. When it reaches zero it discards the packet and sends a notification back to the source monitor. Chaining these responses together suggests a likely forward path.

	hop 1	hop 2	hop 3	hop 4	destination
trace 1	13.5.1.8	5.5.1.28			5.5.1.28
trace 2	13.5.1.8	10.0.1.5	10.0.2.3		10.0.2.3
trace 3	13.5.1.8	10.0.1.5	10.0.1.1	9.0.1.1	9.0.1.1

BGP Collectors

data sources

Collecting and sharing global routing [Border Gateway Protocol (BGP)] data:

- University of Oregon
  - 6 collectors
  - http://www.routeviews.org
- RIPE NCC (Regional Internet Registry for Europe/Middle East)
  - 13 collectors
  - http://www.ripe.net/data-tools/stats/ris/ris-raw-data
- used to map IP addresses to ASes

origin AS

BGP dump

```
TABLE_DUMP2 | 127  1 649600 | B | 157.130.10.233 |   701  | 4.21.103.0/24 |  70 1 3549 46
TABLE_DUMP2 | 127  1 649600 | B | 203.62.252.186 |   122  | 4.21.103.0/24 | 122 1  4637 3549 46
TABLE_DUMP2 | 127  1 649600 | B | 12.0.1.63      |  7018  | 4.21.103.0/24 | 7018 3549 46133
                                  source IP        source AS   prefix          AS path
```

BGP Routes

summary

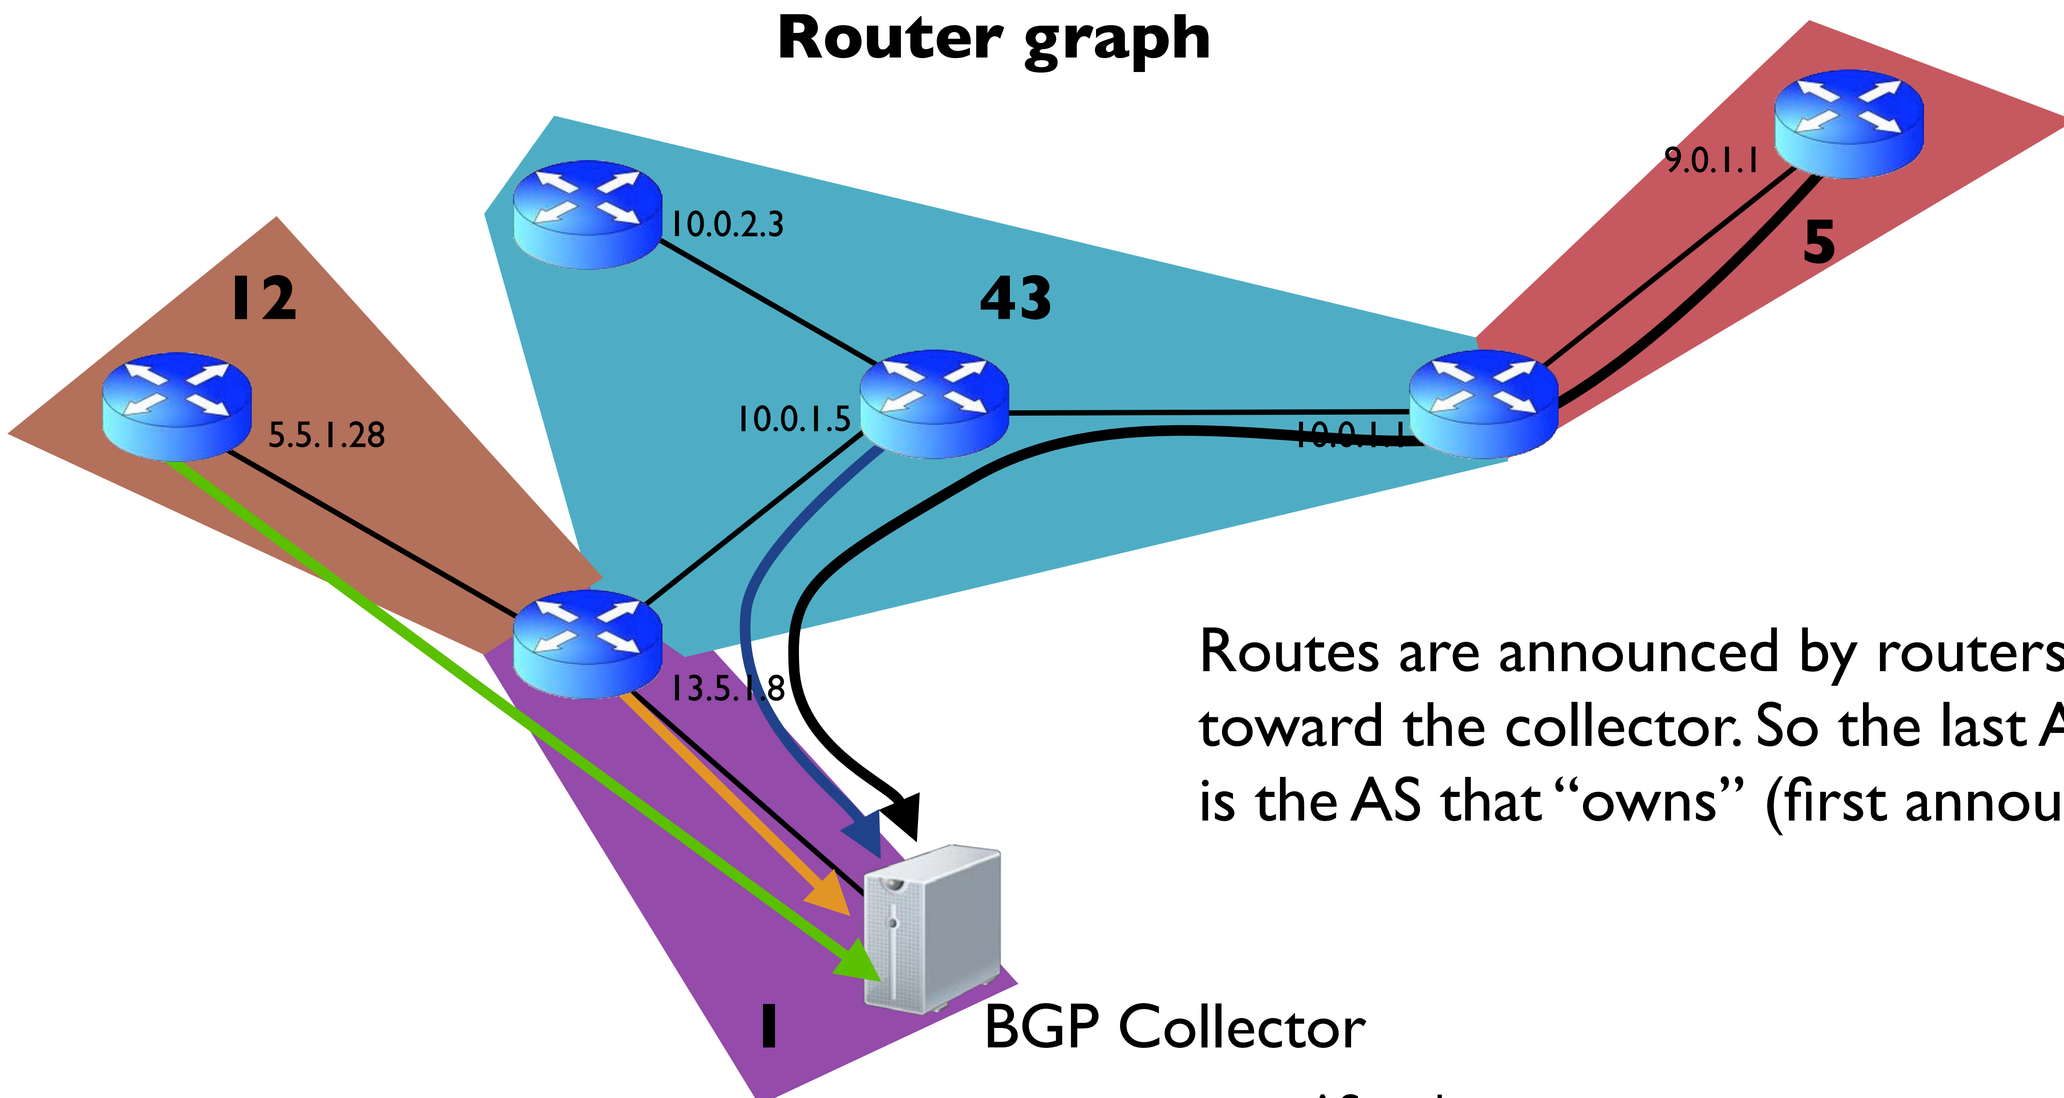

Routes are announced by routers and forwarded toward the collector. So the last AS, the "origin" AS, is the AS that "owns" (first announces) the prefix.

	prefix	AS path		origin AS
route 1	5.5.1.0/24		1	12
route 2	13.5.1.0/24			1
route 3	10.0.0.0/16		1	43
route 5	9.0.1.0/24	1	43	5

Netacuity

- Digital Envoy's commercial geolocation server
- Geolocation
  - identification of real-world geographic location of Internet identifiers
- MaxMind GeoLite is a free service
  - http://www.maxmind.com/app/geoip_country

Netacuity geographic dump

IP first	IP last	country	state	city	latitude	longtiude
192.172.226.0	192.127.226.255	usa	ca	la jolla	32.855	-117.249
137.164.23.0	137.164.23.255	usa	ca	tustin	33.736	-117.823
137.164.46.0	137.164.46.255	usa	ca	los angeles	33.973	-118.248
74.125.49.0	74.125.49.255	usa	il	chicago	41.886	-87.623

Whois

data sources

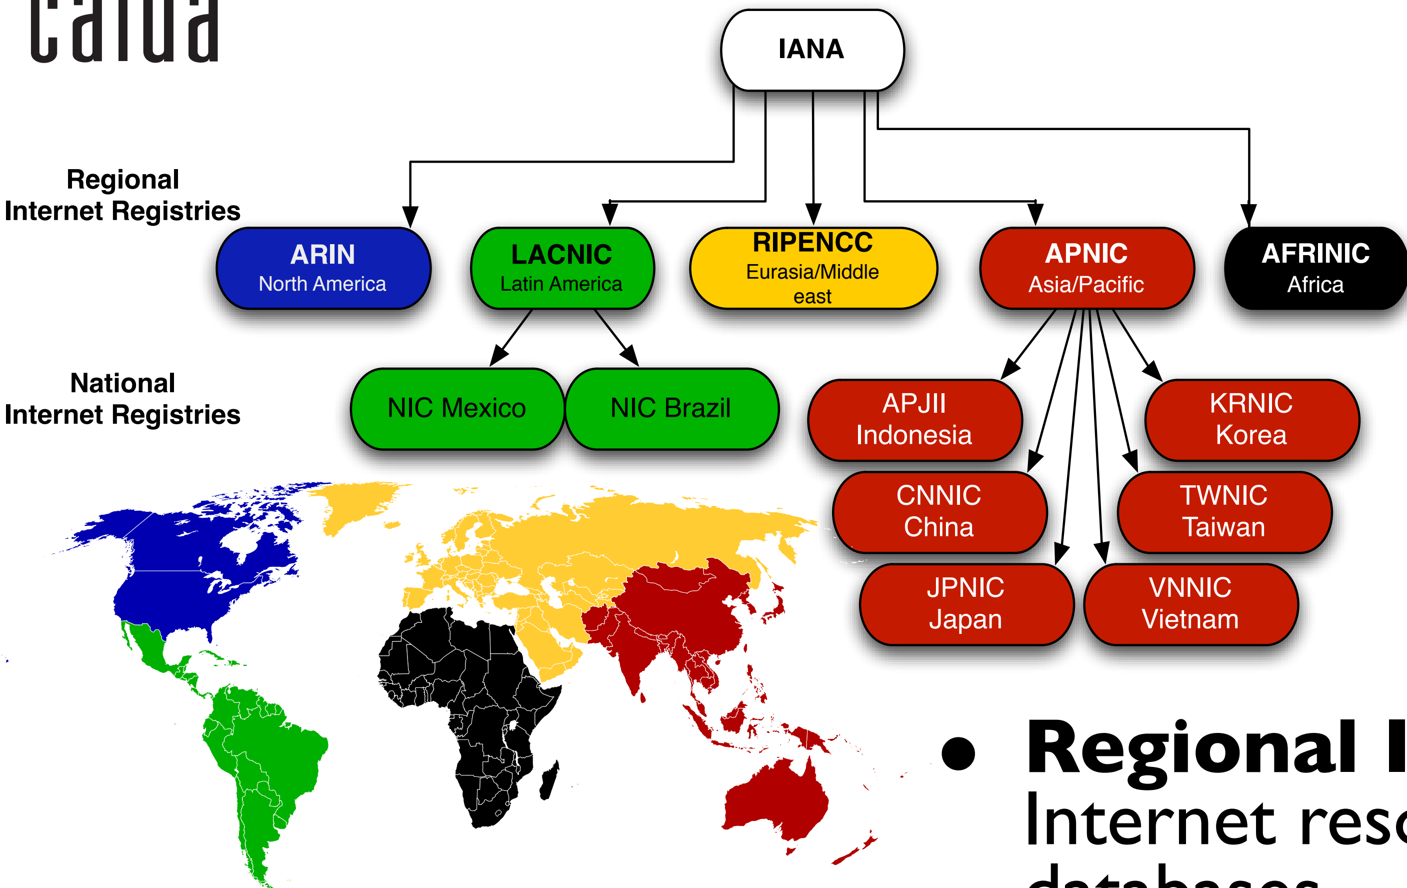

- **Regional Internet Registries** (RIRs) assign Internet resources and maintain the WHOIS databases.
- **WHOIS** databases store information about Internet registered users or assignees.

whois dump

```
ASNumber:        1909
OrgId:           SDSC

OrgId:           SDSC
OrgName:         San Diego Supercomputer Center
Address:         9500 Gilman Drive
```

whois

- *whois* command tools
  - whois is a command line client used to access the RIR servers
    whois -h whois.<RIR>.net <resource>
    <RIR> - afrinic, apnic, arin, lacnic, ripe,
    <resource> 129.10.1.1, AS12
  - start with ARIN, unless you know which region the allocation is in.

```
 > whois -h whois.arin.net AS43

ASNumber:       43
ASName:         BNL-AS
ASHandle:       AS43
RegDate:        1985-04-11
Updated:        2003-07-24
Ref:            http://whois.arin.net/rest/asn/AS43

OrgName:        Brookhaven National Laboratory
OrgId:          BNL
Address:        61 Brookhaven Ave
Address:        Bldg. 515
City:           Upton
StateProv:      NY
PostalCode:     11973
Country:        US
RegDate:        1984-09-13
Updated:        2007-02-01
Comment:        Brookhaven National Laboratory
Ref:            http://whois.arin.net/rest/org/BNL

OrgTechHandle: JB3159-ARIN
OrgTechName:   Bigrow, John
```

building AS paths

data process

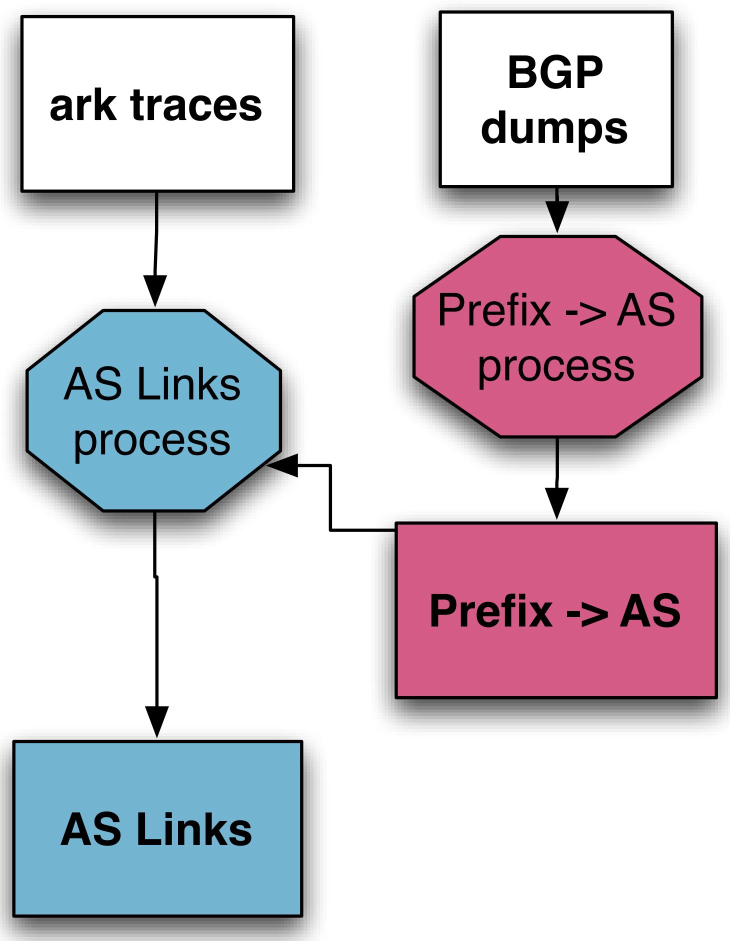

We take the IP-level topology generated by ark and convert it to a AS-level topology.

We first map the IP address to the AS announcing the address space that contains it.

IP Paths to AS Paths

data process

	prefix	AS path		origin AS
route 1	5.5.1.0/24		1	12
route 2	10.0.0.0/16			1
route 3	13.5.1.0/24		1	43
route 5	9.0.1.0/24	1	43	5

Map the IP address to the longest matching prefix and the those prefixes to their origin AS.

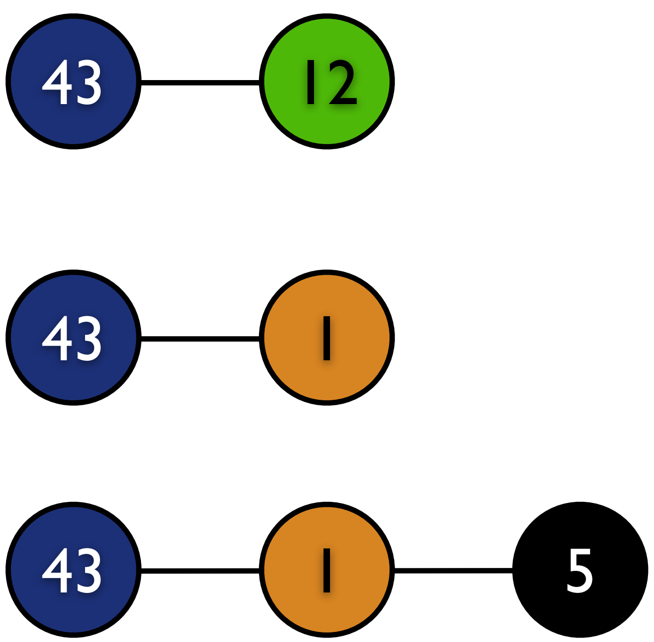

	hop 1	hop 2	hop 3	hop 4
trace 1	13.5.1.8	5.5.1.28		
prefix path	13.5.1.0/24	5.5.1.0/24		
AS path	**43**	**12**		

trace 2	13.5.1.8	10.0.1.5	10.0.2.3	
prefix path	13.5.1.0/24	10.0.0.0/16		
AS path	**43**	**1**		

trace 3	13.5.1.8	10.0.1.5	10.0.1.1	9.0.1.1
prefix path	13.5.1.0/24	10.0.0.0/16		9.0.1.0/24
AS path	**43**	**1**		**5**

Fill in neighbors

data process

paths

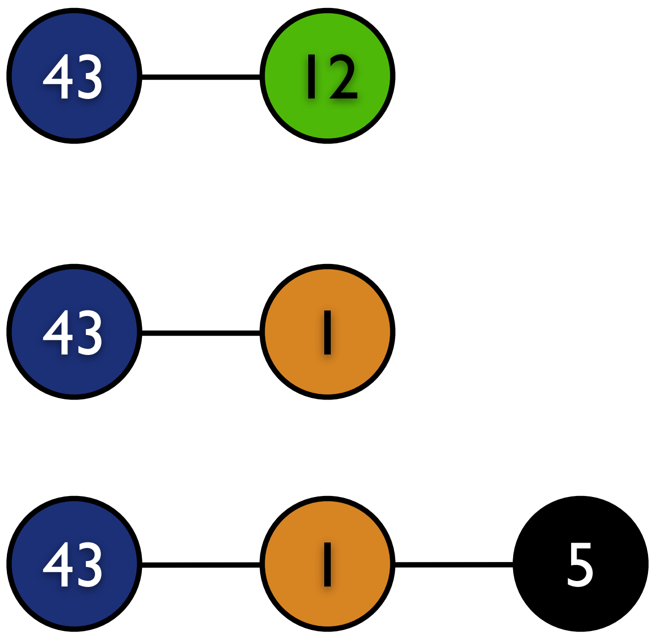

Graph

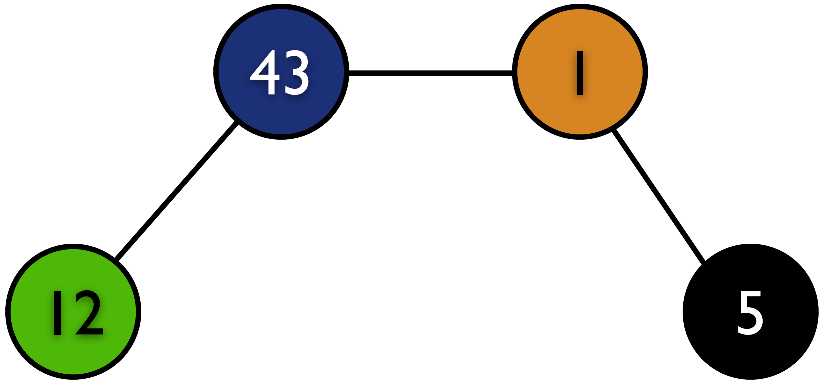

AS	AS's name	AS's longitude	AS's neighbors	degree
1			5, 43	2
5			1	1
12			43	1
43			1, 12	2

AS geography/ownership

data process

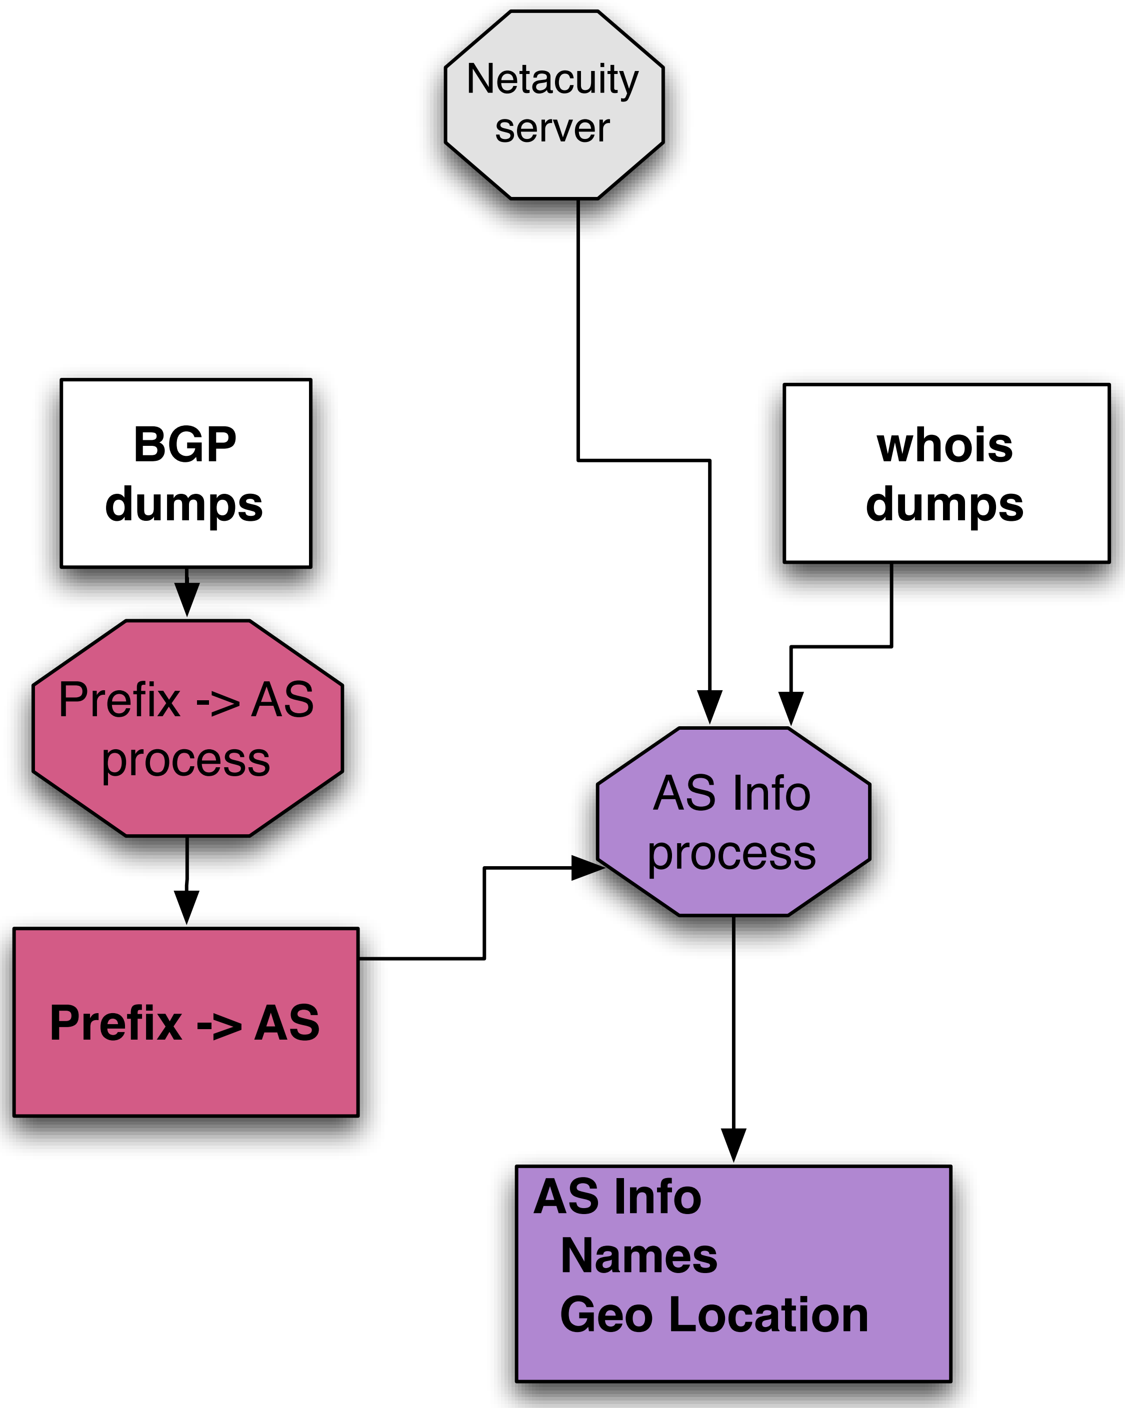

We take the organization name directly from the WHOIS dumps.

Geographic location will be harder, since our geolocation database does not provide locations for ASes, only IP addresses.

Geolocation to longitude

data process

We assign an AS's longitude to be equal to the weighted average of the Netacuity address blocks it announces.

$$\frac{\sum_i \text{block}_i.\text{longitude} * \text{block}_i.\text{size}}{\sum_i \text{block}_i.\text{size}}$$

origin AS	prefix	geolocation blocks: IP block	geolocation blocks: longitude	weighted average longitude
12	5.5.1.0/24	5.5.1.0 - 5.5.1.255	-103	-103
1	10.0.0.0/16	10.0.0.0-10.0.127.255	25	37.5
		10.0.128.0-10.0.255.255	50	
43	13.5.1.0/24	13.5.1.0-13.5.1.255	-23	-23
5	9.0.1.0/24	9.0.1.0-9.0.1.255	45	45

Bring it all together

data process

We now have everything we need to build the graph

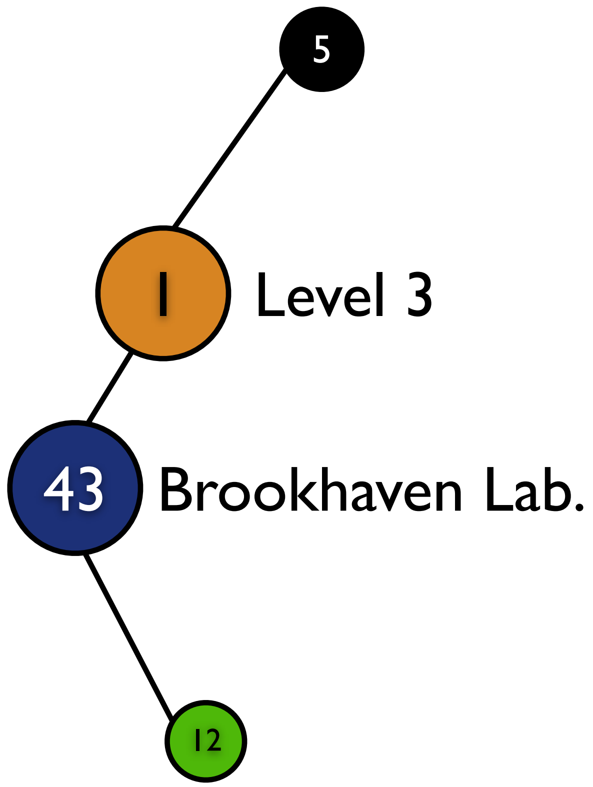

AS	AS's name	AS's longitude	AS's neighbors	degree
1	Level 3	37.5	5, 43	2
5	Symbolics, Inc.	45	1	1
12	New York University	-103	43	1
43	Brookhaven Laboratory	-23	1, 12	2

how is it drawn

visualization breakdown

Each node is a single AS, although ASes with nearby/same degree and longitude will overlap.

node's color/radius	$1 - \log\left(\frac{\text{degree (AS)} + 1}{\text{maxmium.degree} + 1}\right)$
node's size	$\frac{\text{degree (AS)} + 1}{\text{maxmium.degree} + 1}$
node's angle	longitude of the AS's BGP prefixes
link color	node's color with smallest degree

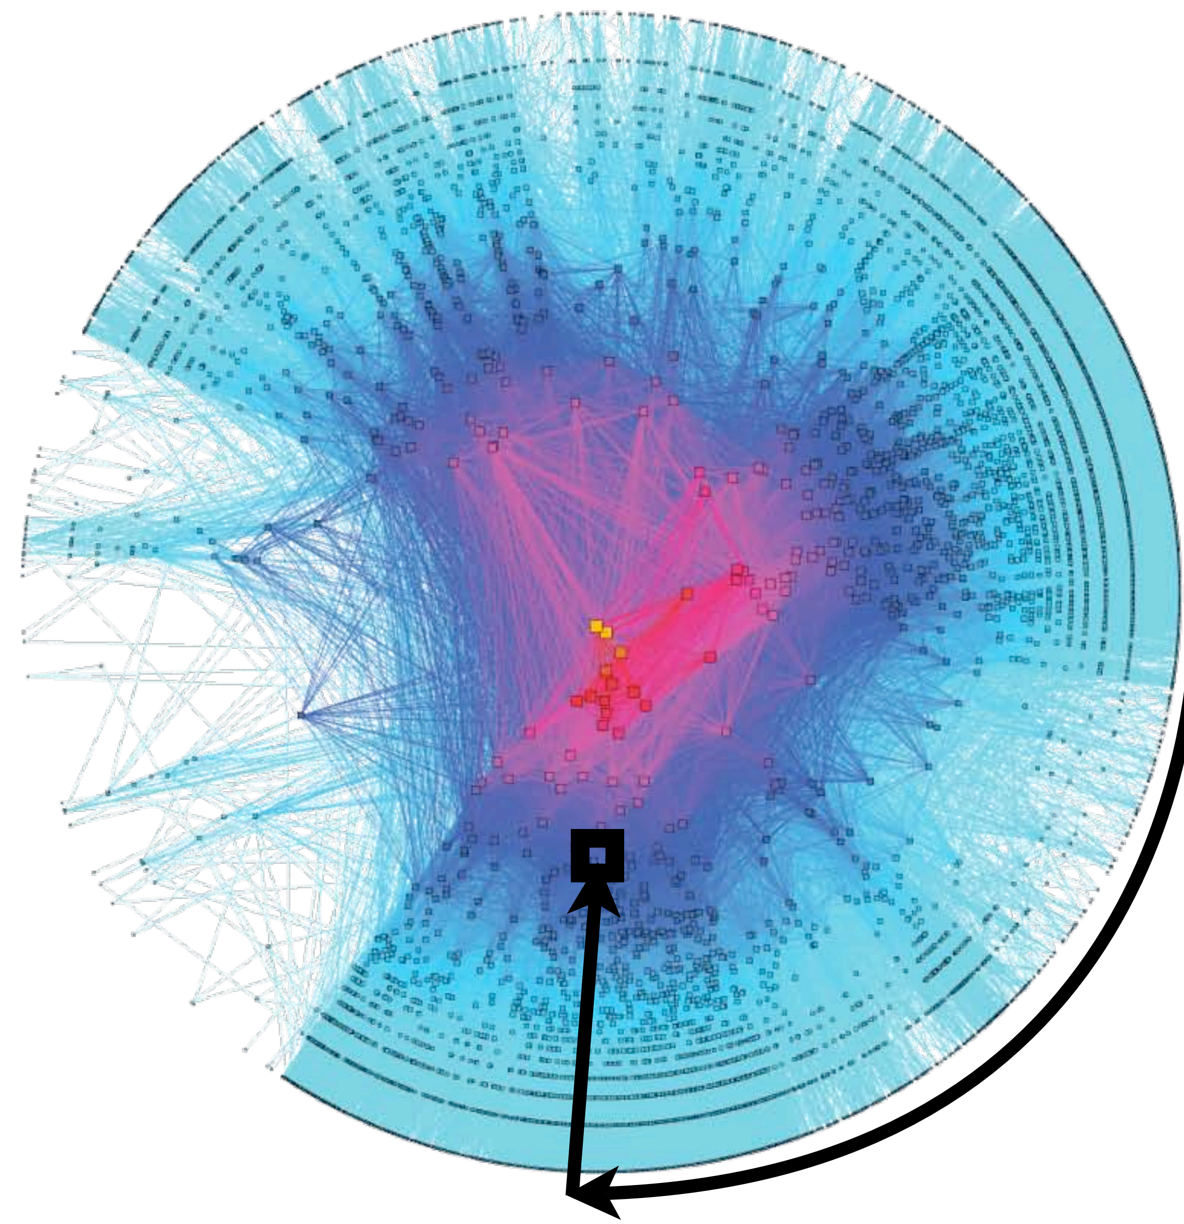

geographic regions

visualization breakdown

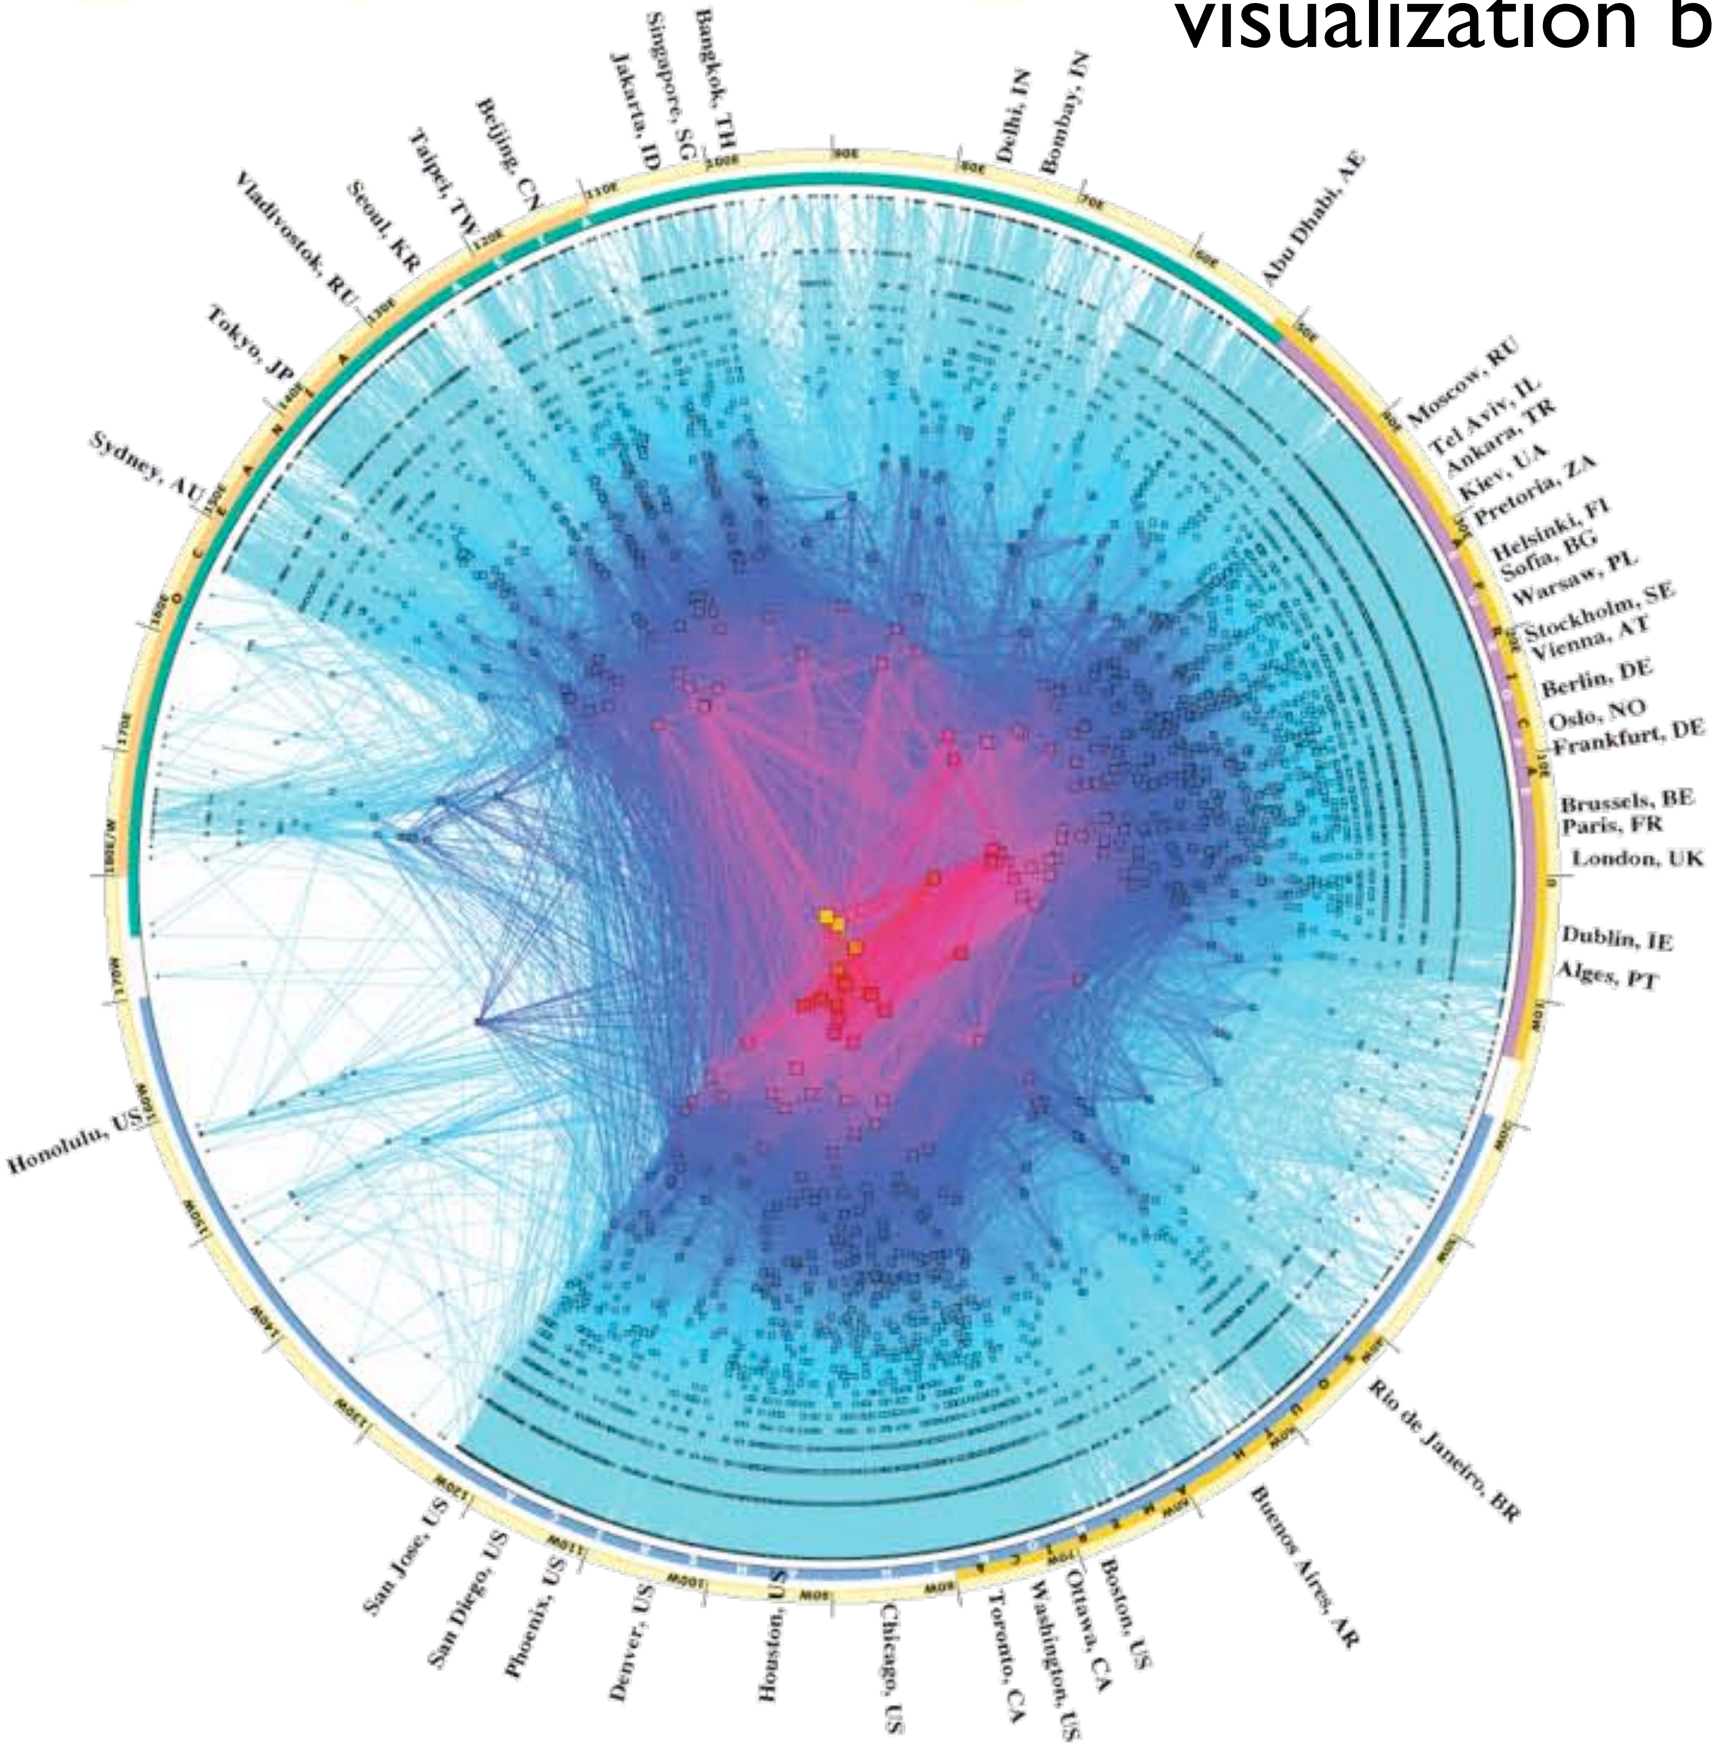

geographic regions

visualization breakdown

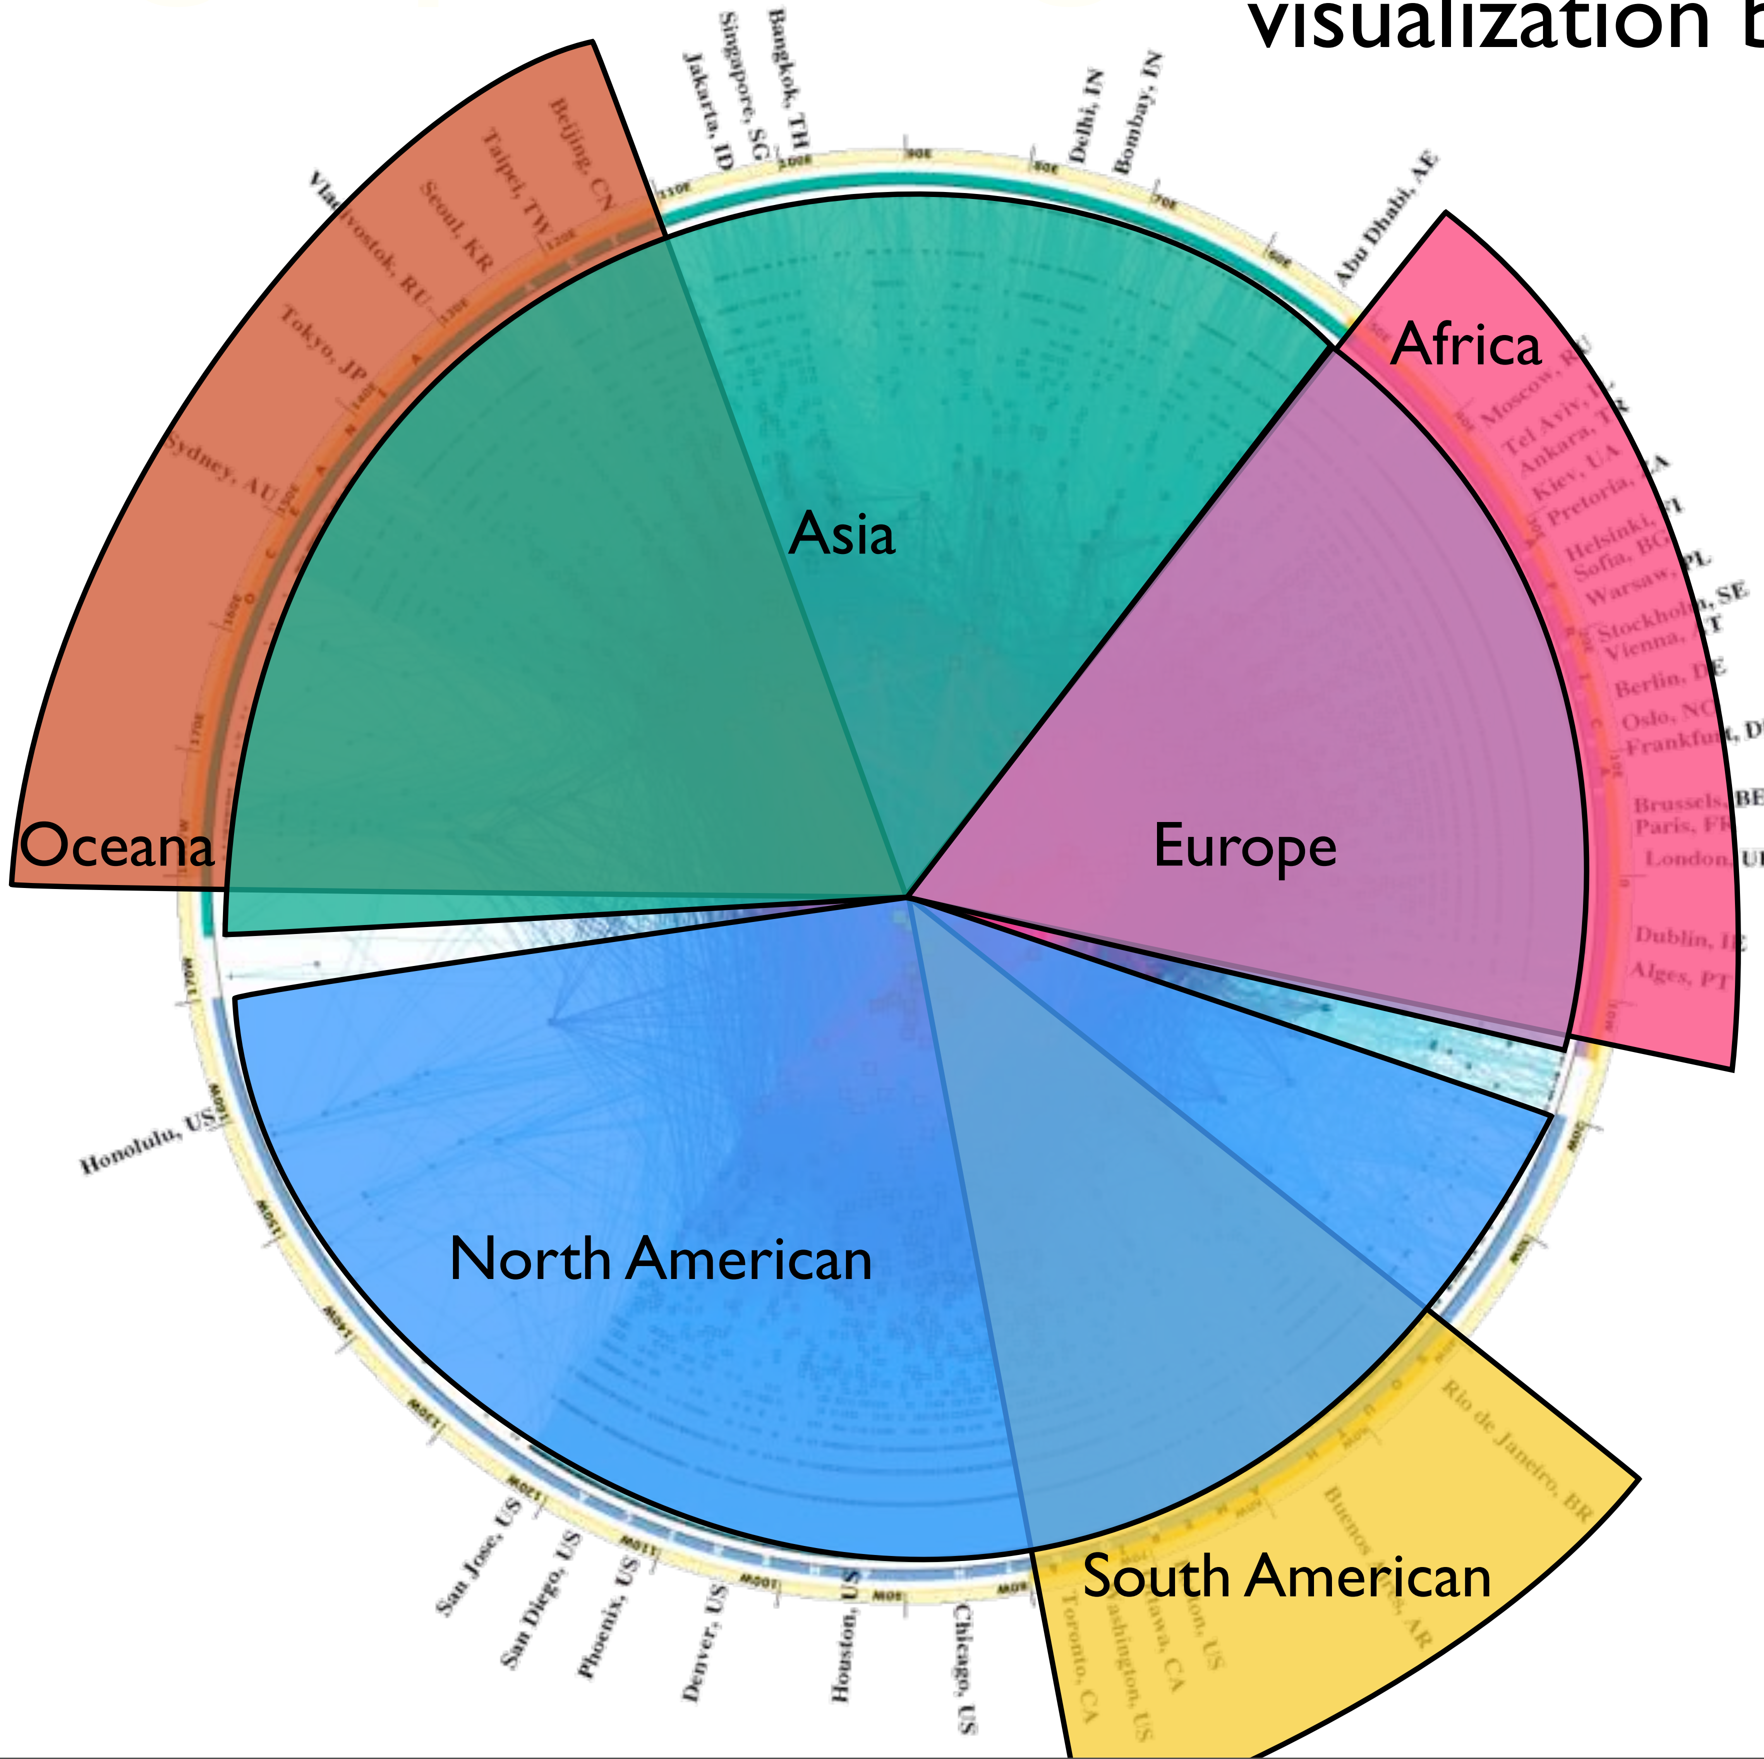

why IPv6?

IPv4 vs IPv6

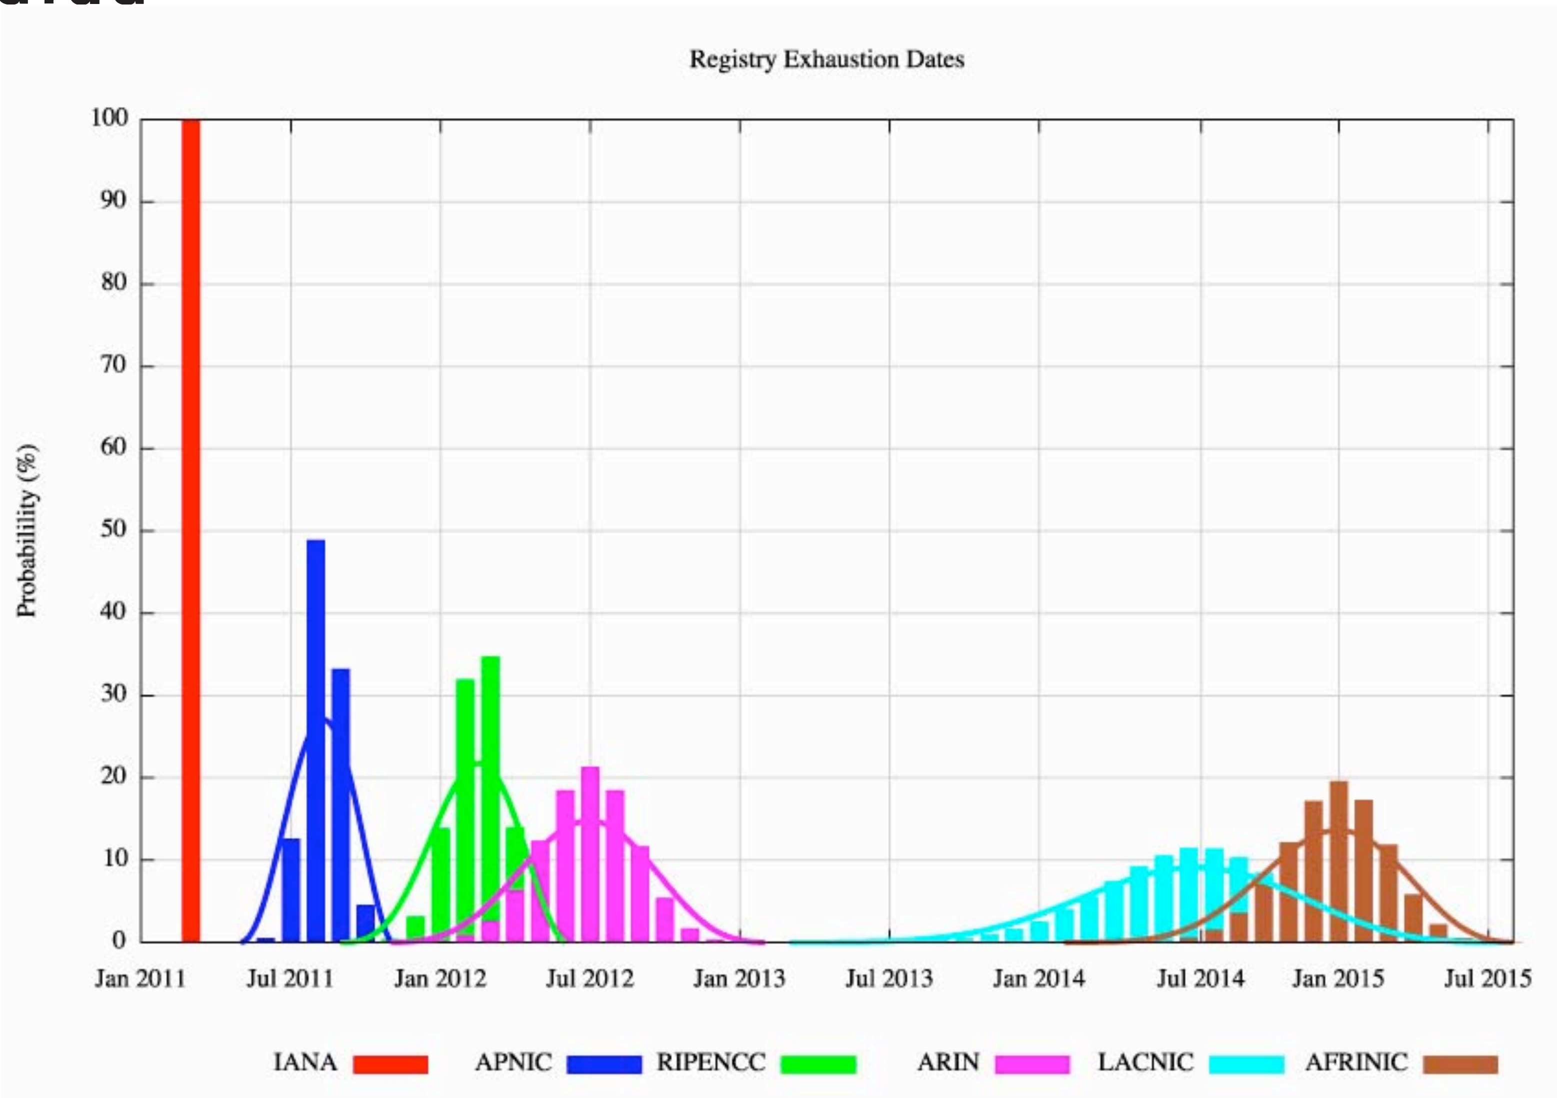

http://www.potaroo.net/ispcol/2010-10/when.html

Internet Assigned Number Authority (IANA) allocated its last /8 to the RIR on 31 January 2011

The RIRs are expected to run out of IPv4 address by no later then July 2015.

Future IANA allocations must come from IPv6 address space.

IPv4 vs IPv6 graphs

IPv4 vs IPv6

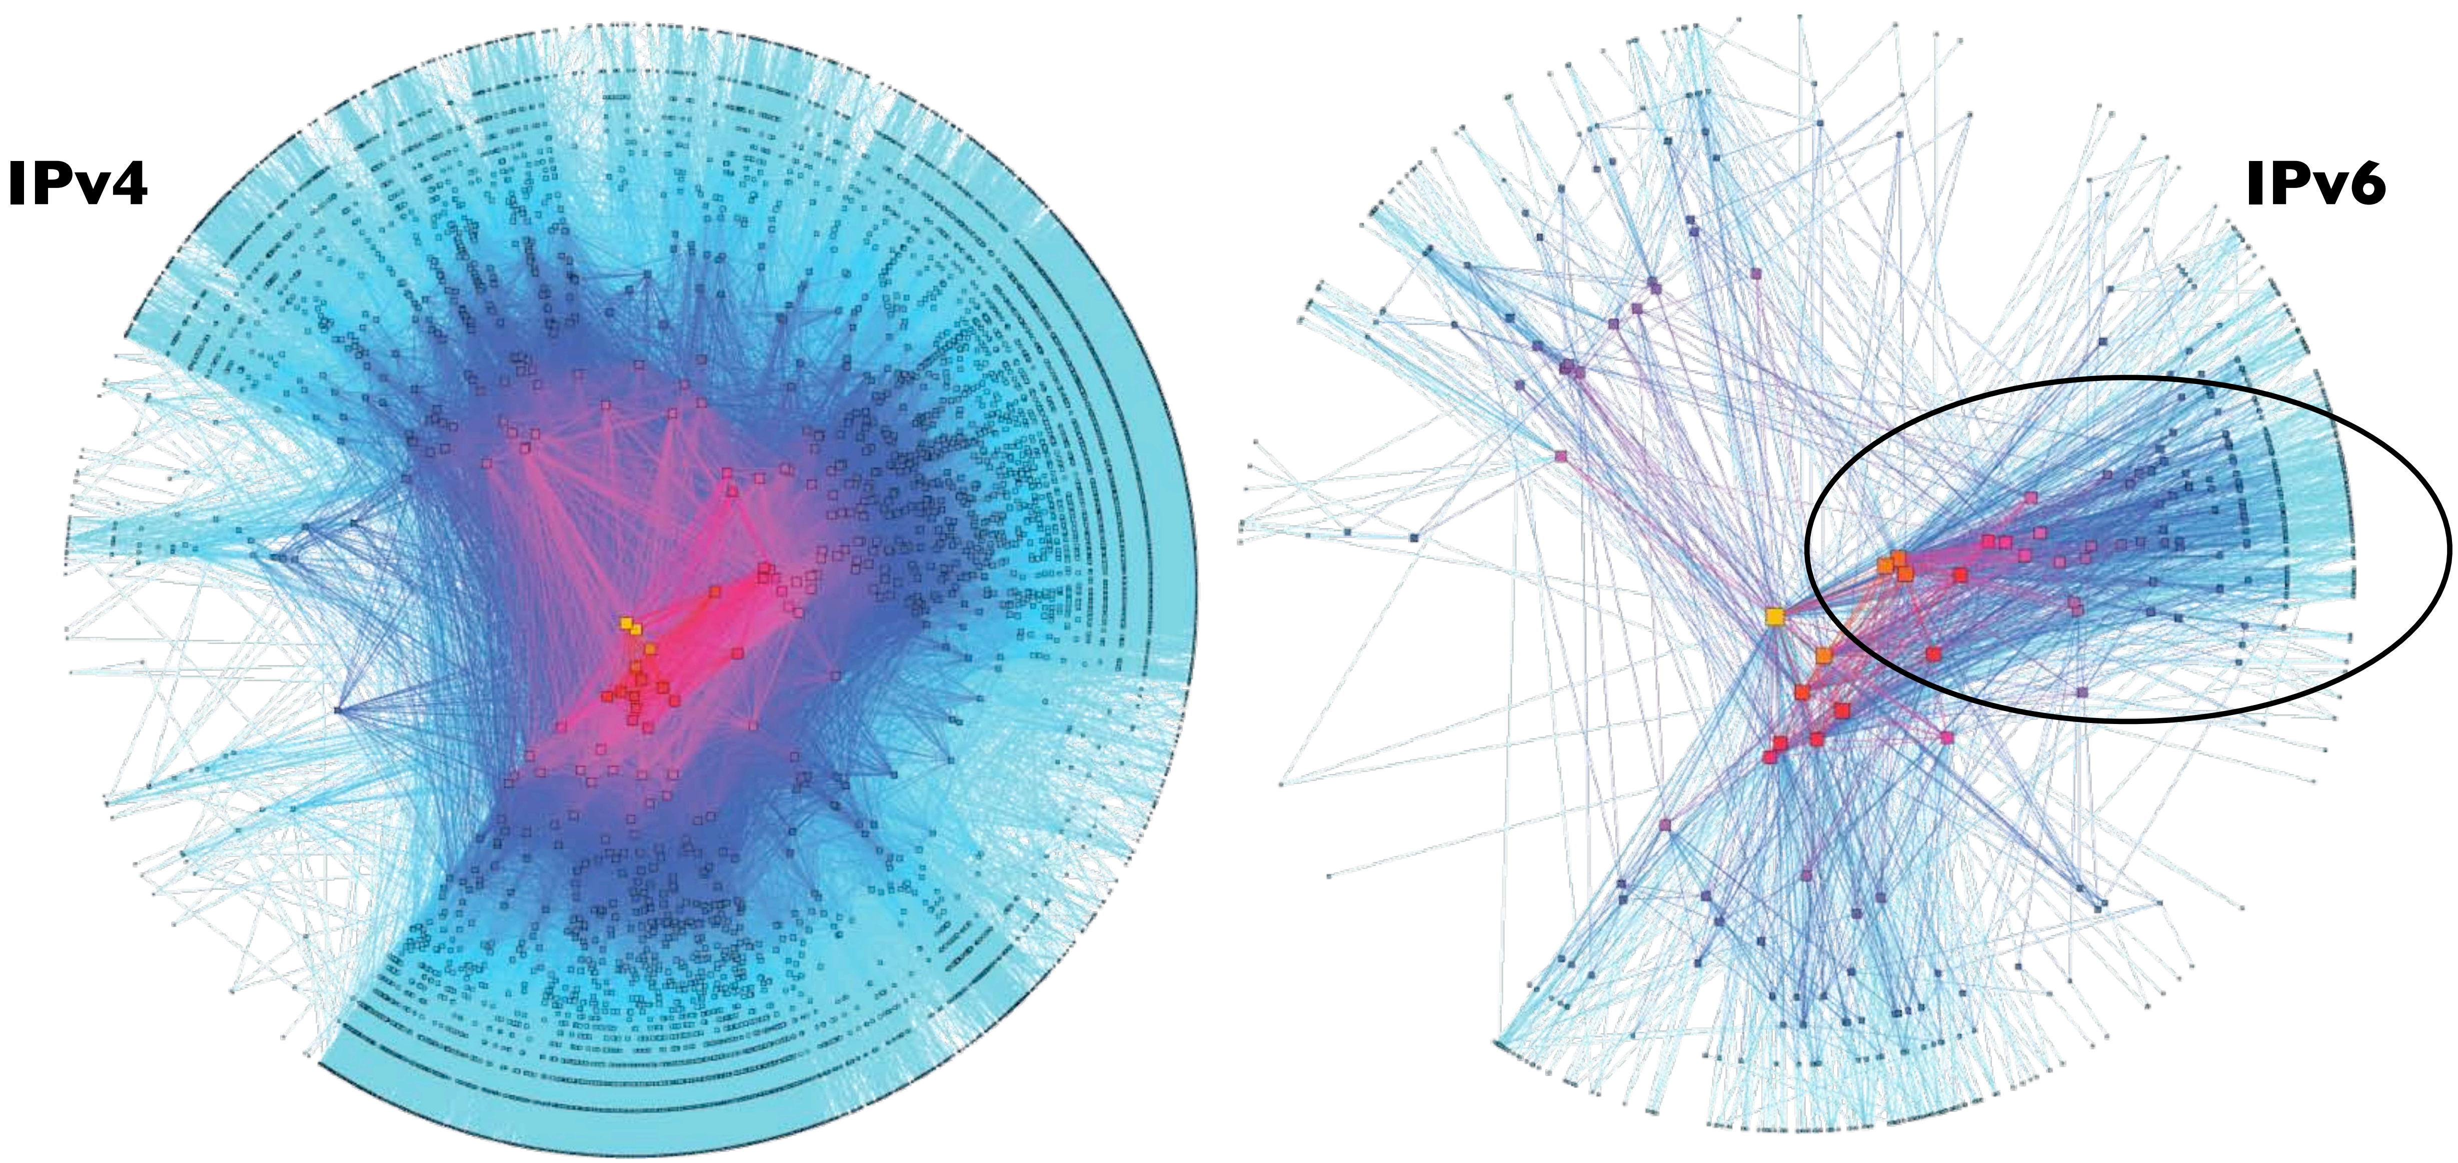

IPv4 vs IPv6 graphs

IPv4 vs IPv6

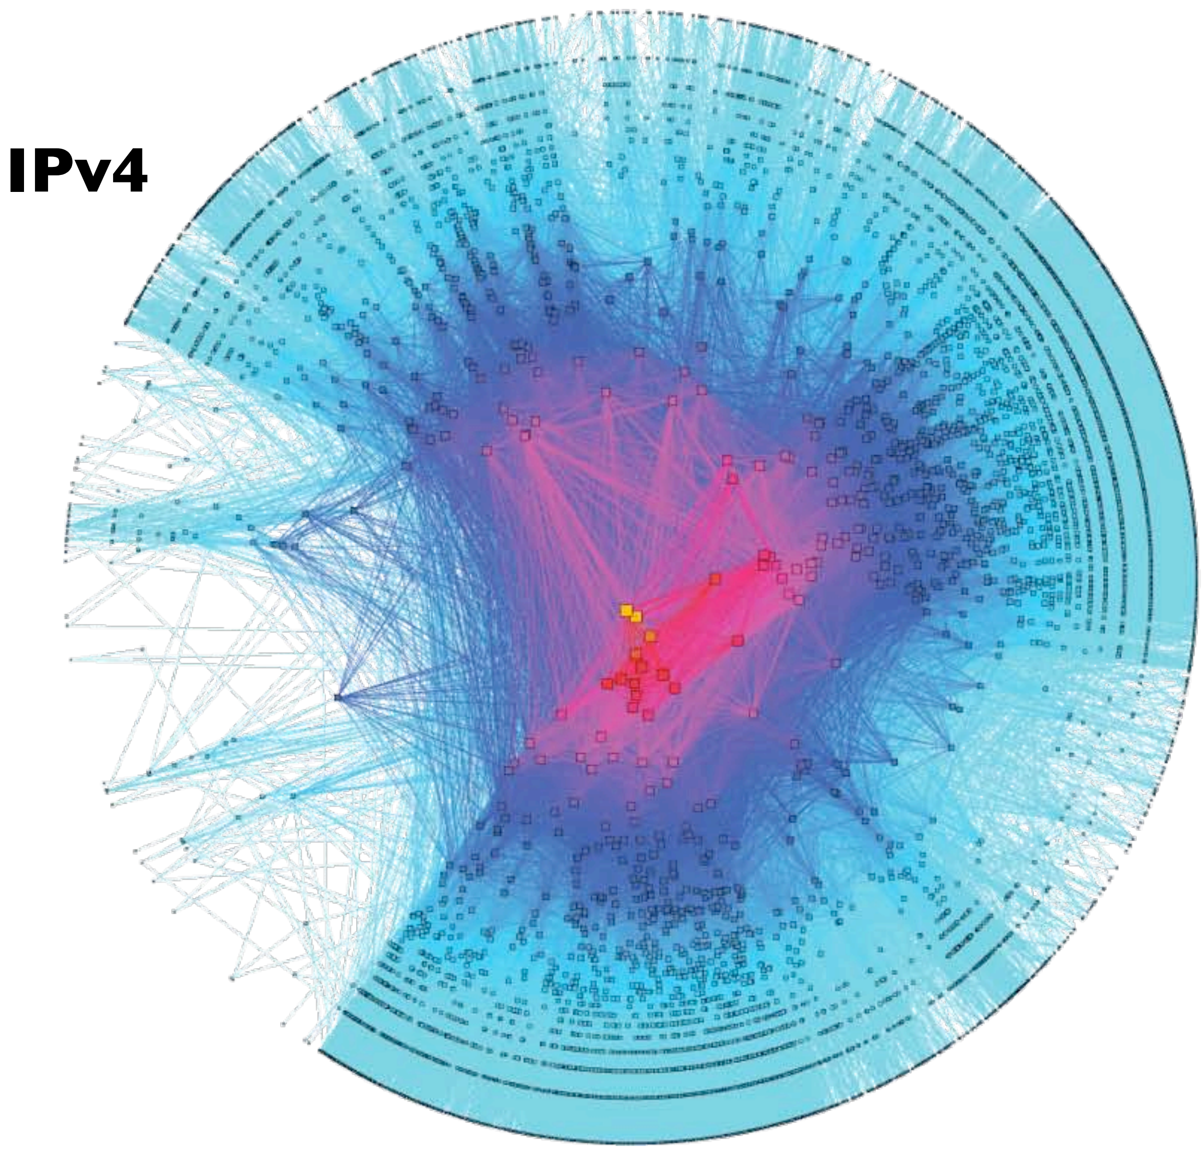

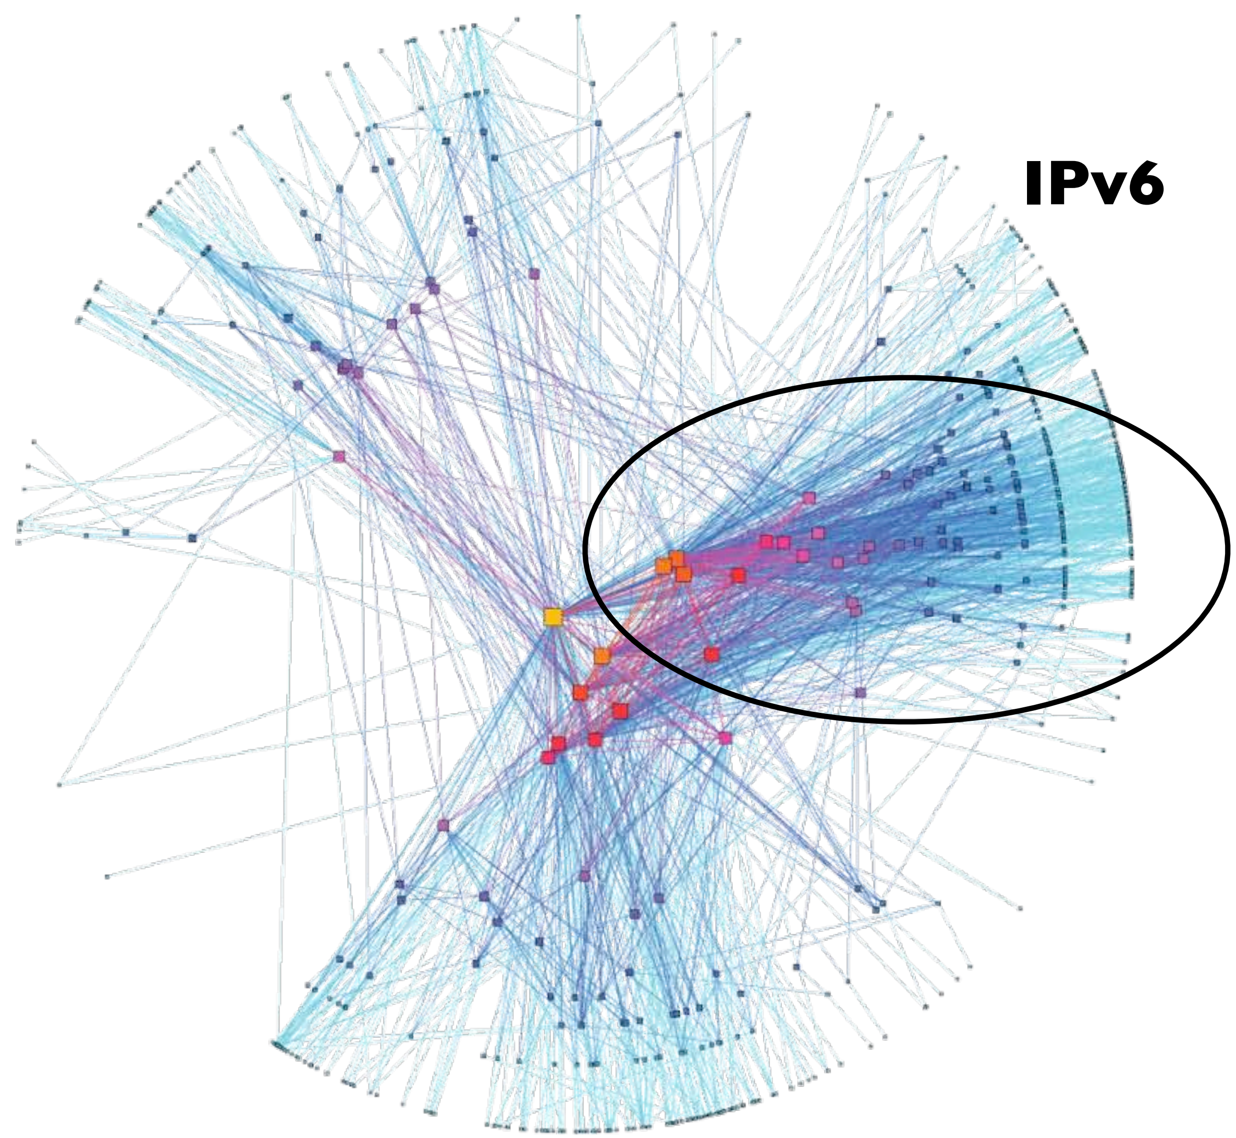

IPv6 highest area of density in Europe

IPv4 vs IPv6 graphs

IPv4 vs IPv6

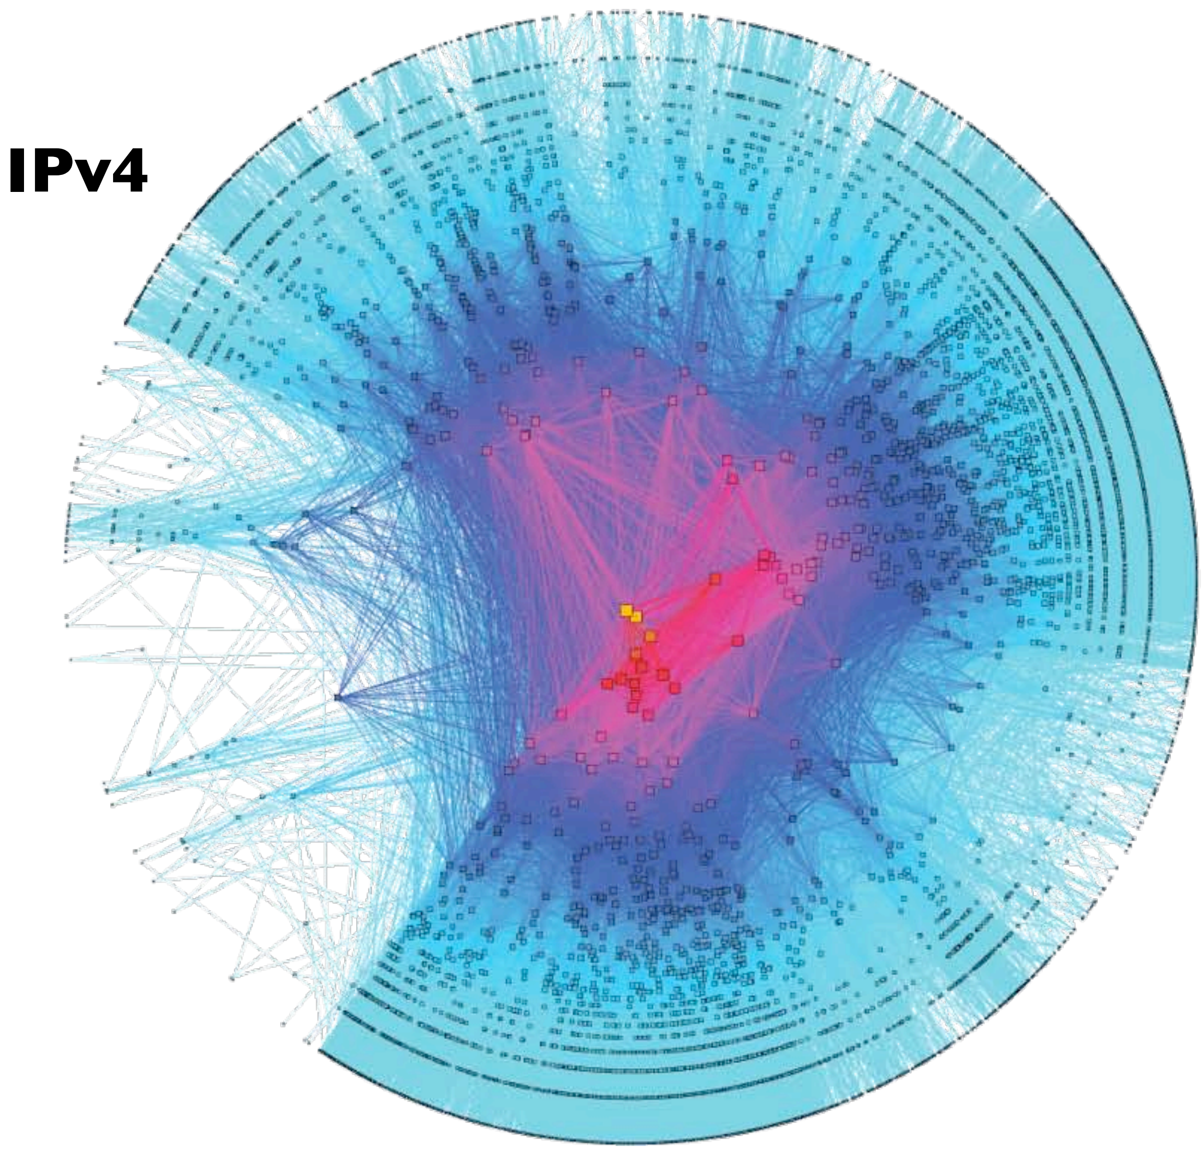

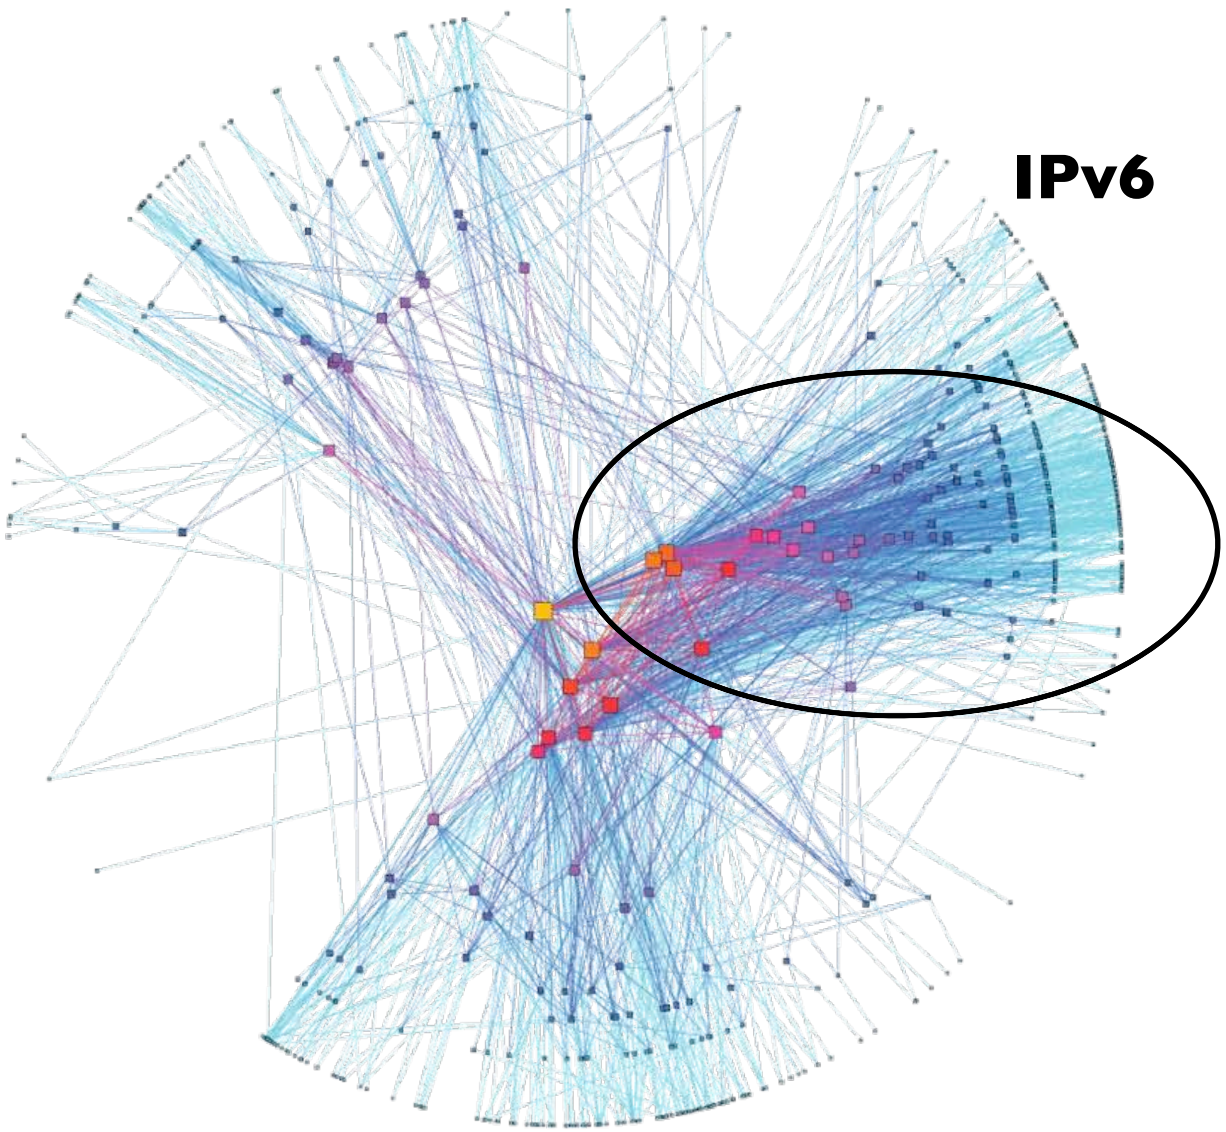

IPv4 high density in Asia, America, and Europe

IPv6 highest area of density in Europe

IPv4 vs IPv6 cores

IPv4 vs IPv6

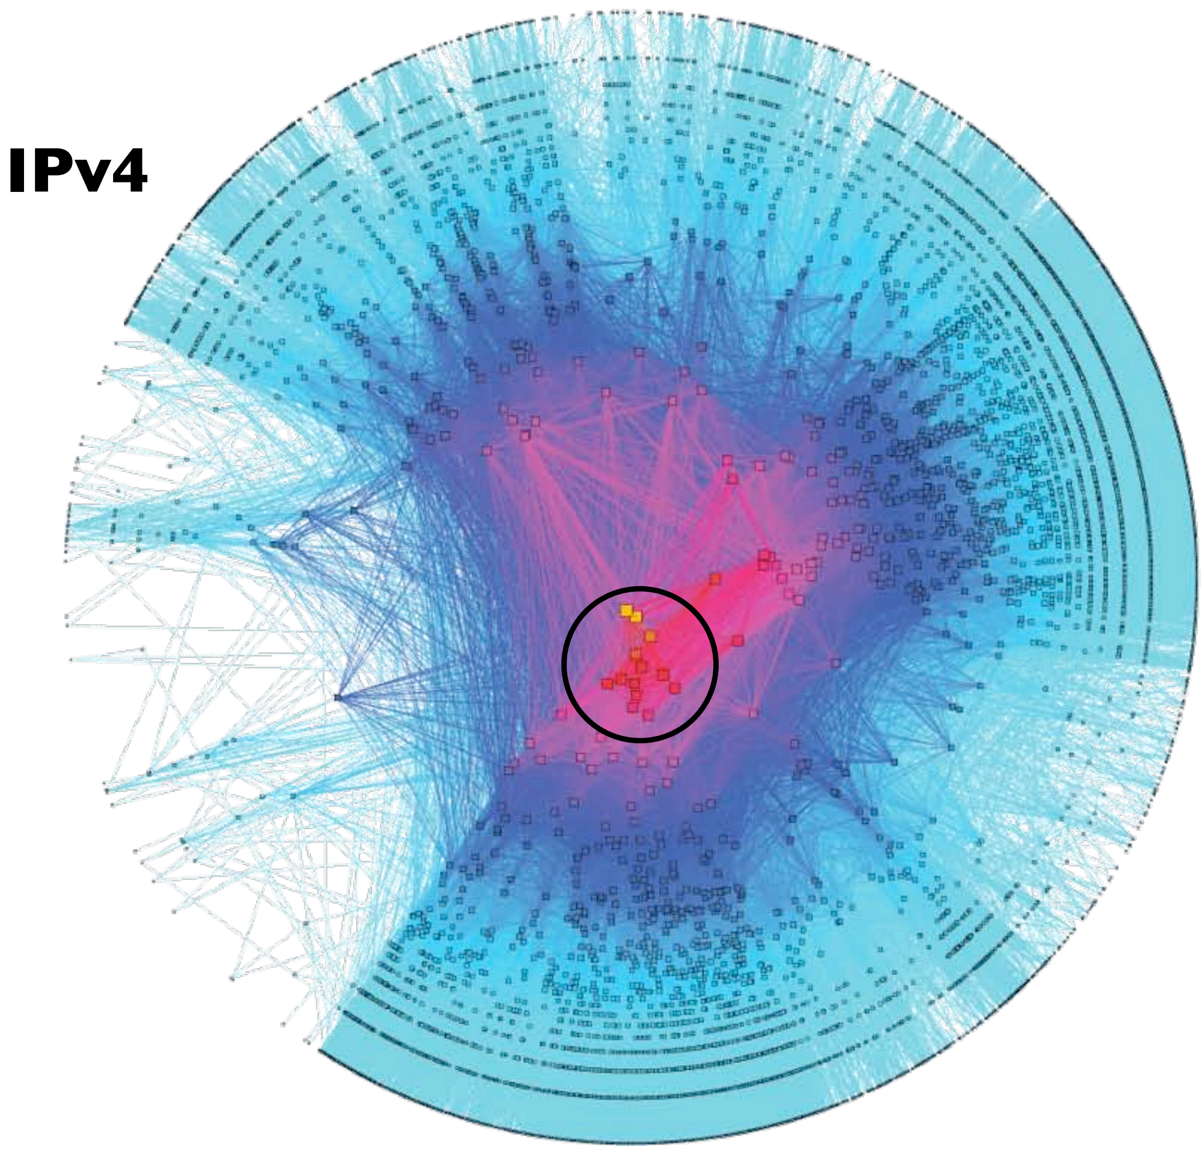

IPv4 core primarily in North America

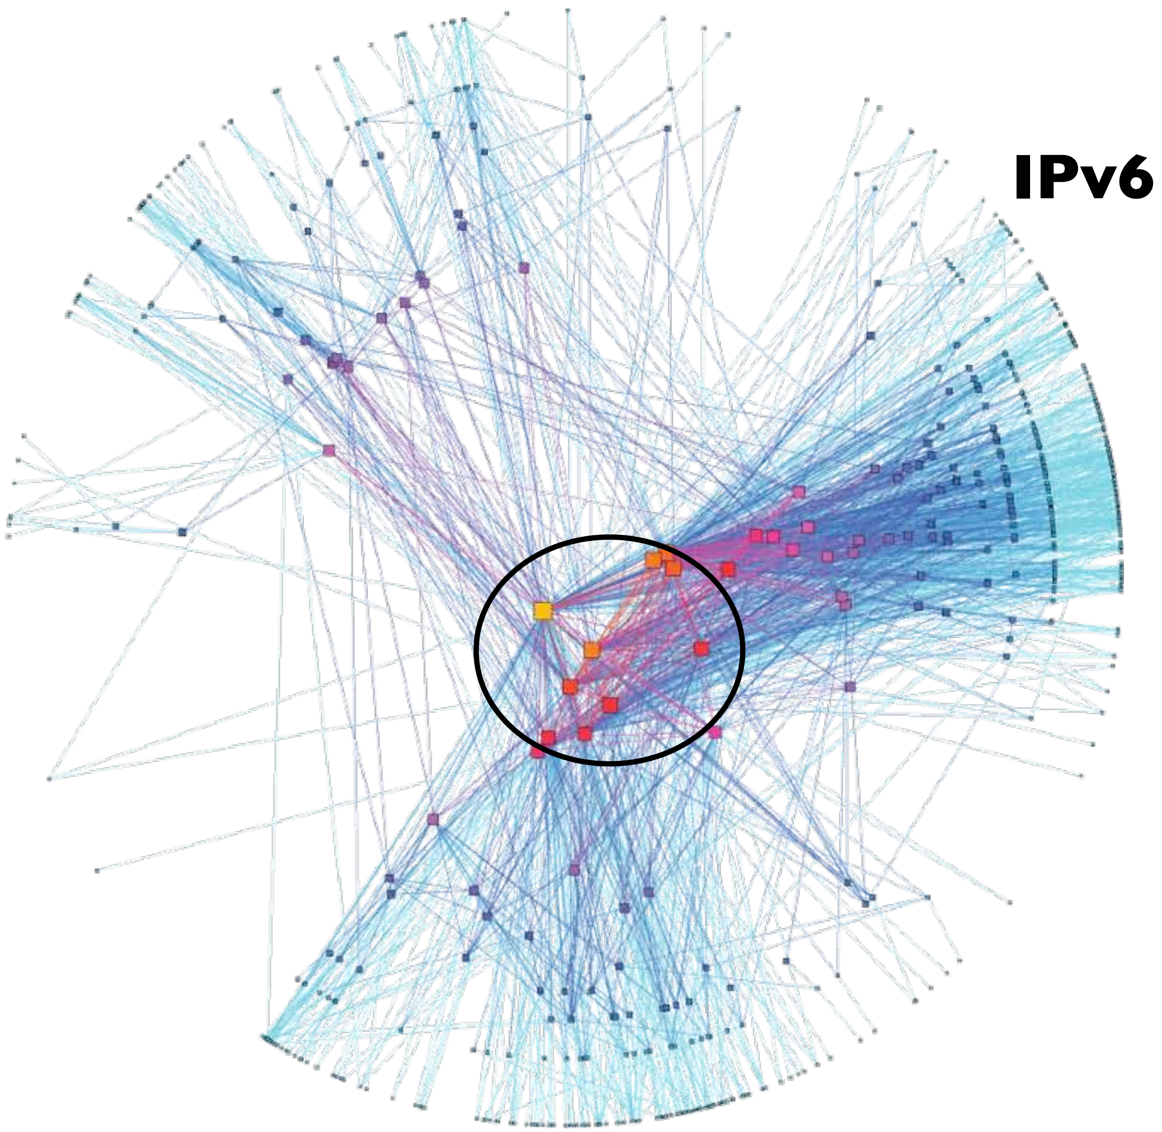

IPv6 core spread between America and Europe

IPv4 vs IPv6 cores

IPv4 vs IPv6

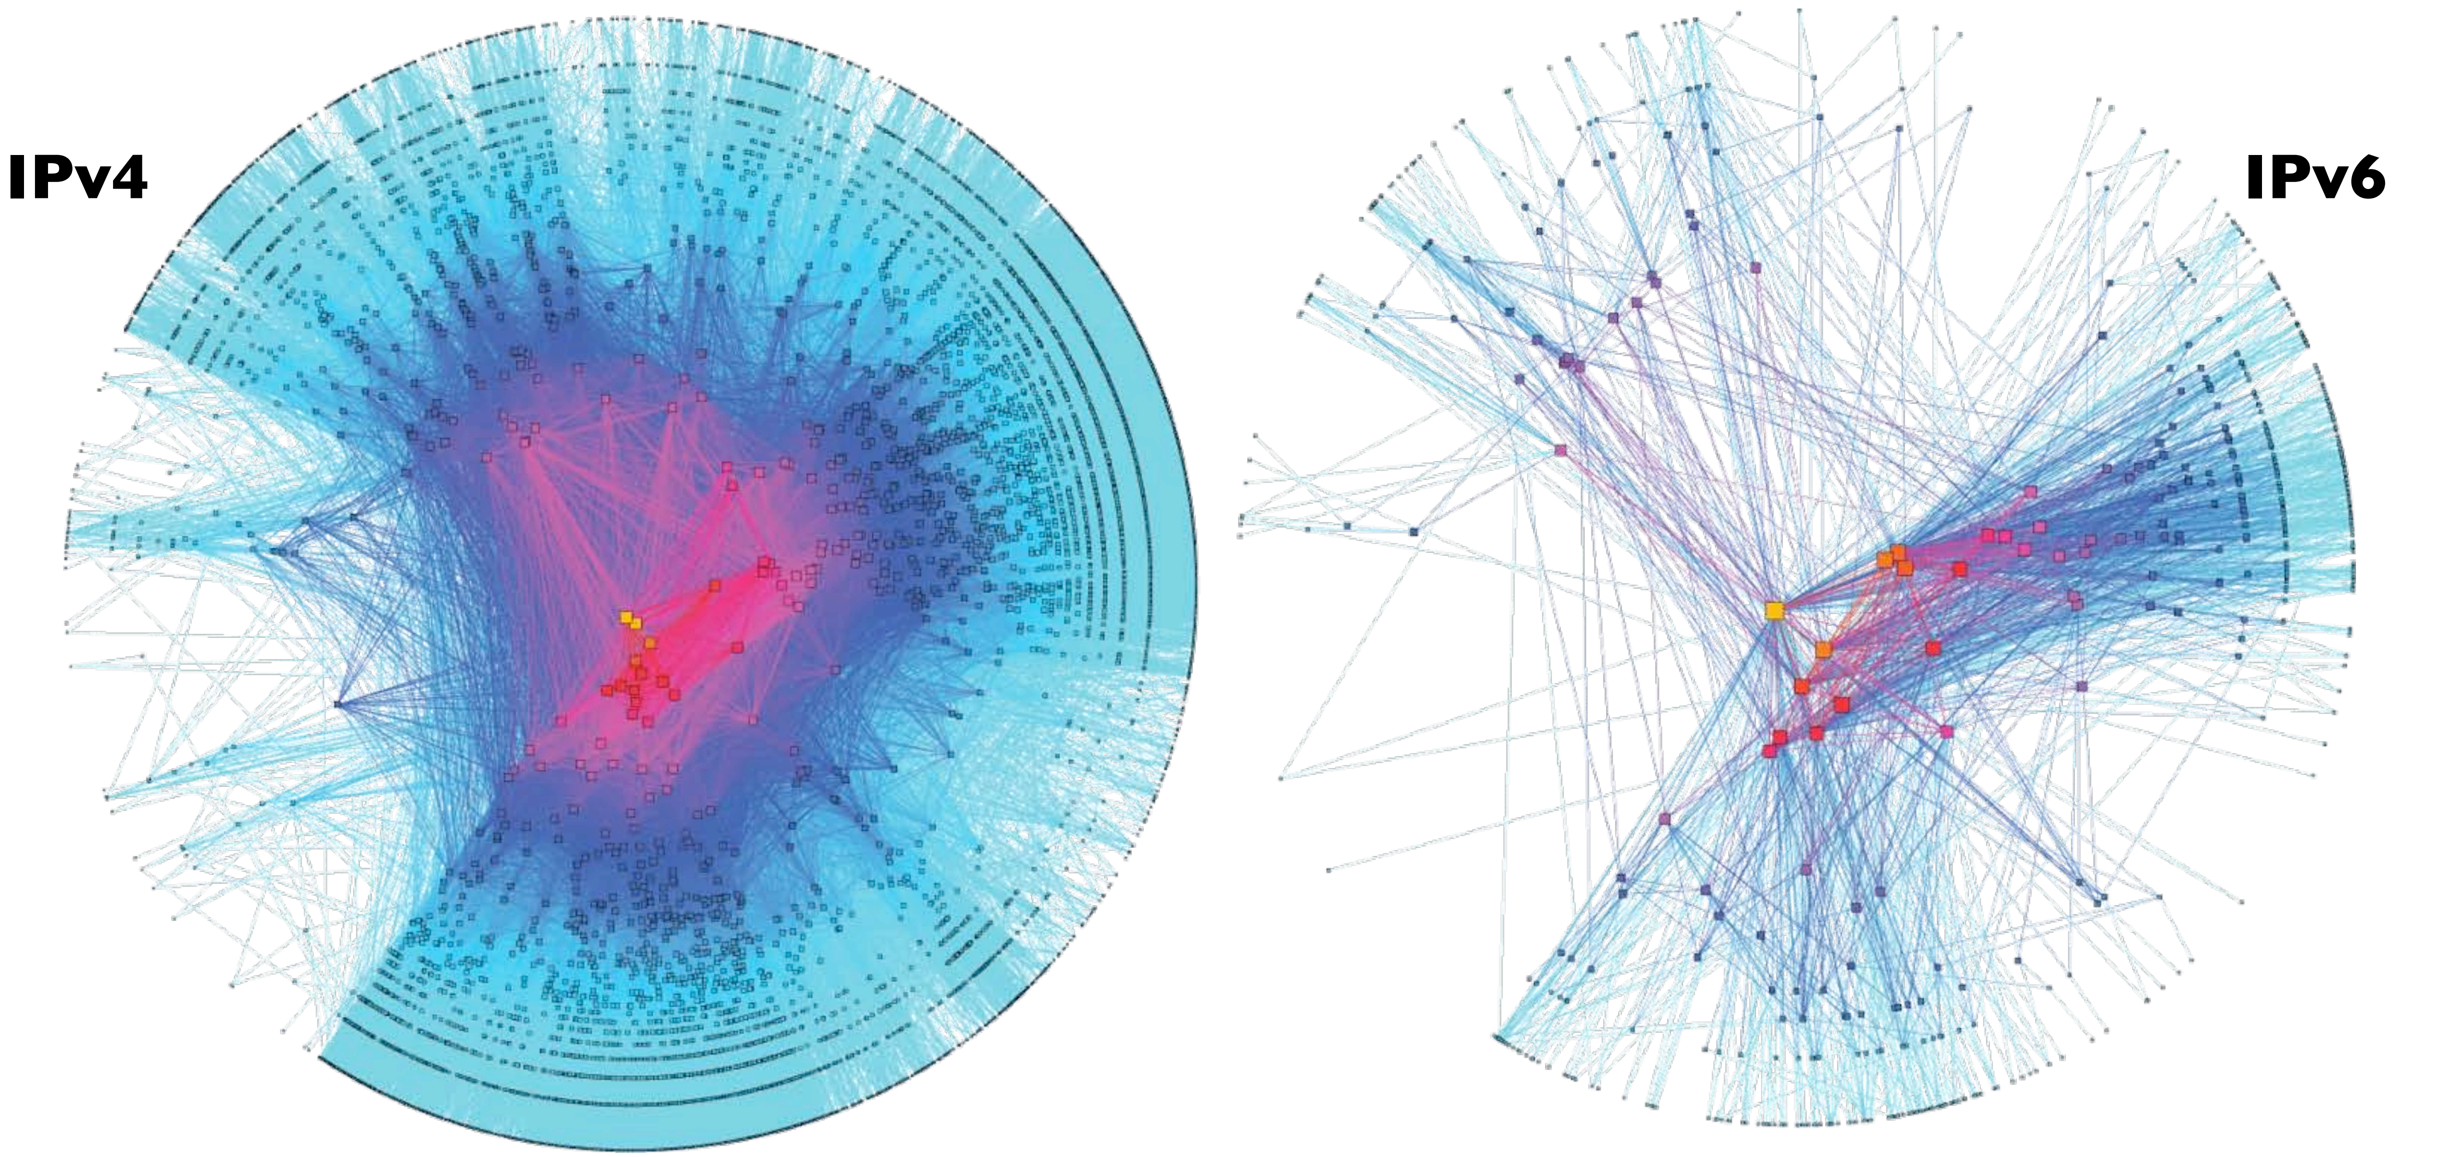

American ISPs have been slower then European ISPs to take up IPv6.

With IPv4 exhaustion finally here, will this change?

URLs

summary

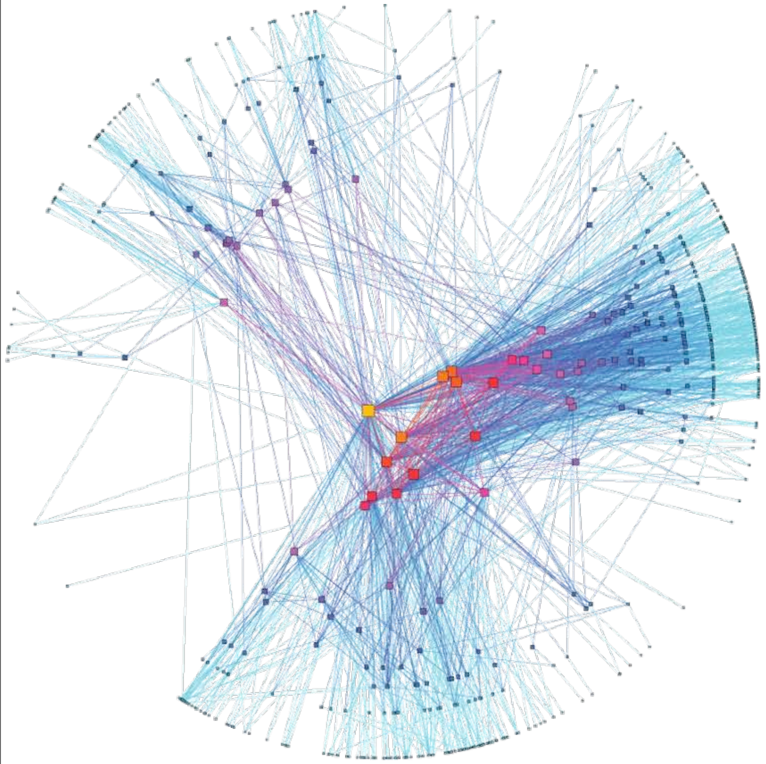

- Archipelago
  - http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml
- BGP collectors
  - http://www.ripe.net/data-tools/stats/ris/ris-raw-data
  - http://www.routeviews.org
- MaxMind GeoLite
  - http://www.maxmind.com/app/geoip_country
- IPv4 RIR exhaustion
  - http://www.potaroo.net/ispcol/2010-10/when.html

Questions?

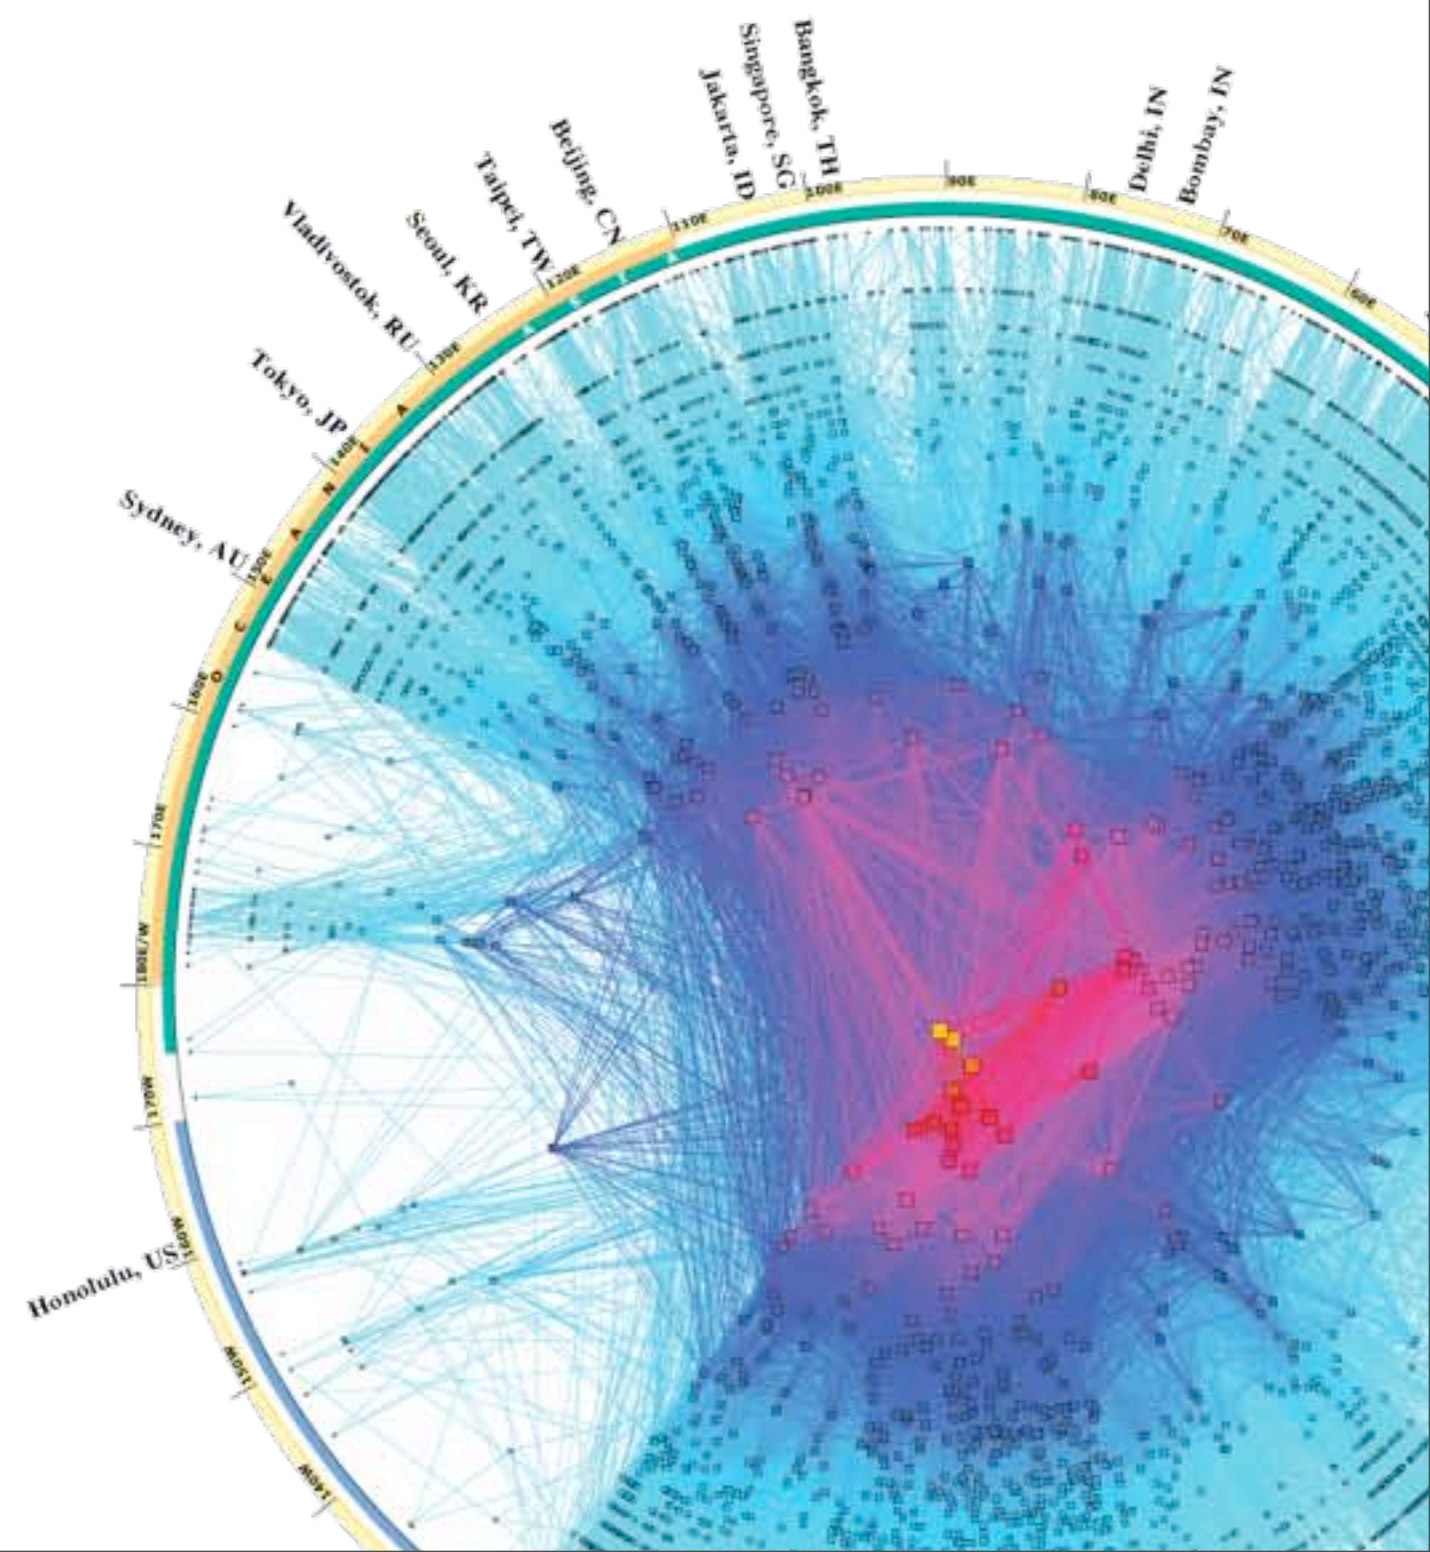

Internships:

http://www.caida.org/home/jobs/