# Autonomous Systems (AS) Introduction and Topology

**Bradley Huffaker**
**CAIDA**
**SDSC/UCSD**

WIND 16
March 2016

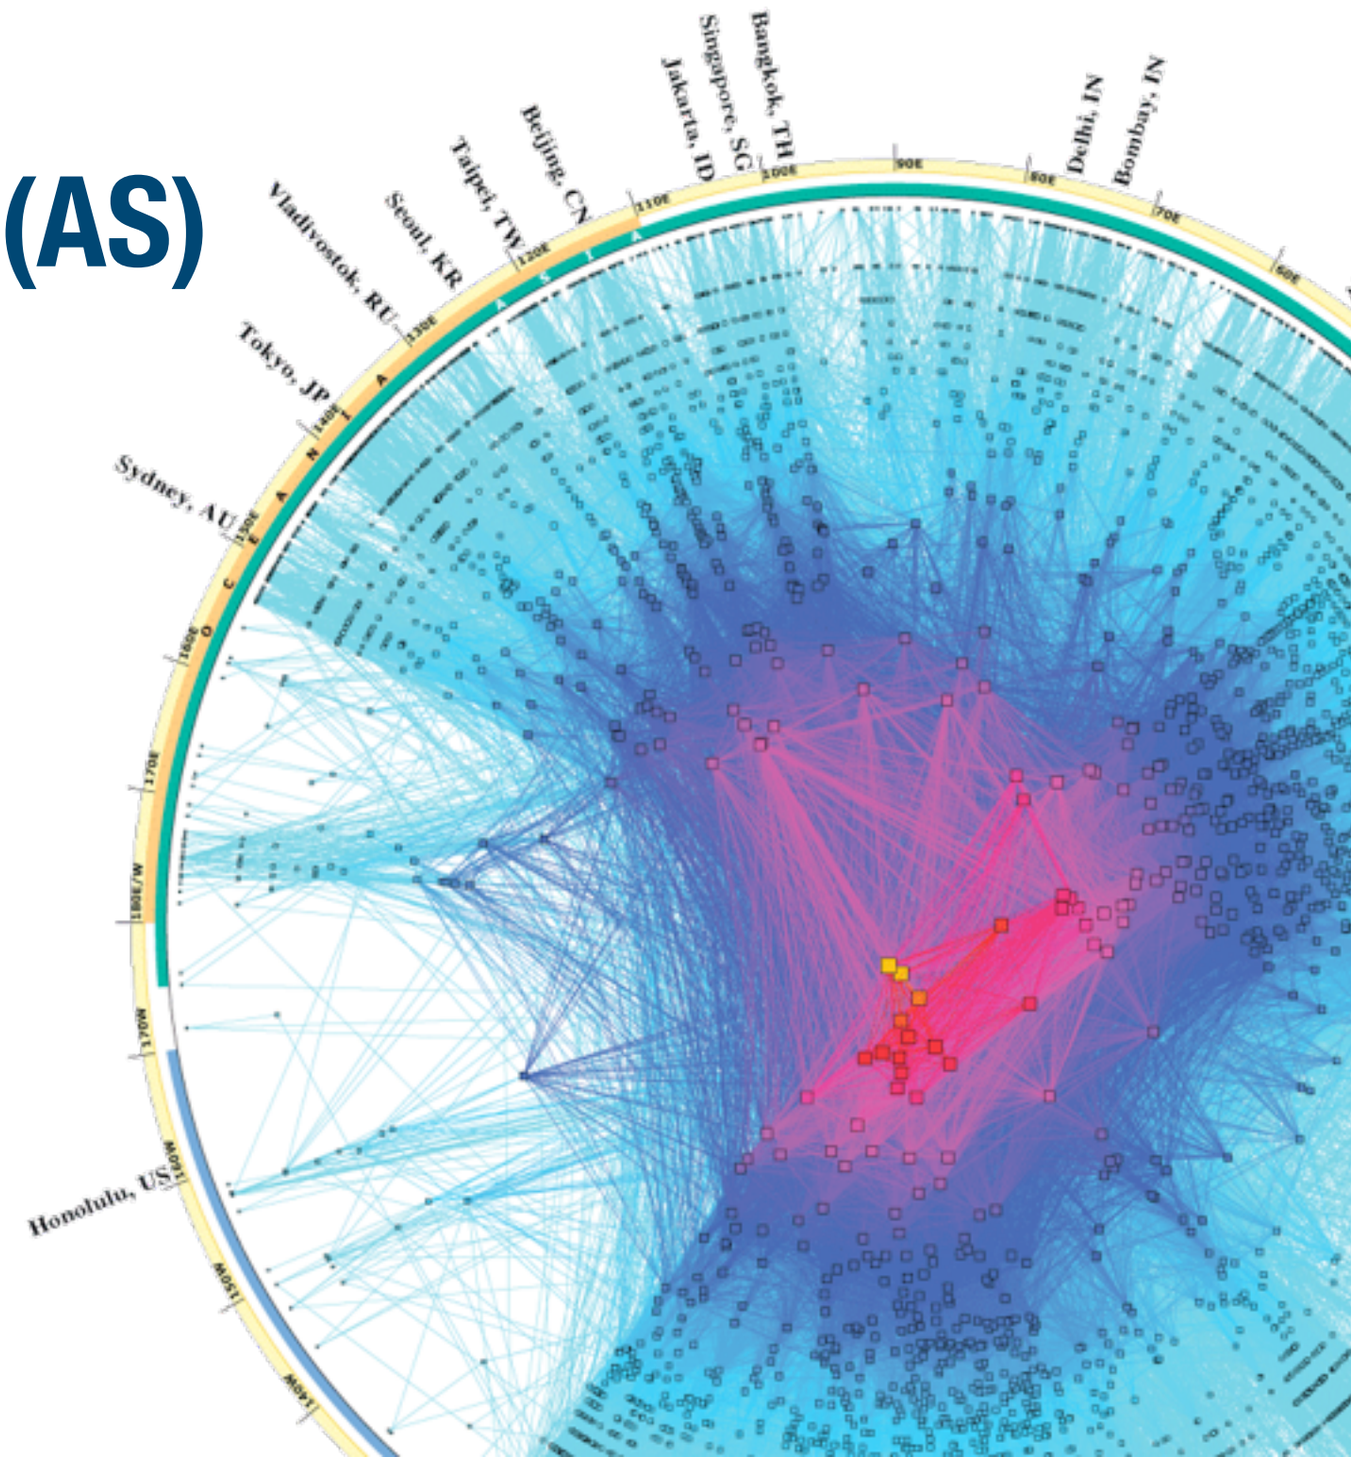

# overview



- **introduction**
- **datasets**
  - paths, locations, organizations, relationships, classifications
- **open questions**
- **summary**

# Internet of layers

overview

**Internet maps are often grouped into four levels.**

**Autonomous Systems (AS)** are numbers used to route groups of **IP addresses.**

**Points of Presence (PoP),** locations where ISP store their **routers.**

**Routers,** machines that route the traffic, interconnect via **IP addresses.**

**IP addresses** that connect devices on to the Internet.

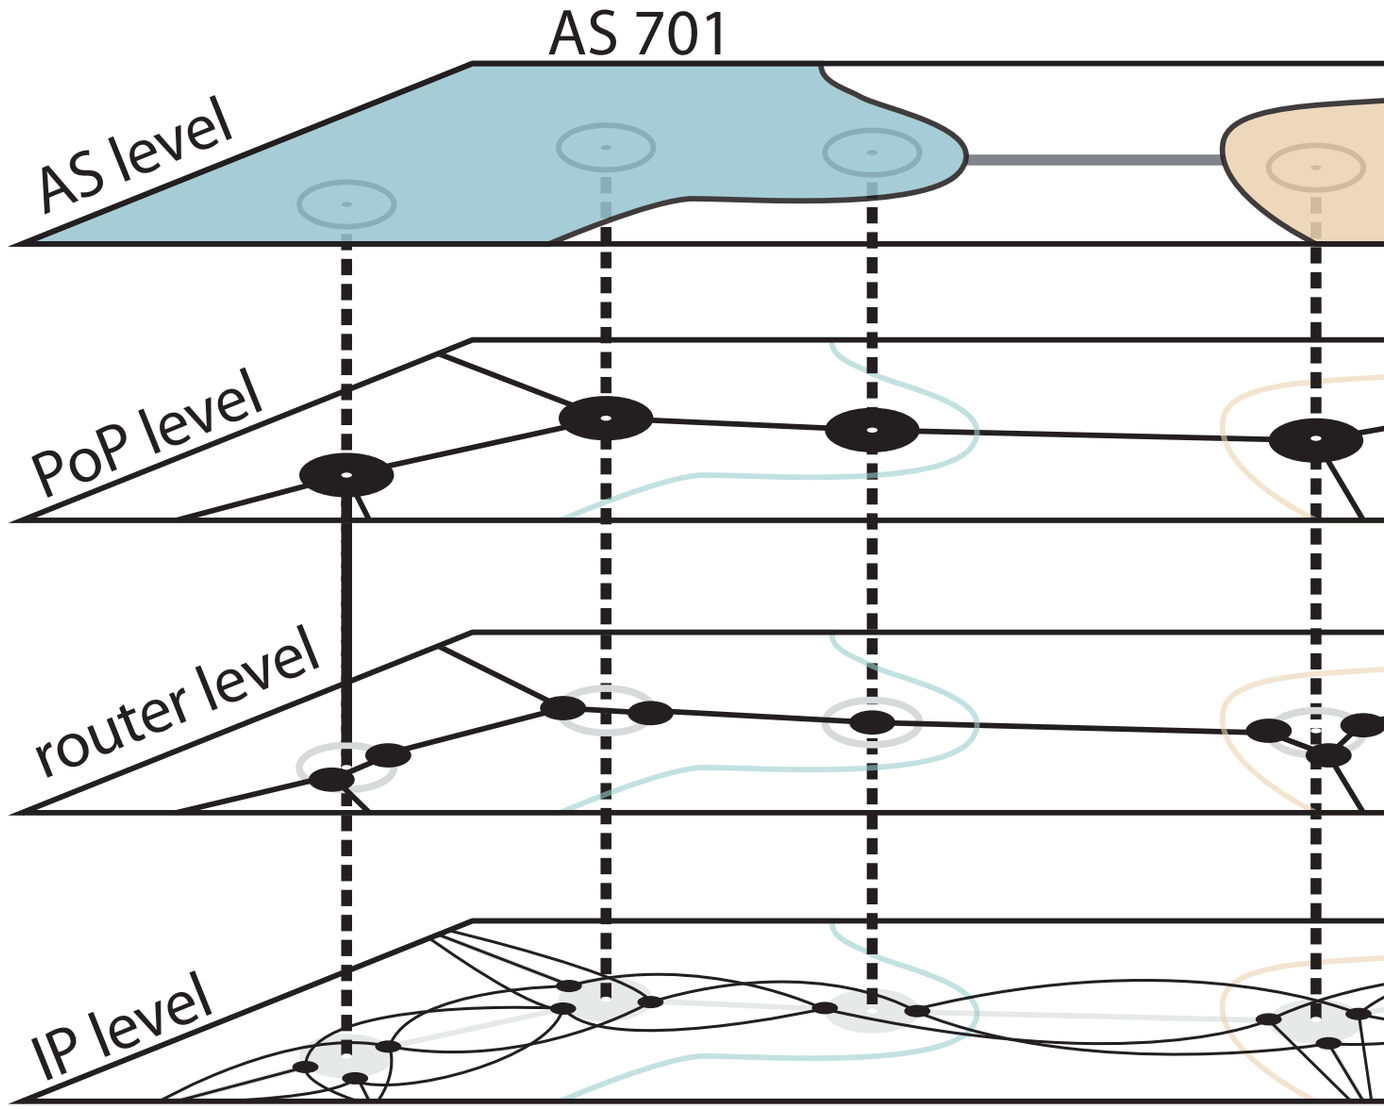

# Focus on AS

overview

**Autonomous Systems (AS)** are numbers used to route groups of **IP addresses.**

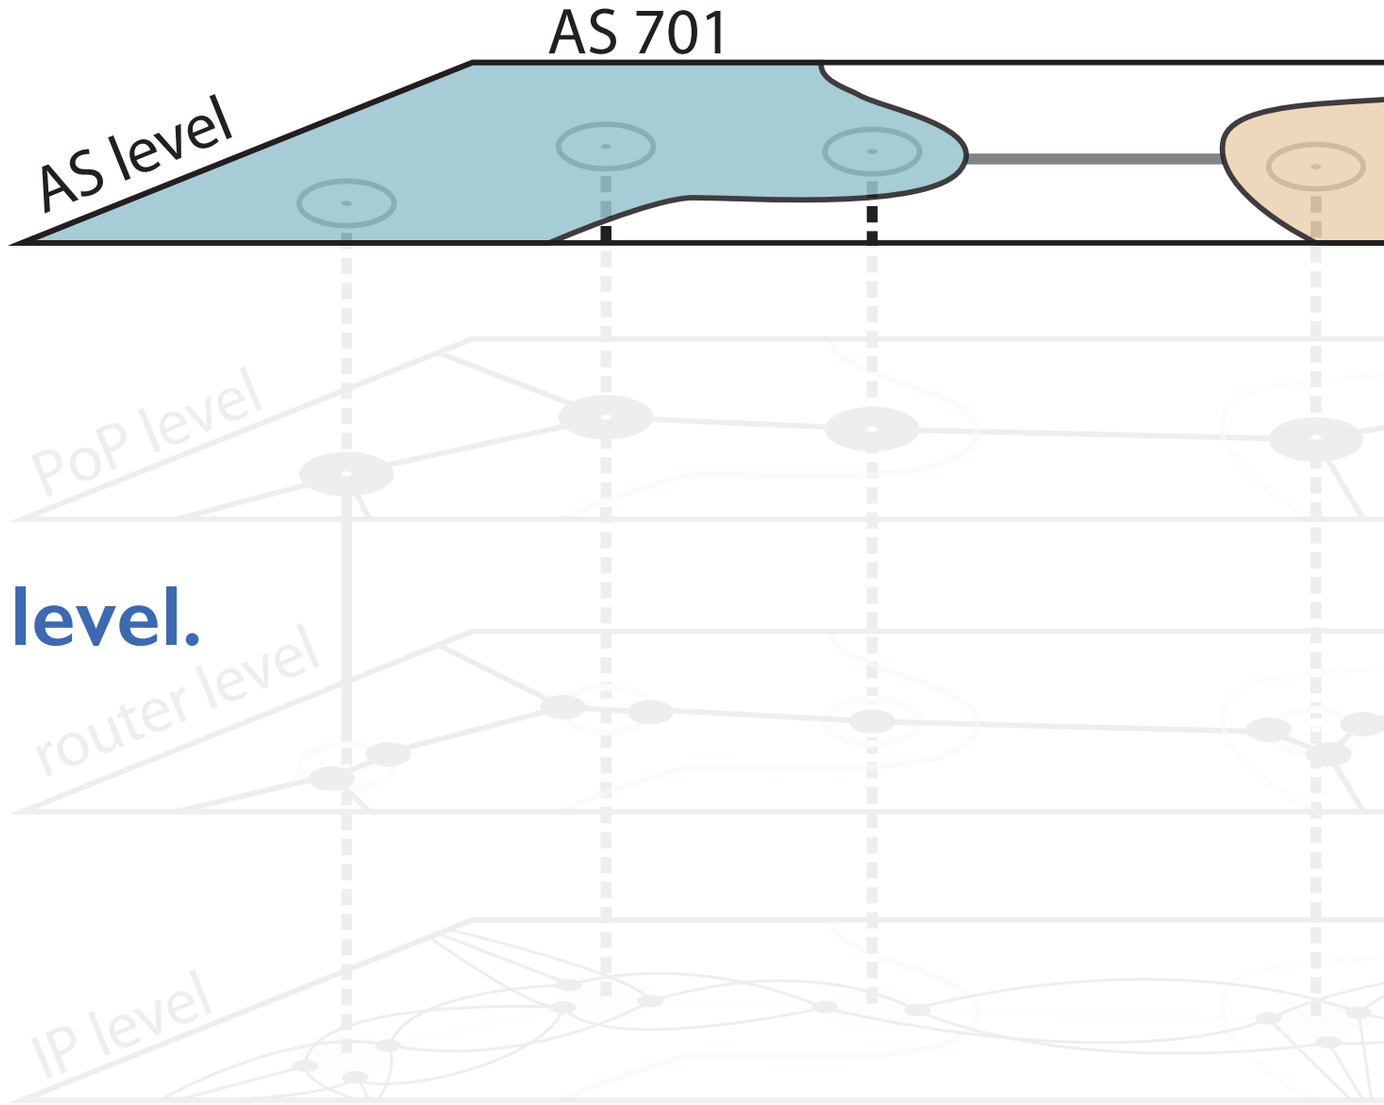


## We will focus on the AS level.

# what is an AS?

overview

**An AS can roughly be thought of as a single organization.**

Some companies use multiple ASes.

**AS topology (observed January 2016)**

- 51,000+ ASes
  - organization location
  - country of registration
- 418,000+ AS links
  - business relationship

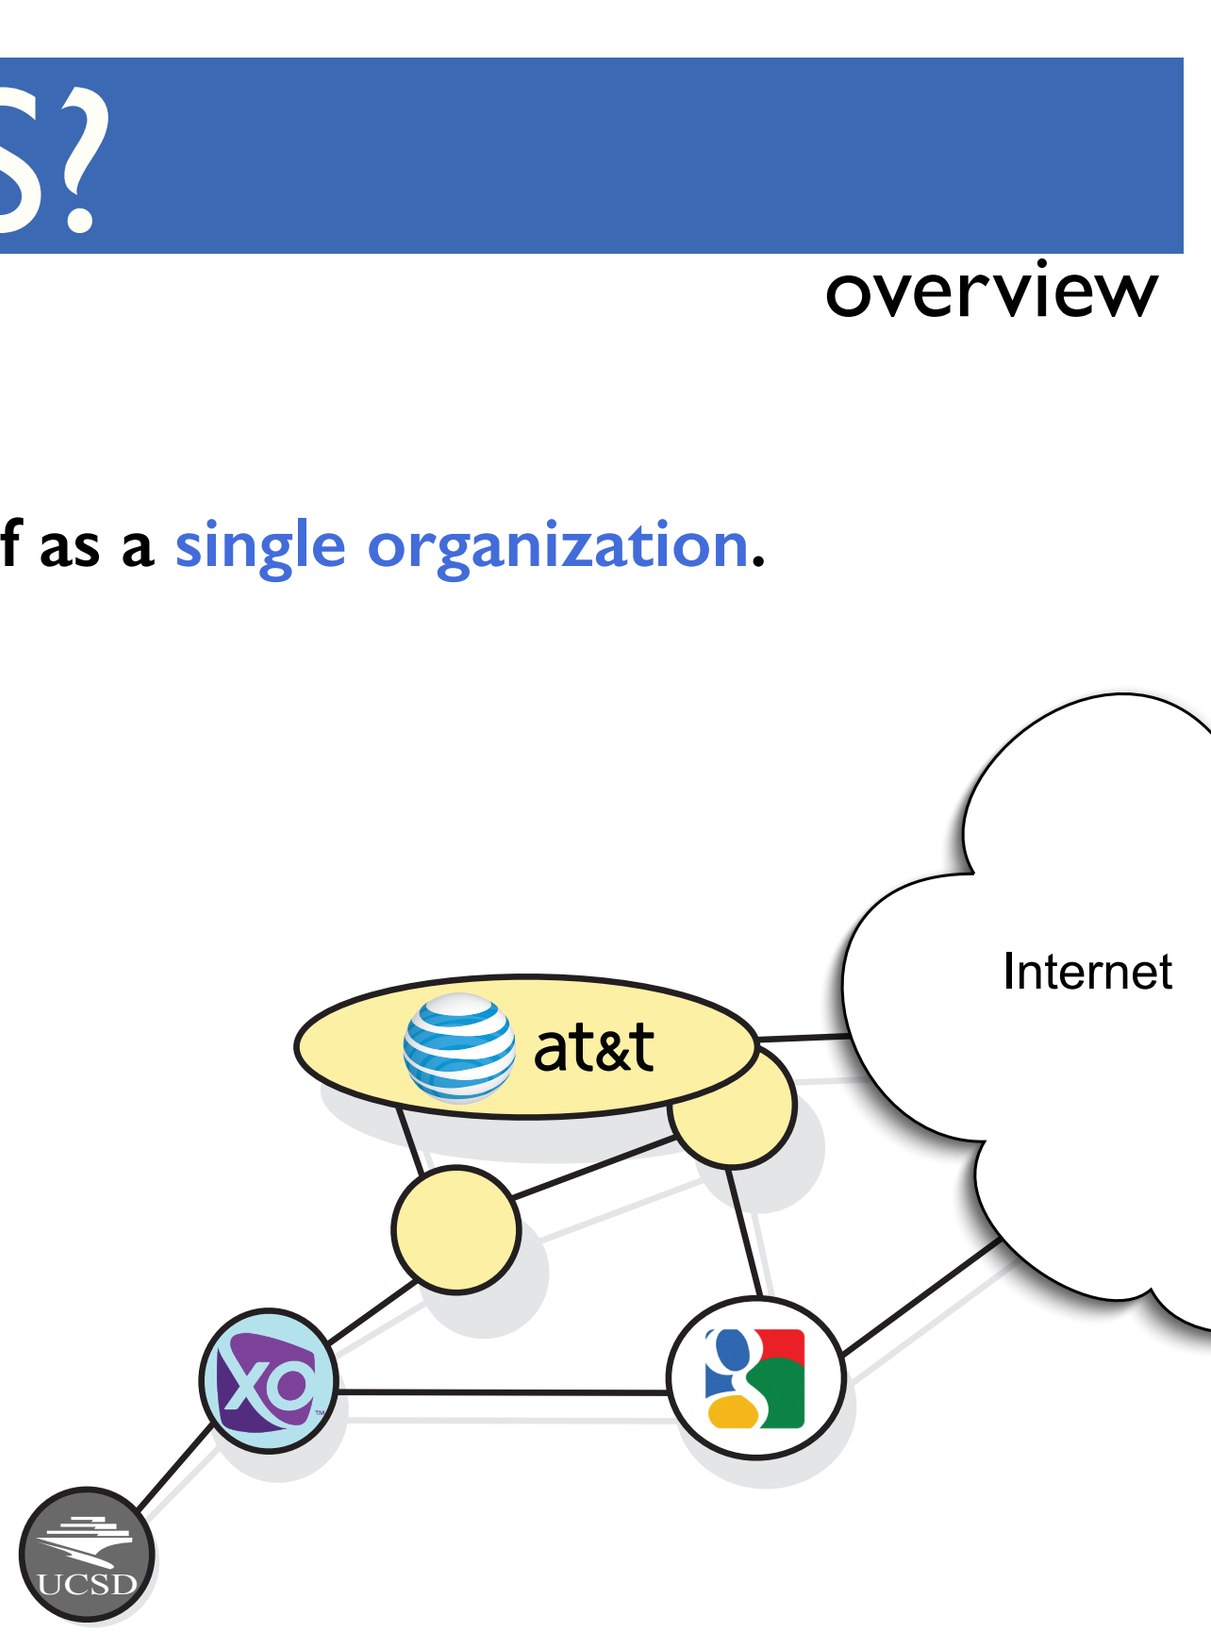

# Collecting AS paths (BGP)

paths: datasets

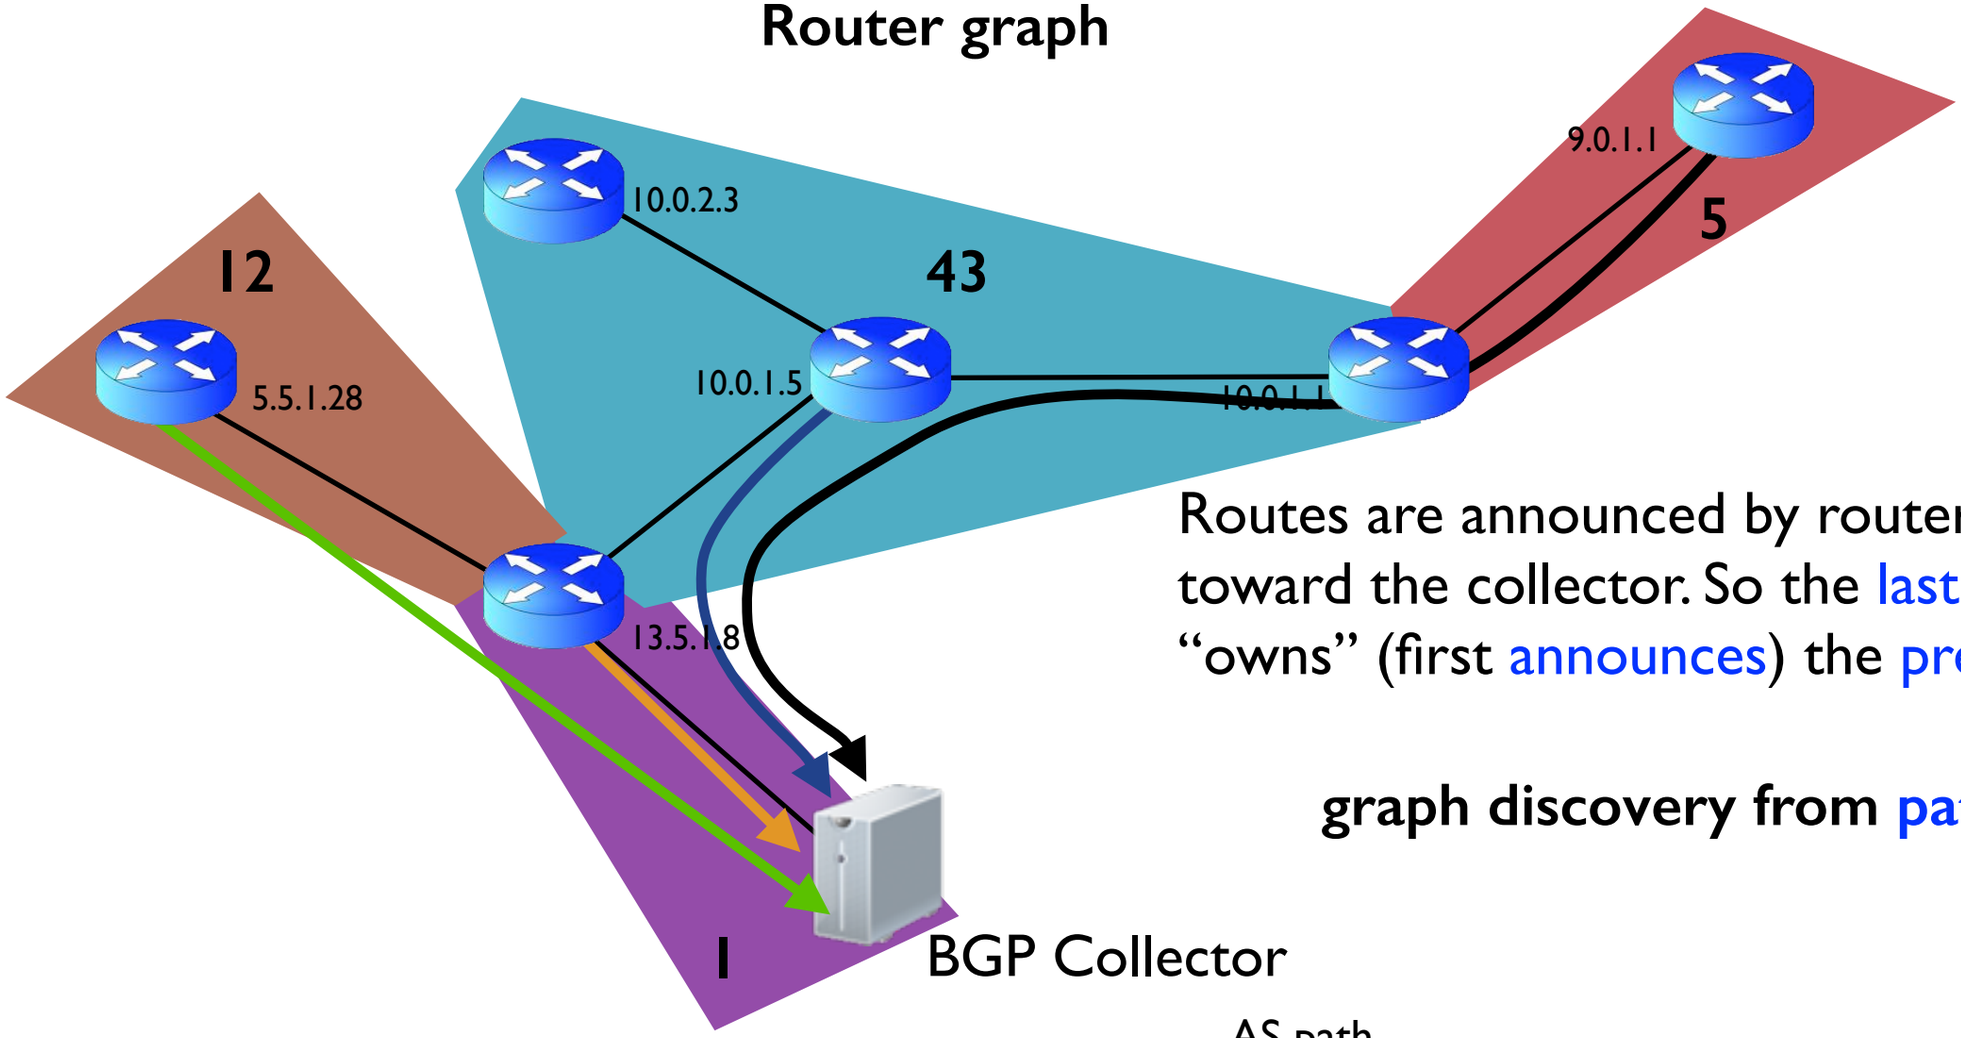


Routes are announced by routers and forwarded toward the collector. So the last AS, the "origin" AS, "owns" (first announces) the prefix.

**graph discovery from paths, not nodes**

| | 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 |

# BGP Collectors (raw)



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

- University of Oregon, Route Views Project
  - http://www.routeviews.org
- RIPE NCC (Regional Internet Registry for Europe/Middle East)
  - http://www.ripe.net/data-tools/stats/ris/ris-raw-data

**BGP dump**

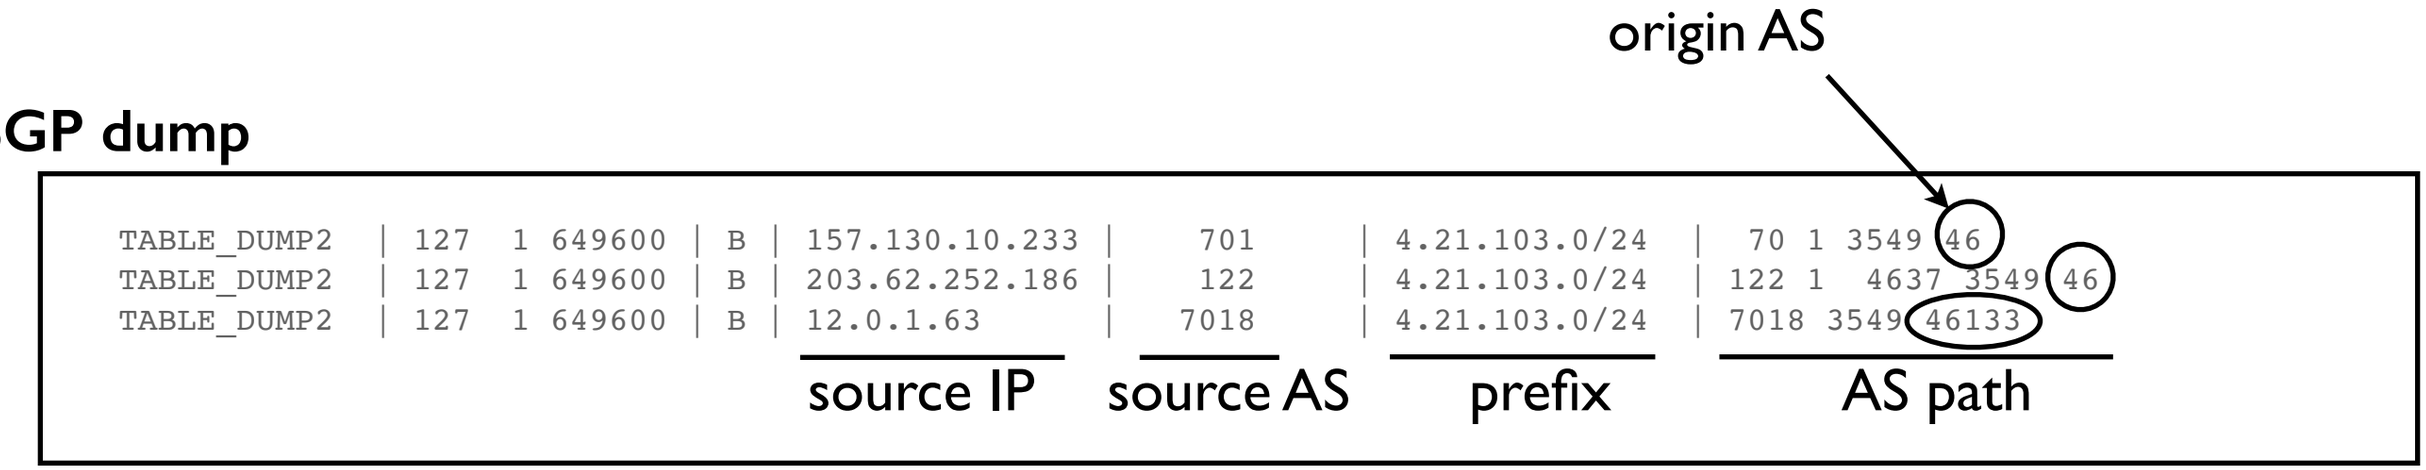

# AS path (file)



20150101.paths.bz2

```
 # source:topology|BGP|20150104|routeviews|wide
# source:topology|BGP|20150105|routeviews|wide
513|3320|1299|24961|13301
28917|3356|22773|16653
28571|1251|20080|6939|4766|38420
25220|1299|2914|4648|4610|17746
25091|2914|286|8529
202109|1299|12778|49725
9002|9304|17408|131149
14840|6453|4755|45820|22853
```

The top of the file lists the sources used to generate the file.
# source:<type>|<source>|<date>|<organization>|<server>

The actual AS path. <source>|<hop0>|<hop1>….|<destination>
The **source AS** is where the path was observed.

Your view of the topology is limited by the number of sources and origin ASes. Depending on what you are doing with the data, it often helps to filter paths to only sources that have all or almost all origin AS.

# AS degree types



We have found that transit degree is a better metric for capturing the "importance" of an AS, since it encodes some of the information in the observed paths.

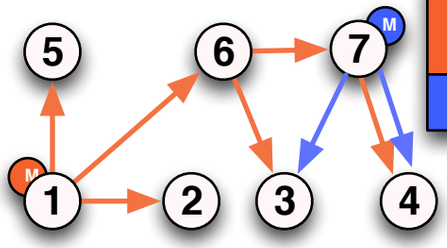


| monitors | undirected paths | | directed paths | | triplets | |
|---|---|---|---|---|---|---|
| M1 | 1-2<br>1-6-3 | 1-5<br>1-6-7-4 | 1→2<br>1→6→3 | 1→5<br>1→6→7→4 | 1-6-3 | 1-6-7<br>6-7-4 |
| M7 | 7-3 | 7-4 | 7→3 | 7→4 | | |

**degrees**

| AS | global | out | transit |
|---|---|---|---|
| 1 | 3 | 3 | 0 |
| 2 | 1 | 0 | 0 |
| 3 | 2 | 0 | 0 |
| 4 | 1 | 0 | 0 |
| 5 | 2 | 0 | 0 |
| 6 | 3 | 2 | 3 |
| 7 | 3 | 2 | 2 |

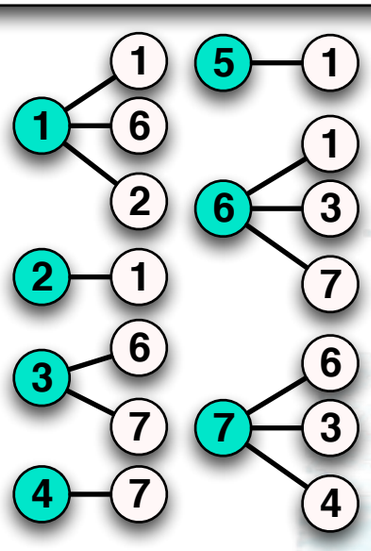

global degree

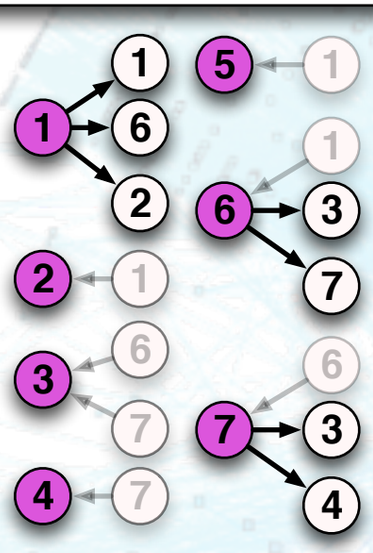

out degree

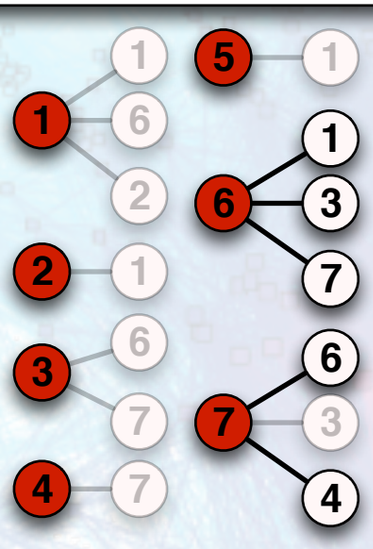

transit degree

# AS geolocation



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

**Netacuity geographic dump**

```
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
```

IP first | IP last | country state city | latitude longitude

# geolocation to longitude

caida

locations: datasets

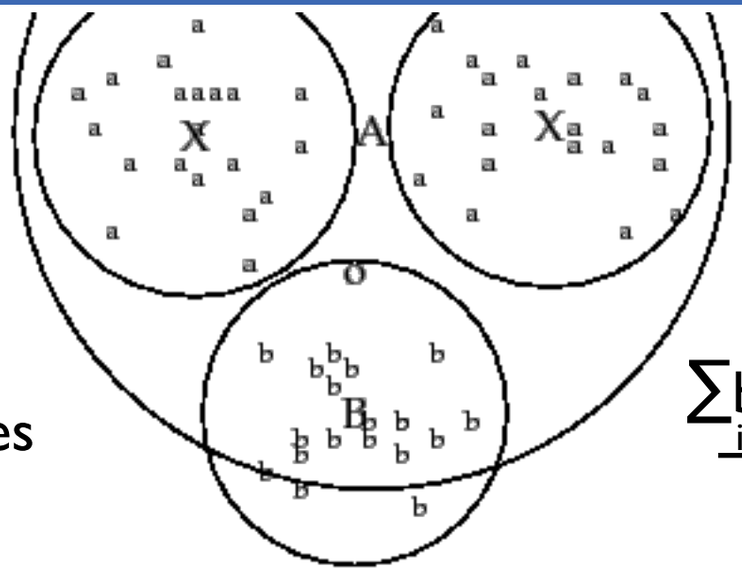


An AS's longitude equals the weighted average of the longitudes assigned by Netacuity to the addresses it announces.

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

| 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 |

# AS centroid (file)

locations: datasets
open question

20150101.as2loc.txt

```
 #format:as|continent|country|region|city|lat|long|total-prefixes
1|||||36.023026628203|-88.3649892886005|55
2|||||39.9882562979862|-75.4179414754833|2
3|||||42.3643803602359|-71.1006790537166|33
```

number of prefixes announced by AS

longitude

latitude

continent, country, state, city
empty in this file

AS number

# AS organization (file)

organizations: datasets
open question

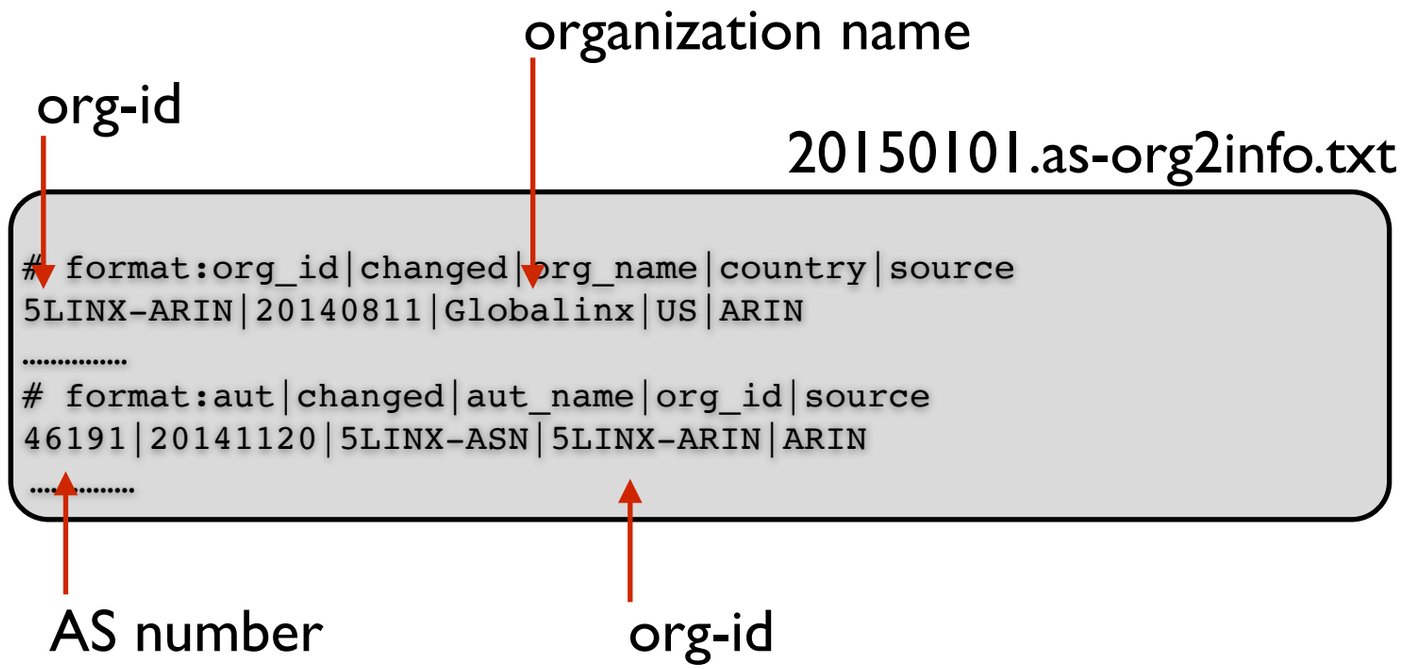


This file contains both organization and ASes. The type changes at every line that starts with "# format:<type>". Connect between entries with org-id. In this example the AS 46191 belongs to organization Globalinx.



http://www.caida.org/data/as-organizations/

# AS relationships

organizations: datasets

**provider**

- you pay them to transit your traffic
  (XO pays AT&T)

**peer**

- unpaid exchange of traffic
  (between XO and Google)

**customer**

- they pay you to transit their traffic
  (UCSD pays XO)

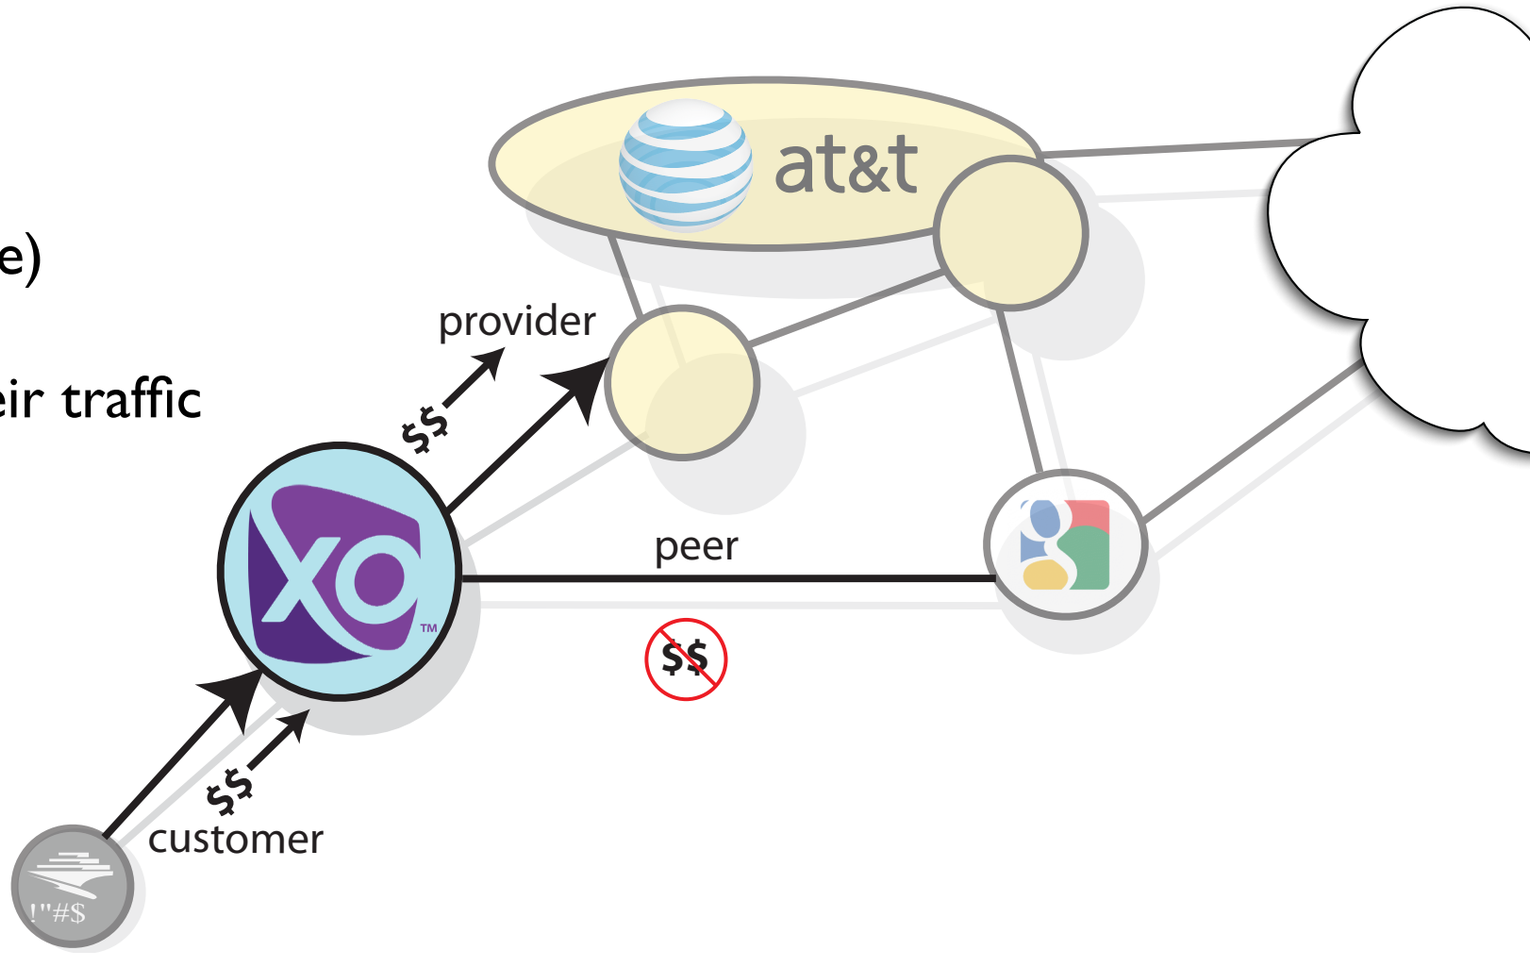

# AS hierarchy

relationships:datasets

**ASes can be organized into a hierarchical structure based on the type of business relationships they form between themselves.**

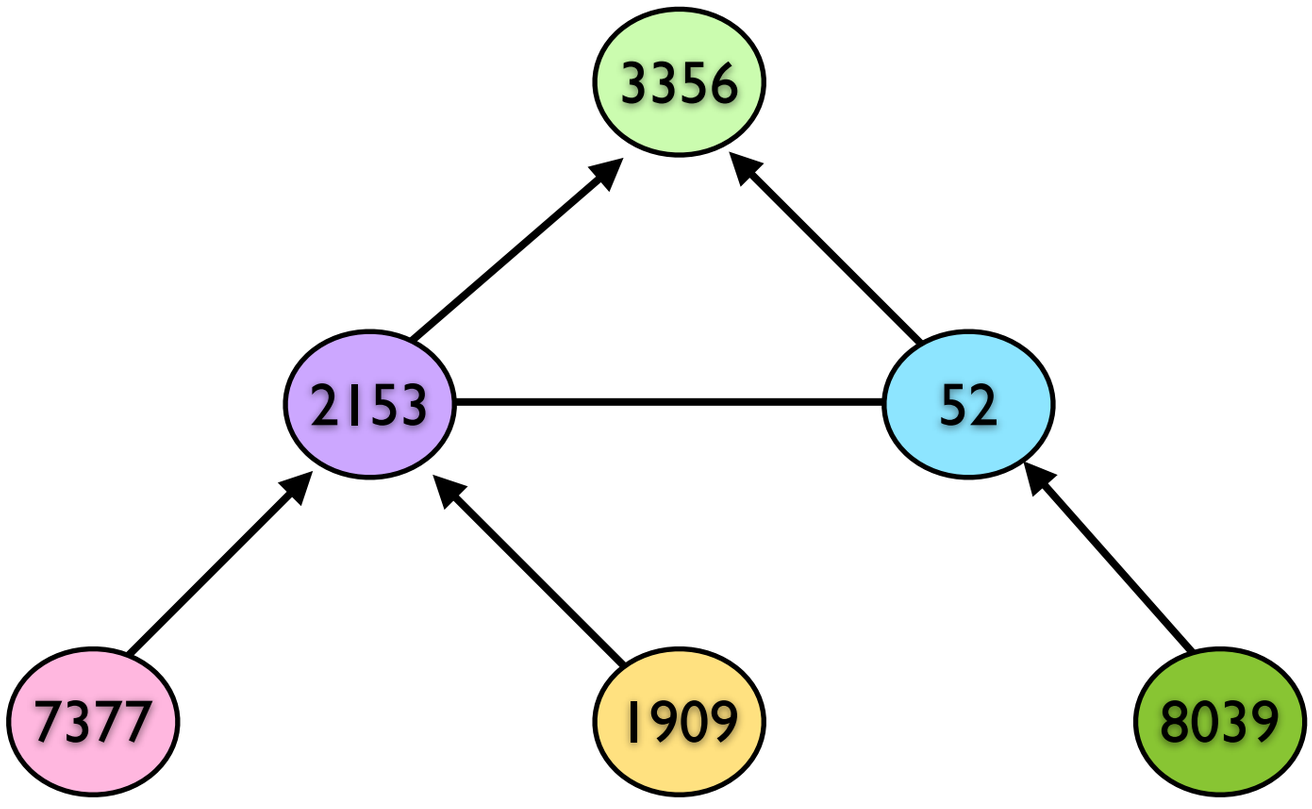


$$ →

customer ——→ provider

customer pays provider for transit

$$

peer —— peer

peers do not pay to accept each other's traffic

# AS hierarchy



**ASes at the top of the hierarchy are call Tier-1.**
**Together they form the backbone of the Internet.**

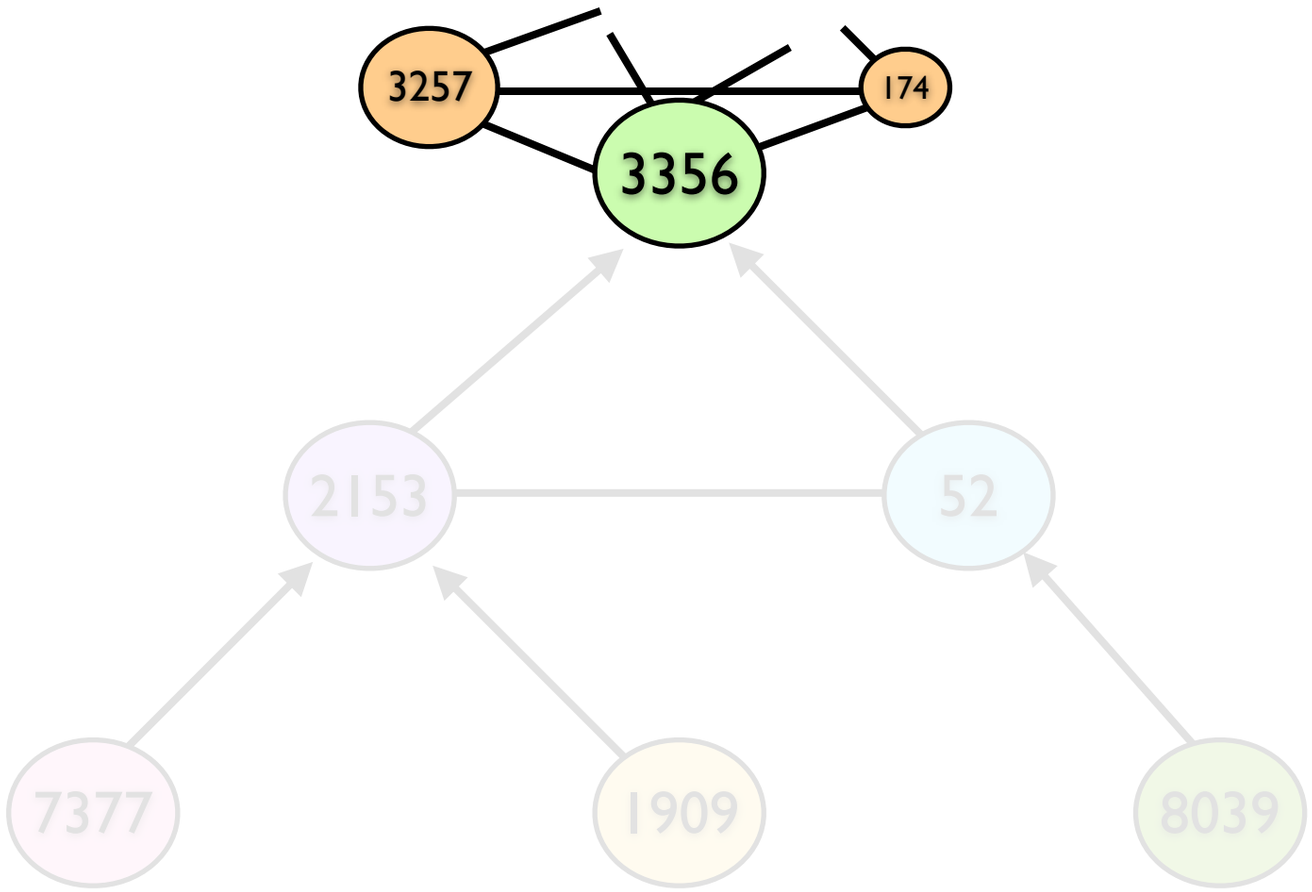

# AS relationship (file)

relationships:datasets

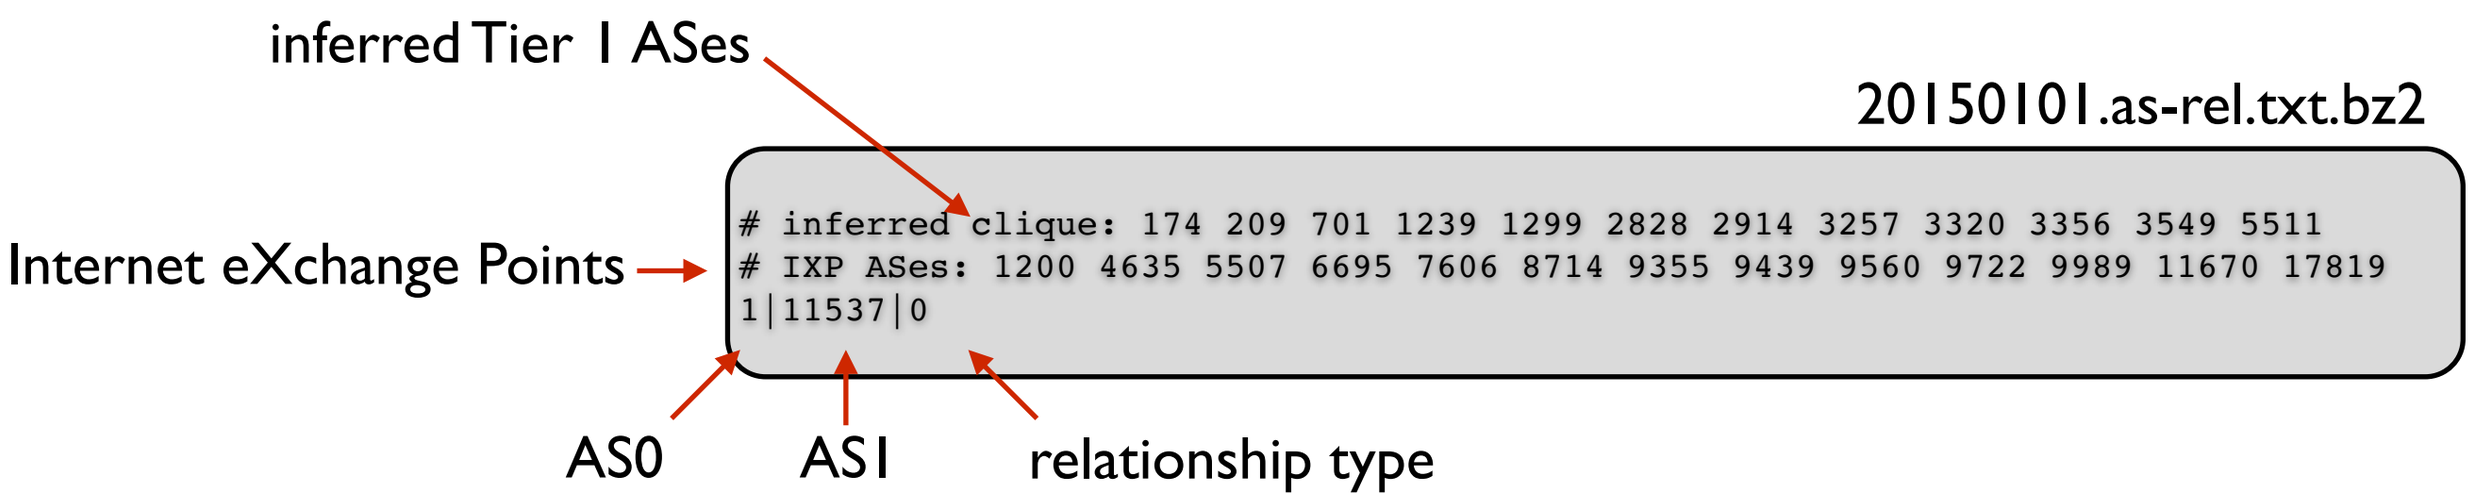


| | | |
|---|---|---|
| -1 | AS0 is a provider of AS1 | AS0 ← AS1 |
| 0 | AS0 and AS1 are peers | AS0 — AS1 |
| 1 | AS0 is a customer of AS1 | AS0 → AS1 |



http://www.caida.org/data/as-relationships/

# AS customer cone

relationships: datasets

**An AS's customer cone contains the set of ASes we observe the AS announce to its peers or providers. It is important that this is more constrained than, but similar to, the set of ASes it can reach through its customers.**

| AS | size |
|---|---|
| A | 6 |
| B | 3 |
| D | 1 |

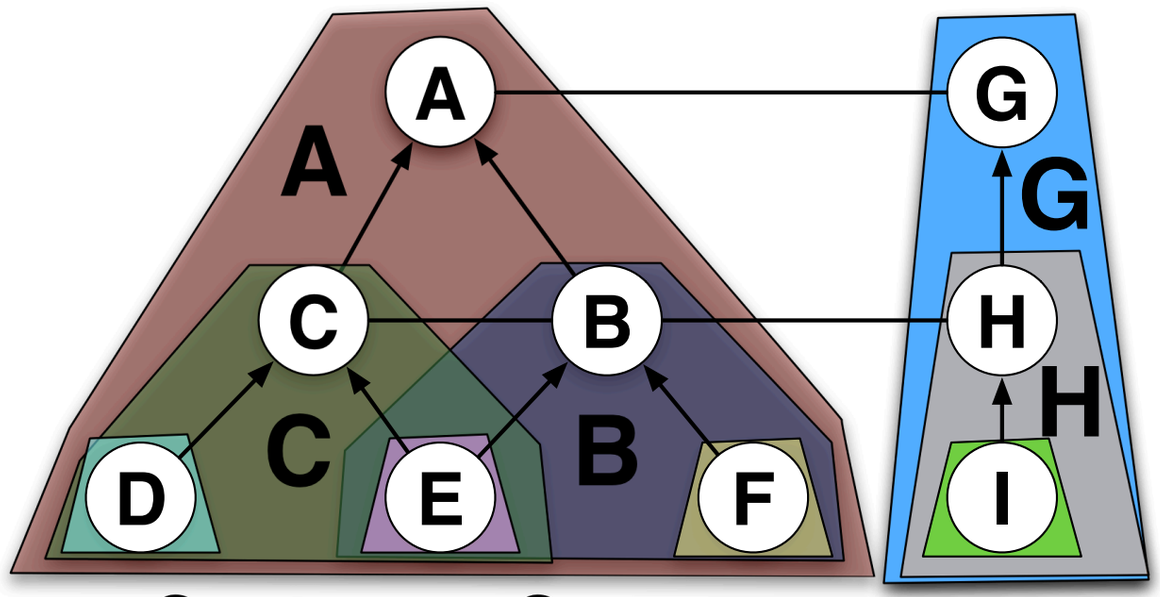


**Customer Cones**

# AS customer cone (file)

relationships: datasets
open question

20150101.ppdc-ases.txt.bz2

```
2345  2345
14    14 209 701 1239 1299 2828 2914 3257 3320 3356 3549 5511
1200  435 1200 5507 6695 7606 8714 9355 9439 9560 9722
```

AS →

customer cone members

| AS | size | members |
|---|---|---|
| 2345 | 1 | 2345 |
| 14 | 12 | 14 209 701 1239 1299 2828 2914 3257 3320 3356 3549 5511 |
| 1 | 10 | 435 1200 5507 6695 7606 8714 9355 9439 9560 9722 |



http://www.caida.org/data/as-relationships/

# AS classification (file)

classification:datasets
open question

**ASes are often divided into Transit, Content, or Access depending on the type of companies that operate them. Larger organizations can combine all three types.**

- **Transit ASes**, such as AT&T, form the backbone of the Internet carrying traffic between ASes.
- **Content ASes**, such as Netflix, provide the videos, webpages, and other content.
- **Access ASes**, such as UCSD or Time Warner Cable, are used by smaller organizations to get access to the Internet.

20150101.as2classification.txt.bz

```
# name: as2type
# ../as-graph-201409/bin/type-convert-amogh.pl
# types:Access|Content|Transit/Access
1|Transit/Access|CAIDA_class
2|Transit/Access|CAIDA_class
```

AS number | type | inference method

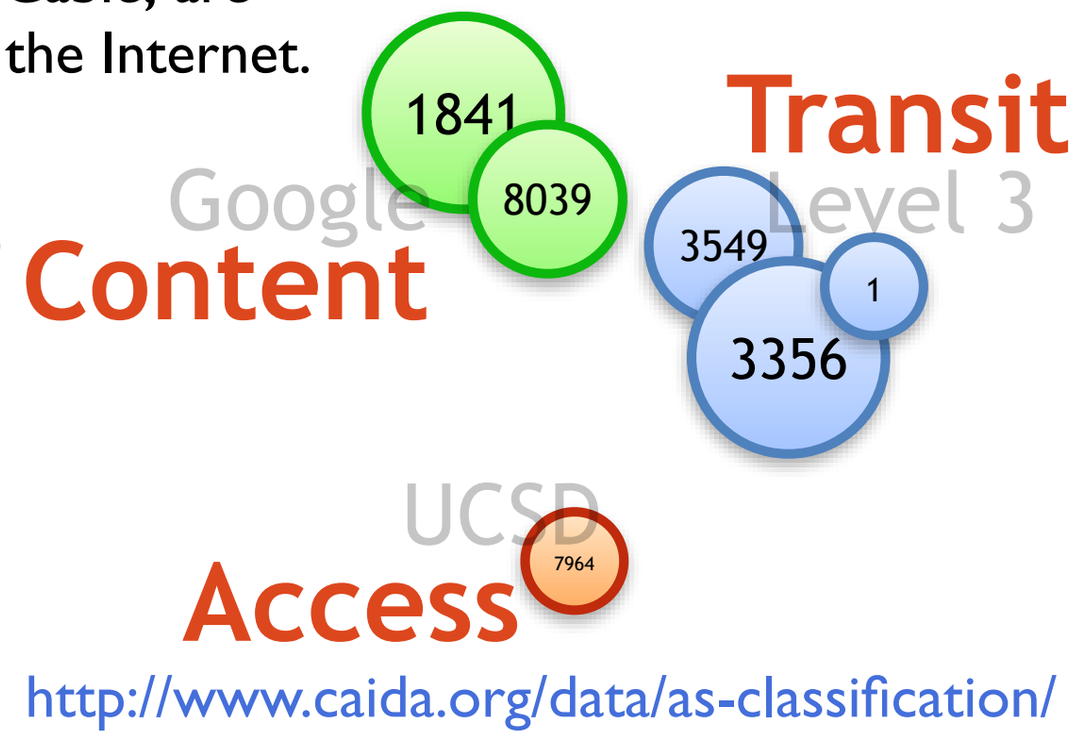


http://www.caida.org/data/as-classification/

# incomplete topology

## open questions

**Given that BGP only forwards the best path, an AS graph generated exclusively from limited number of BGP monitors will be incomplete.**

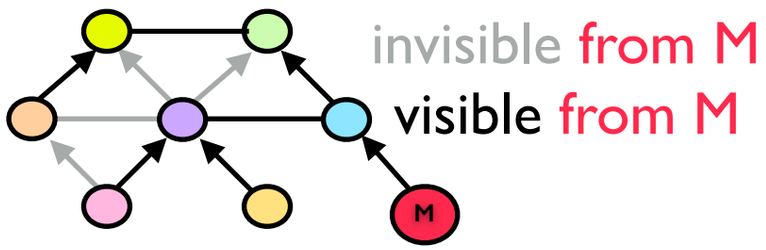


| graphs | links | | | nodes |
|---|---|---|---|---|
| | c/p | peer | total | total |
| BGP-derived | 93,539 | 83,852 | 177,391 | 46,177 |
| traceroute-derived | 0 | 7,166 | 7,166 | 2,432 |
| IX-derived | 0 | 264,803 | 264,803 | 1,555 |
| combined | 93,539 | 325,312 | 418,851 | 46,320 |

AS Topology Dec. 2015 (visibility)

- How incomplete is the topology?
- Can we infer the missing topology from what is visible?
- Given a set of limited resource monitors, can we create targeted measurements that will find the missing components?

# path prediction

open questions

**Due to policy constraints, simple shortest paths algorithms do not work for the AS topology.**

- paths must follow paid transit

valid

B is paid by A

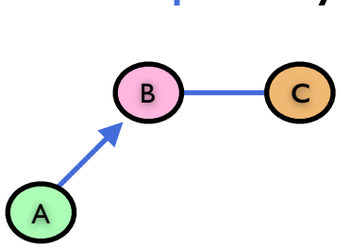


B is paid by C

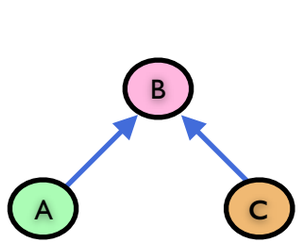


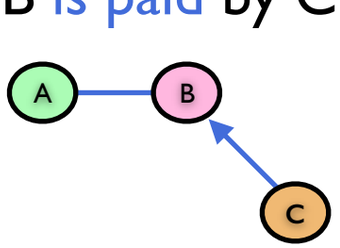


B is paid by A and C

invalid

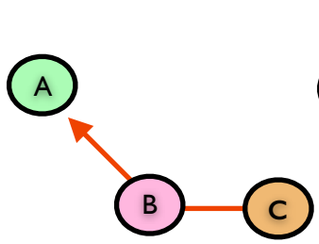


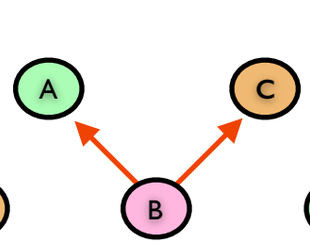


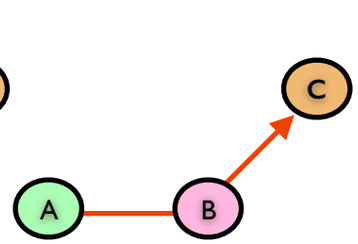


B is not paid by A or C

- current algorithms are between 60% and 80% accurate.
- asymmetric paths are common “hot potato routing”
- backup packs are hidden

# IP to Router to PoP to AS

## open questions

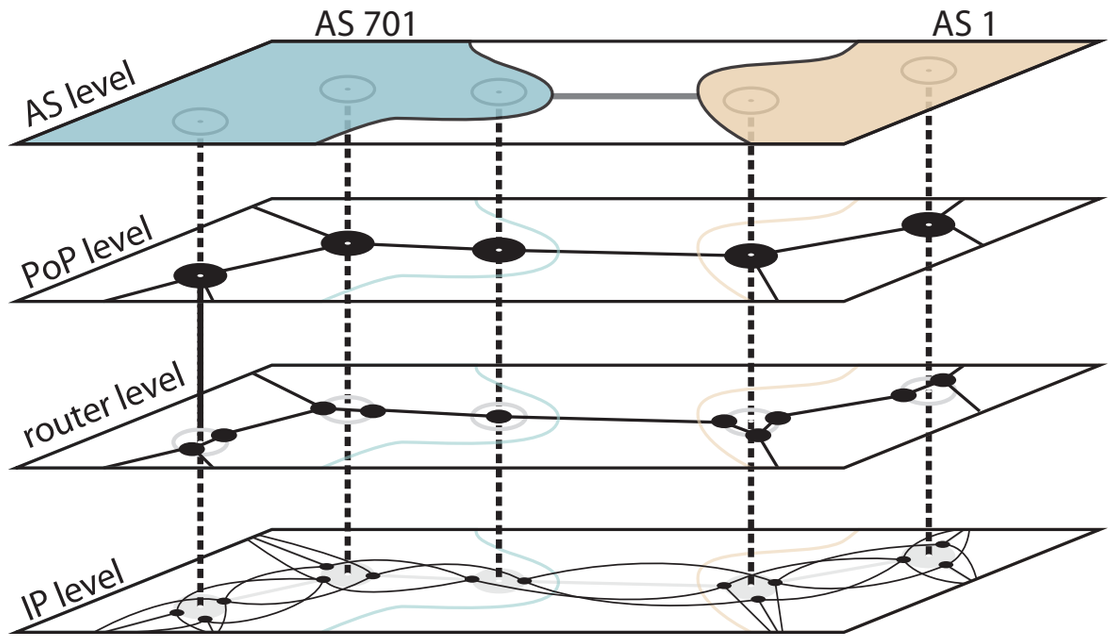


- mapping between interfaces and routers (alias resolution)
- mapping between routers and PoPs (geolocation)
- mapping between PoPs to AS (router ownership)



http://www.caida.org/data/internet-topology-data-kit/

# datasets

## summary

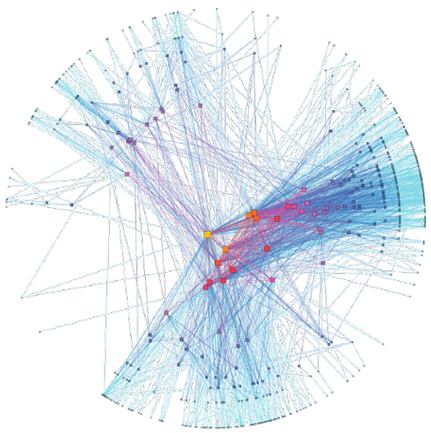

http://www.caida.org/publications/presentations/2016/as_intro_topology_wind/

| | | |
|---|---|---|
| 20150101.paths.txt.bz | AS paths | slide 8 |
| 20150101.as2loc.bz | AS location | slide 12 |
| 20150101.as-org2info.txt.bz | registration country and organization information | slide 13 |
| 20150101.as-rel.txt.bz | types of business relationships between pairs of ASes | slide 17 |
| 20150101.ppdc-ases.bz | the set of ASes that are customers, or customer's customers of an AS | slide 19 |
| 20150101.as2classification.txt.bz | an AS's type of business | slide 20 |

# Questions?

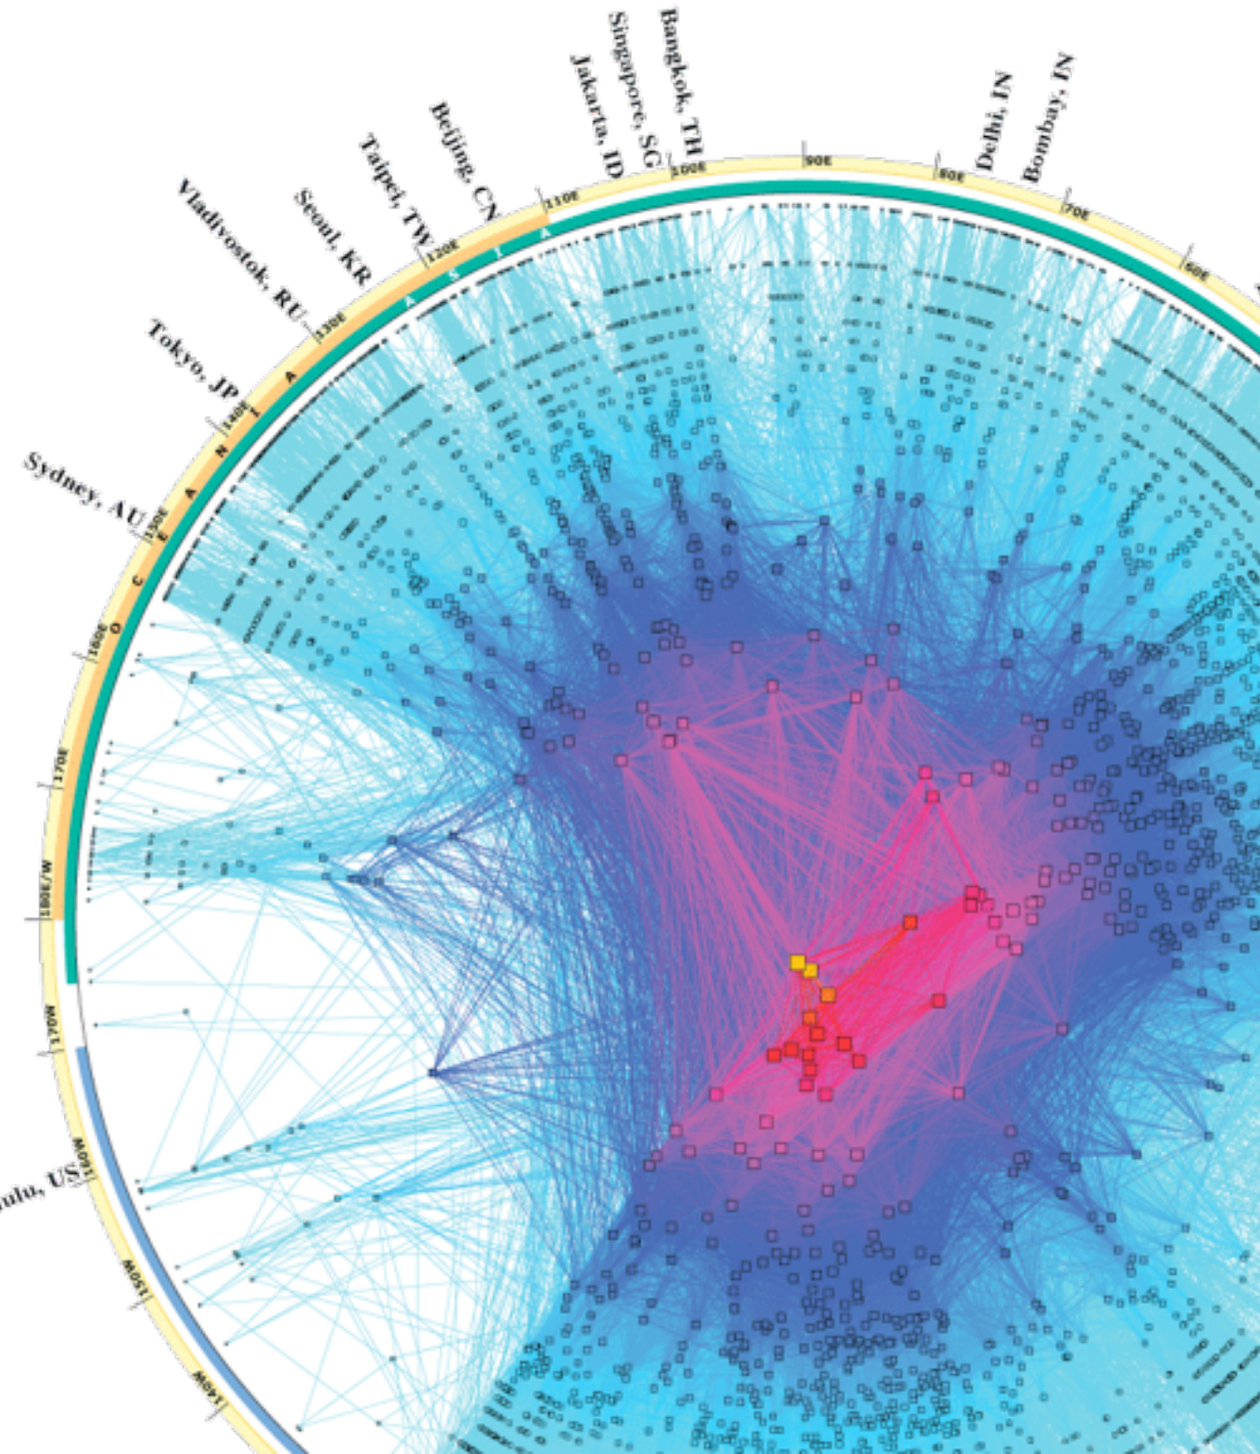


**Internships:**

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

# AS placement

bonus

| | |
|---|---|
| **node label** | center: AS number<br>upper left: rank<br>upper right: country flag |
| **node vertical position** | customer cone size |
| **node size** | customer cone size |
| **link width** | customer cone size of neighbor |
| **link color** | relationship type |

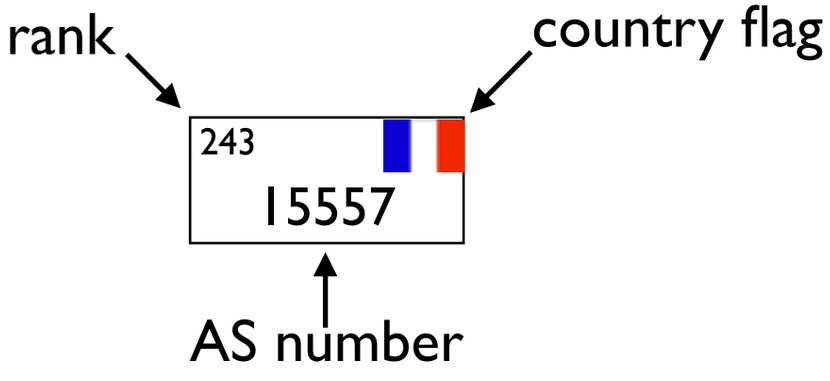


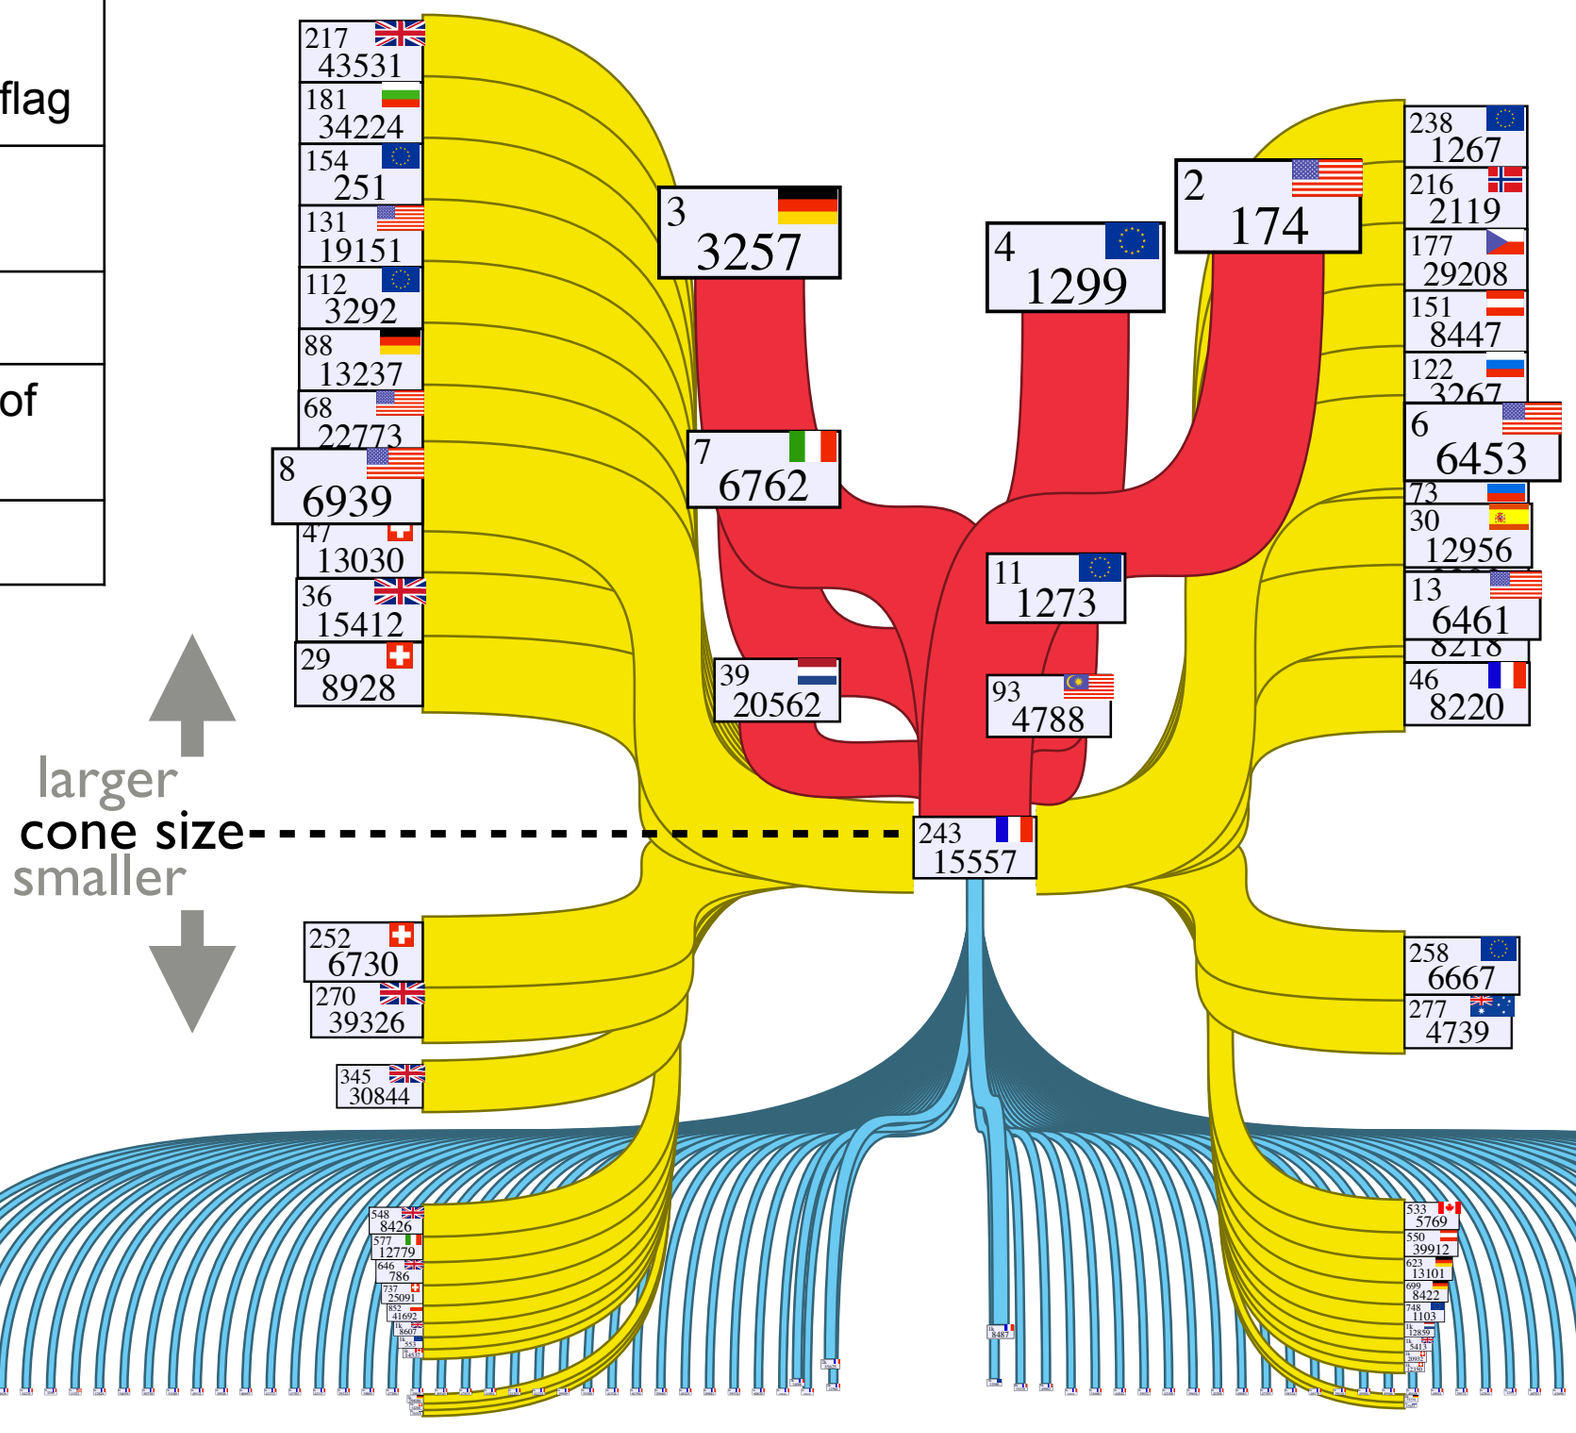

# AS core single AS

bonus

- node location and size set by customer cone size
- node and link color set by relationship to selected AS
- selected AS size is fixed regardless of degree.
- draw order: customers, peers, and providers

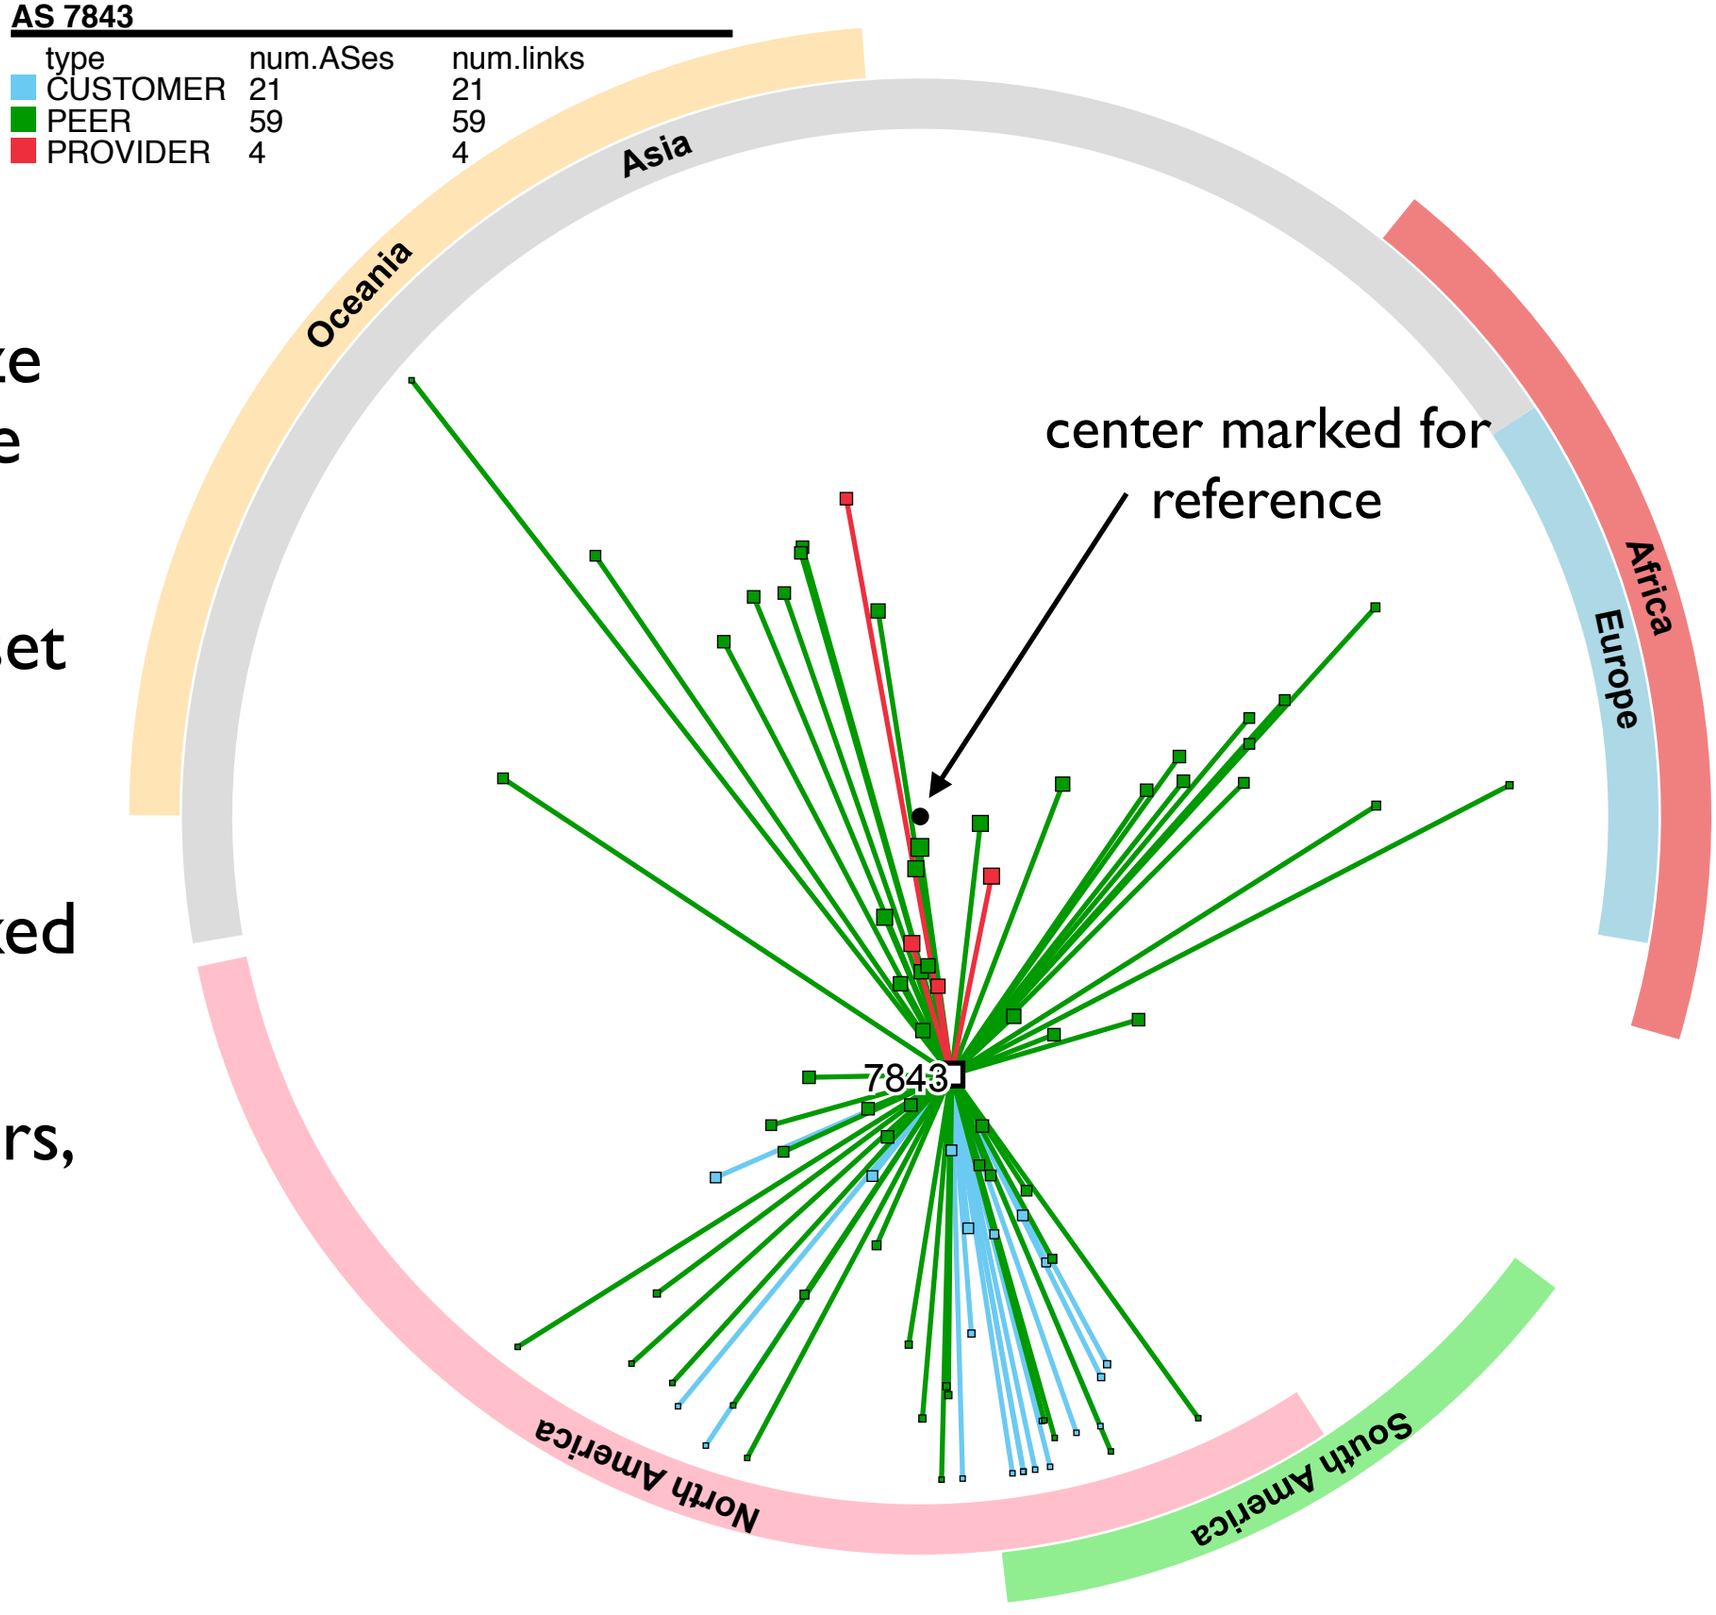

# anycast

caida

## bonus

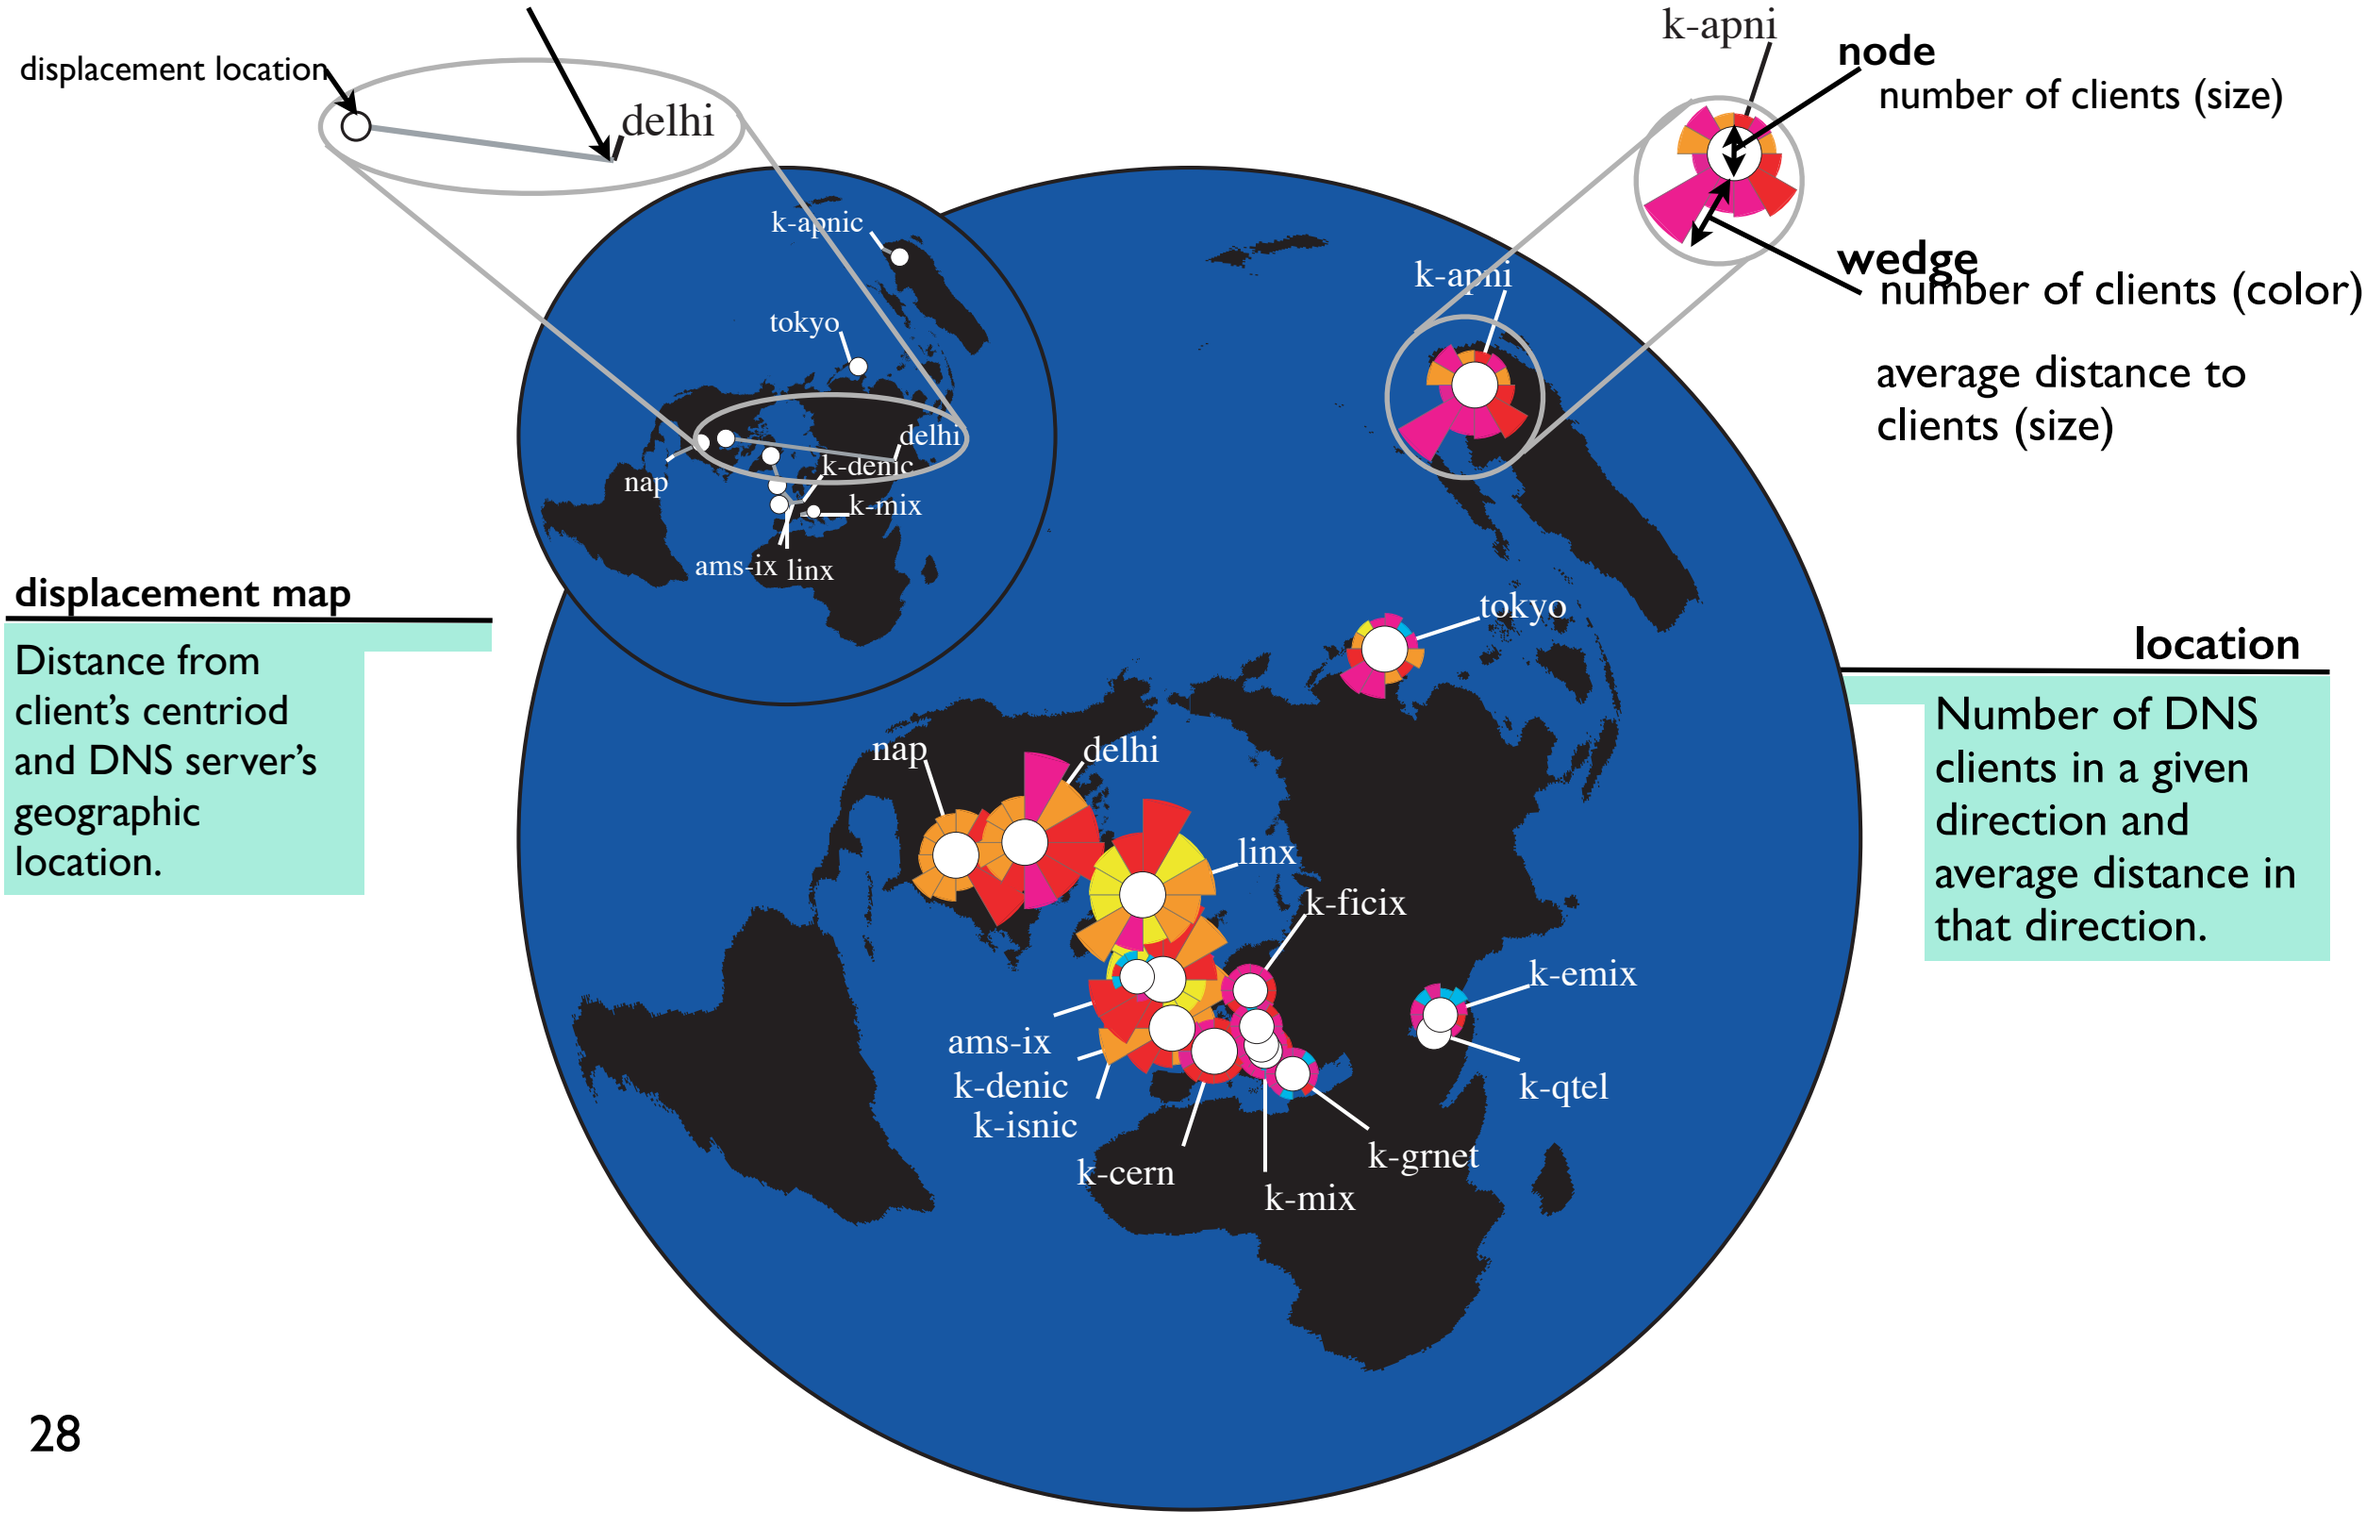


**displacement map**

Distance from client's centriod and DNS server's geographic location.

**location**

Number of DNS clients in a given direction and average distance in that direction.

# Geographic sorted bars

bonus

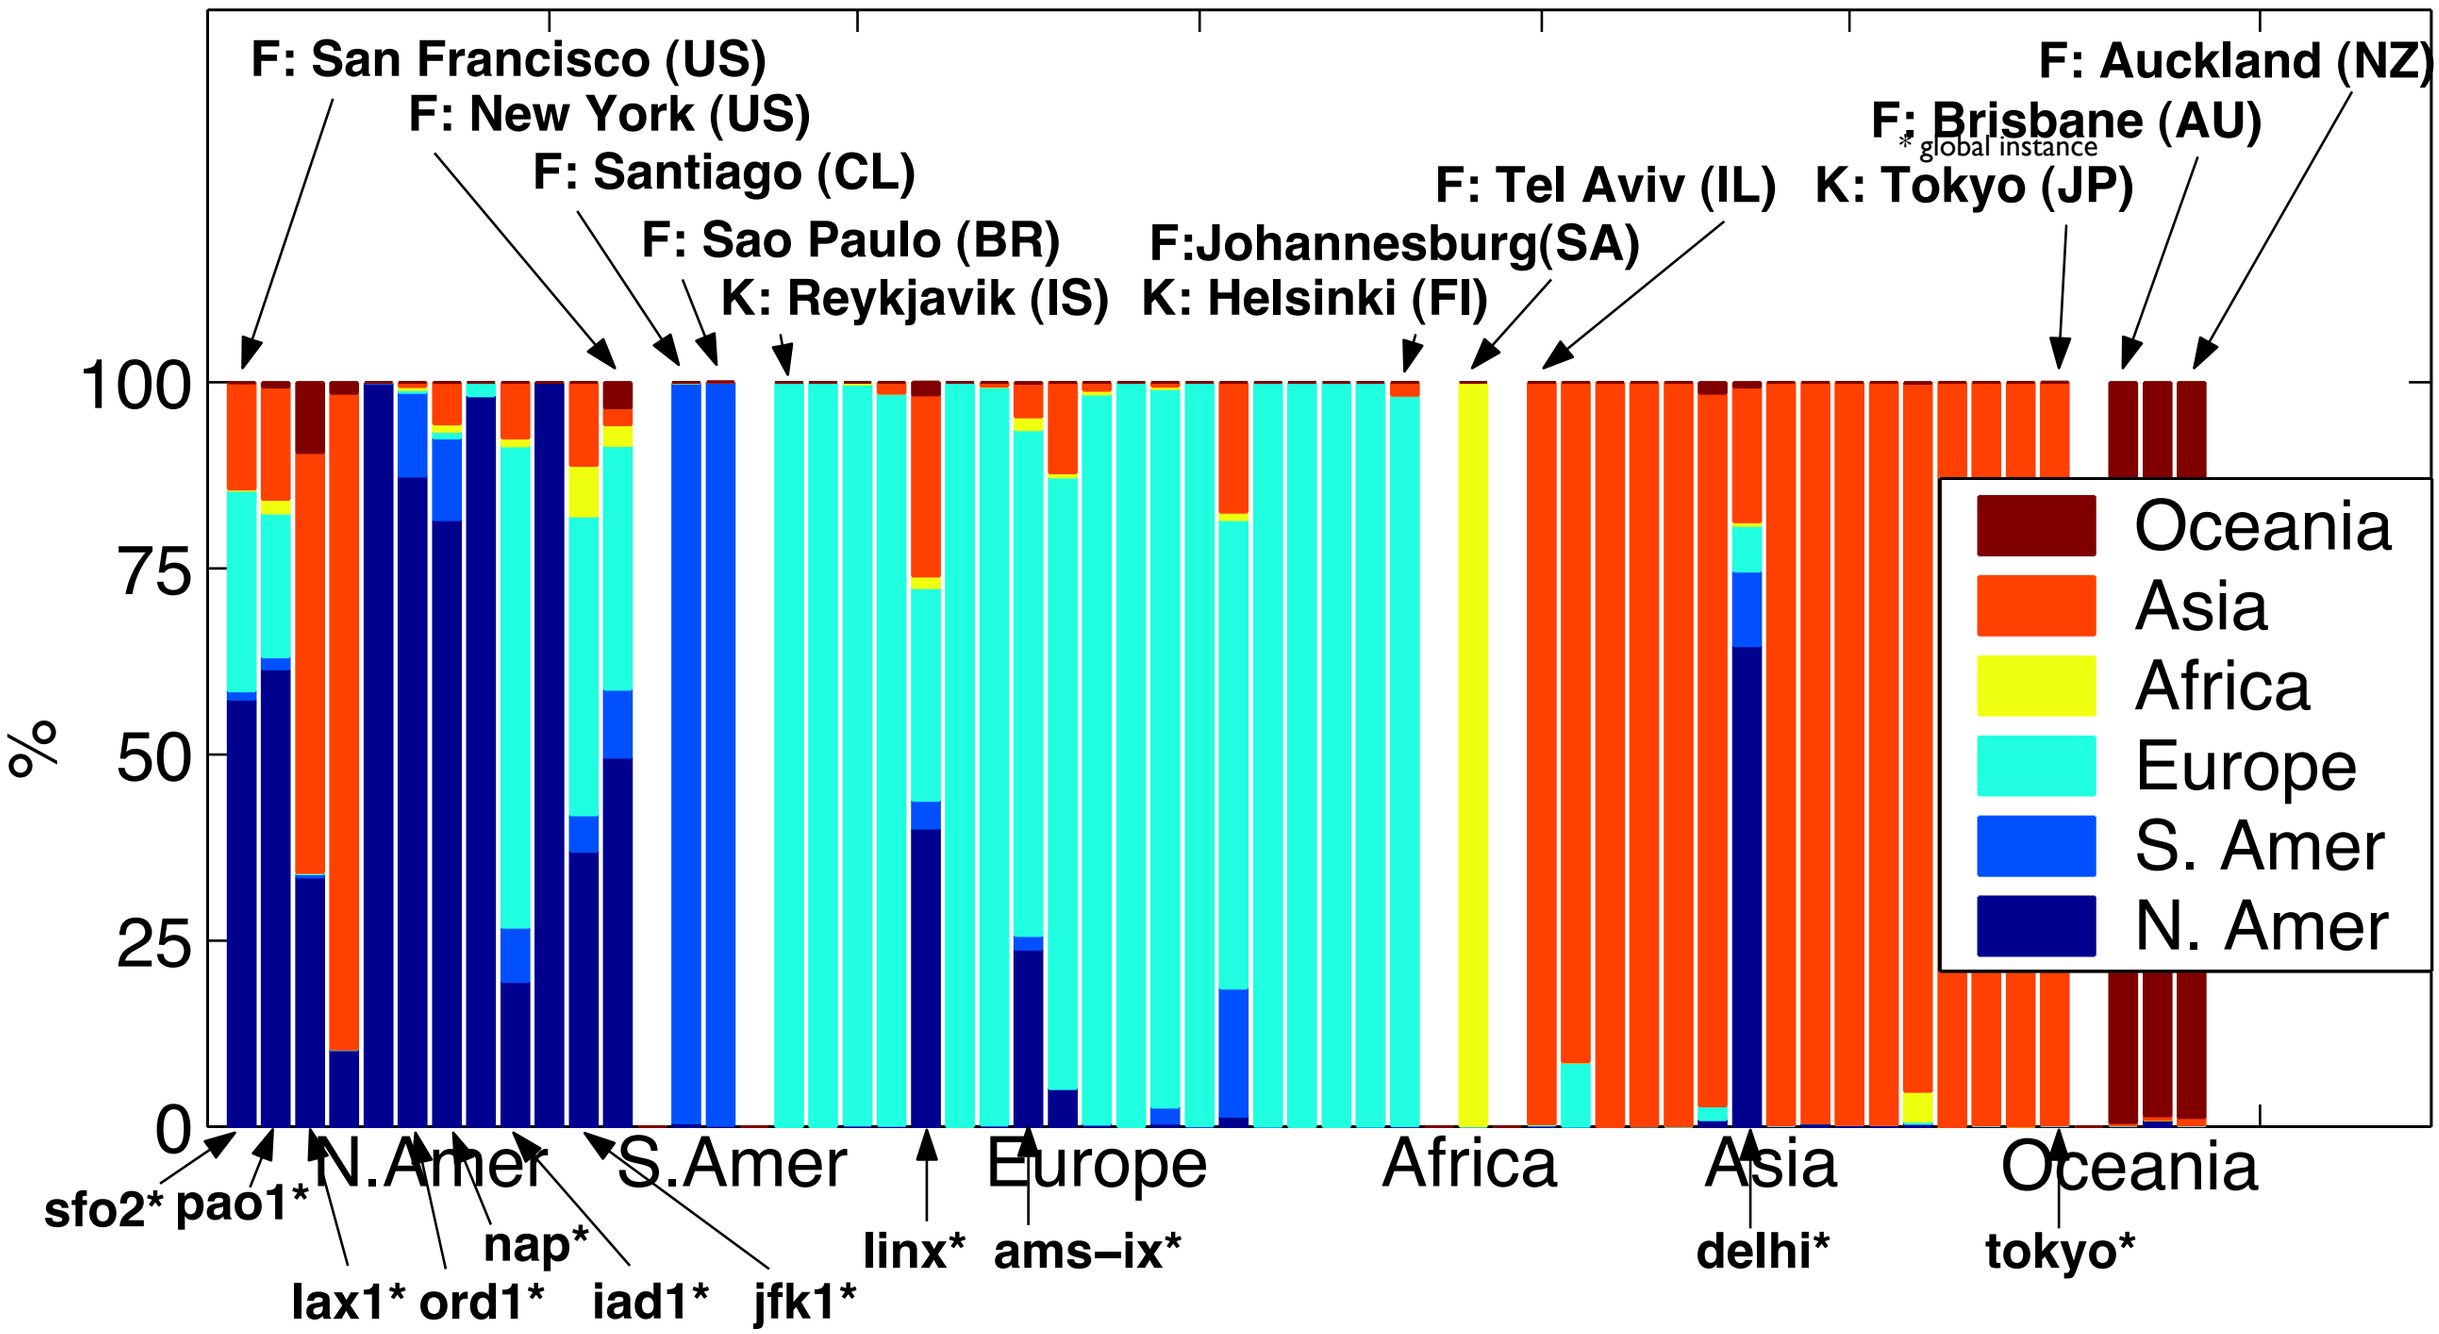

# used address space

bonus

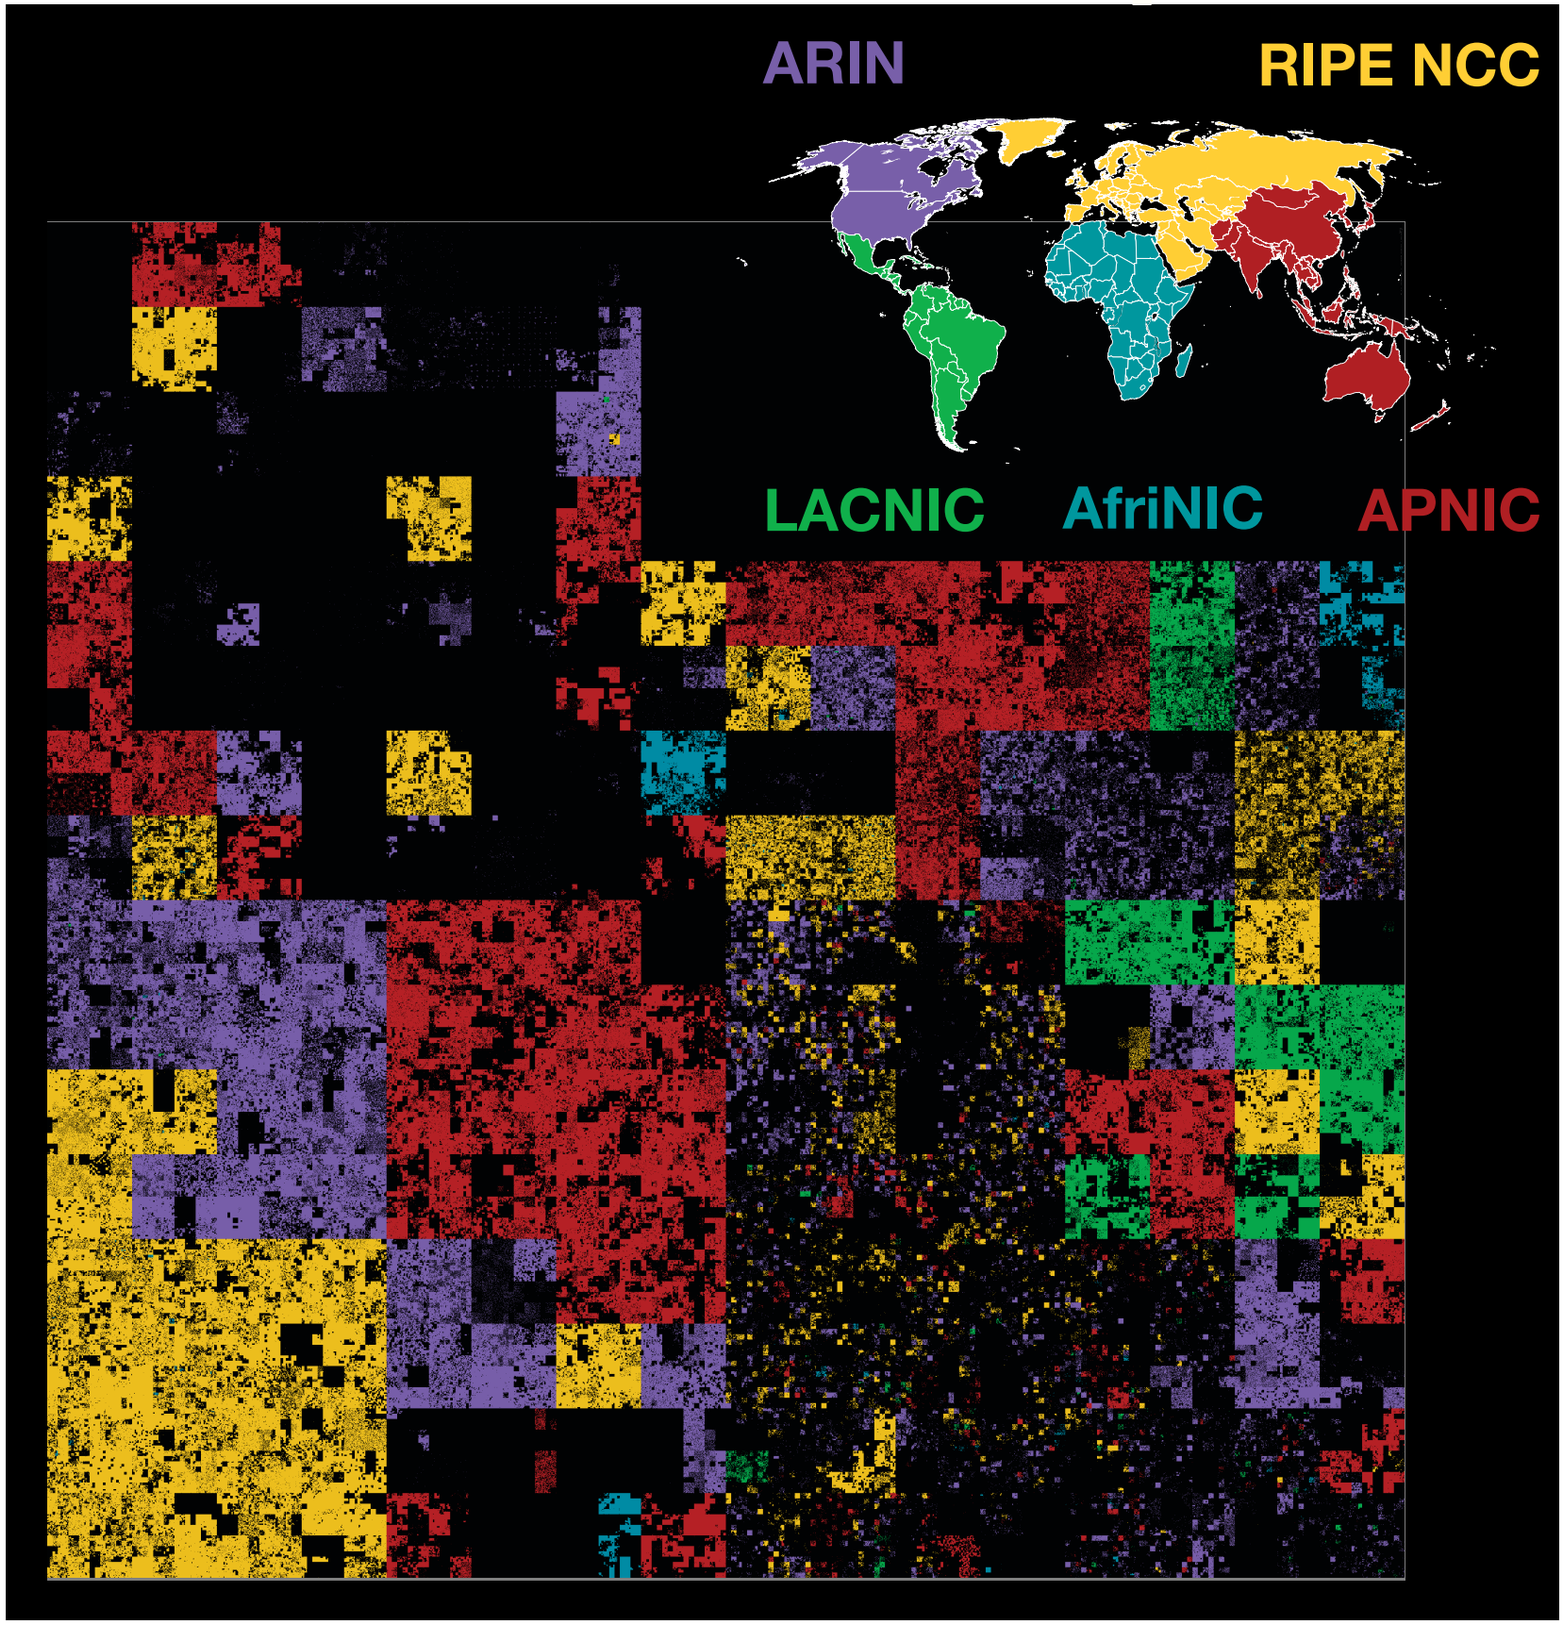

# Code red

bonus