



RIPE NCC

RIPE NETWORK COORDINATION CENTRE

OpenIPmap Geolocating Internet Infra-Structure with Inference Engines and Crowdsourcing

Jasper den Hertog
Research and Development
RIPE NCC

9 March 2018 | DKNOG8 | København



**But now for something
completely different first**

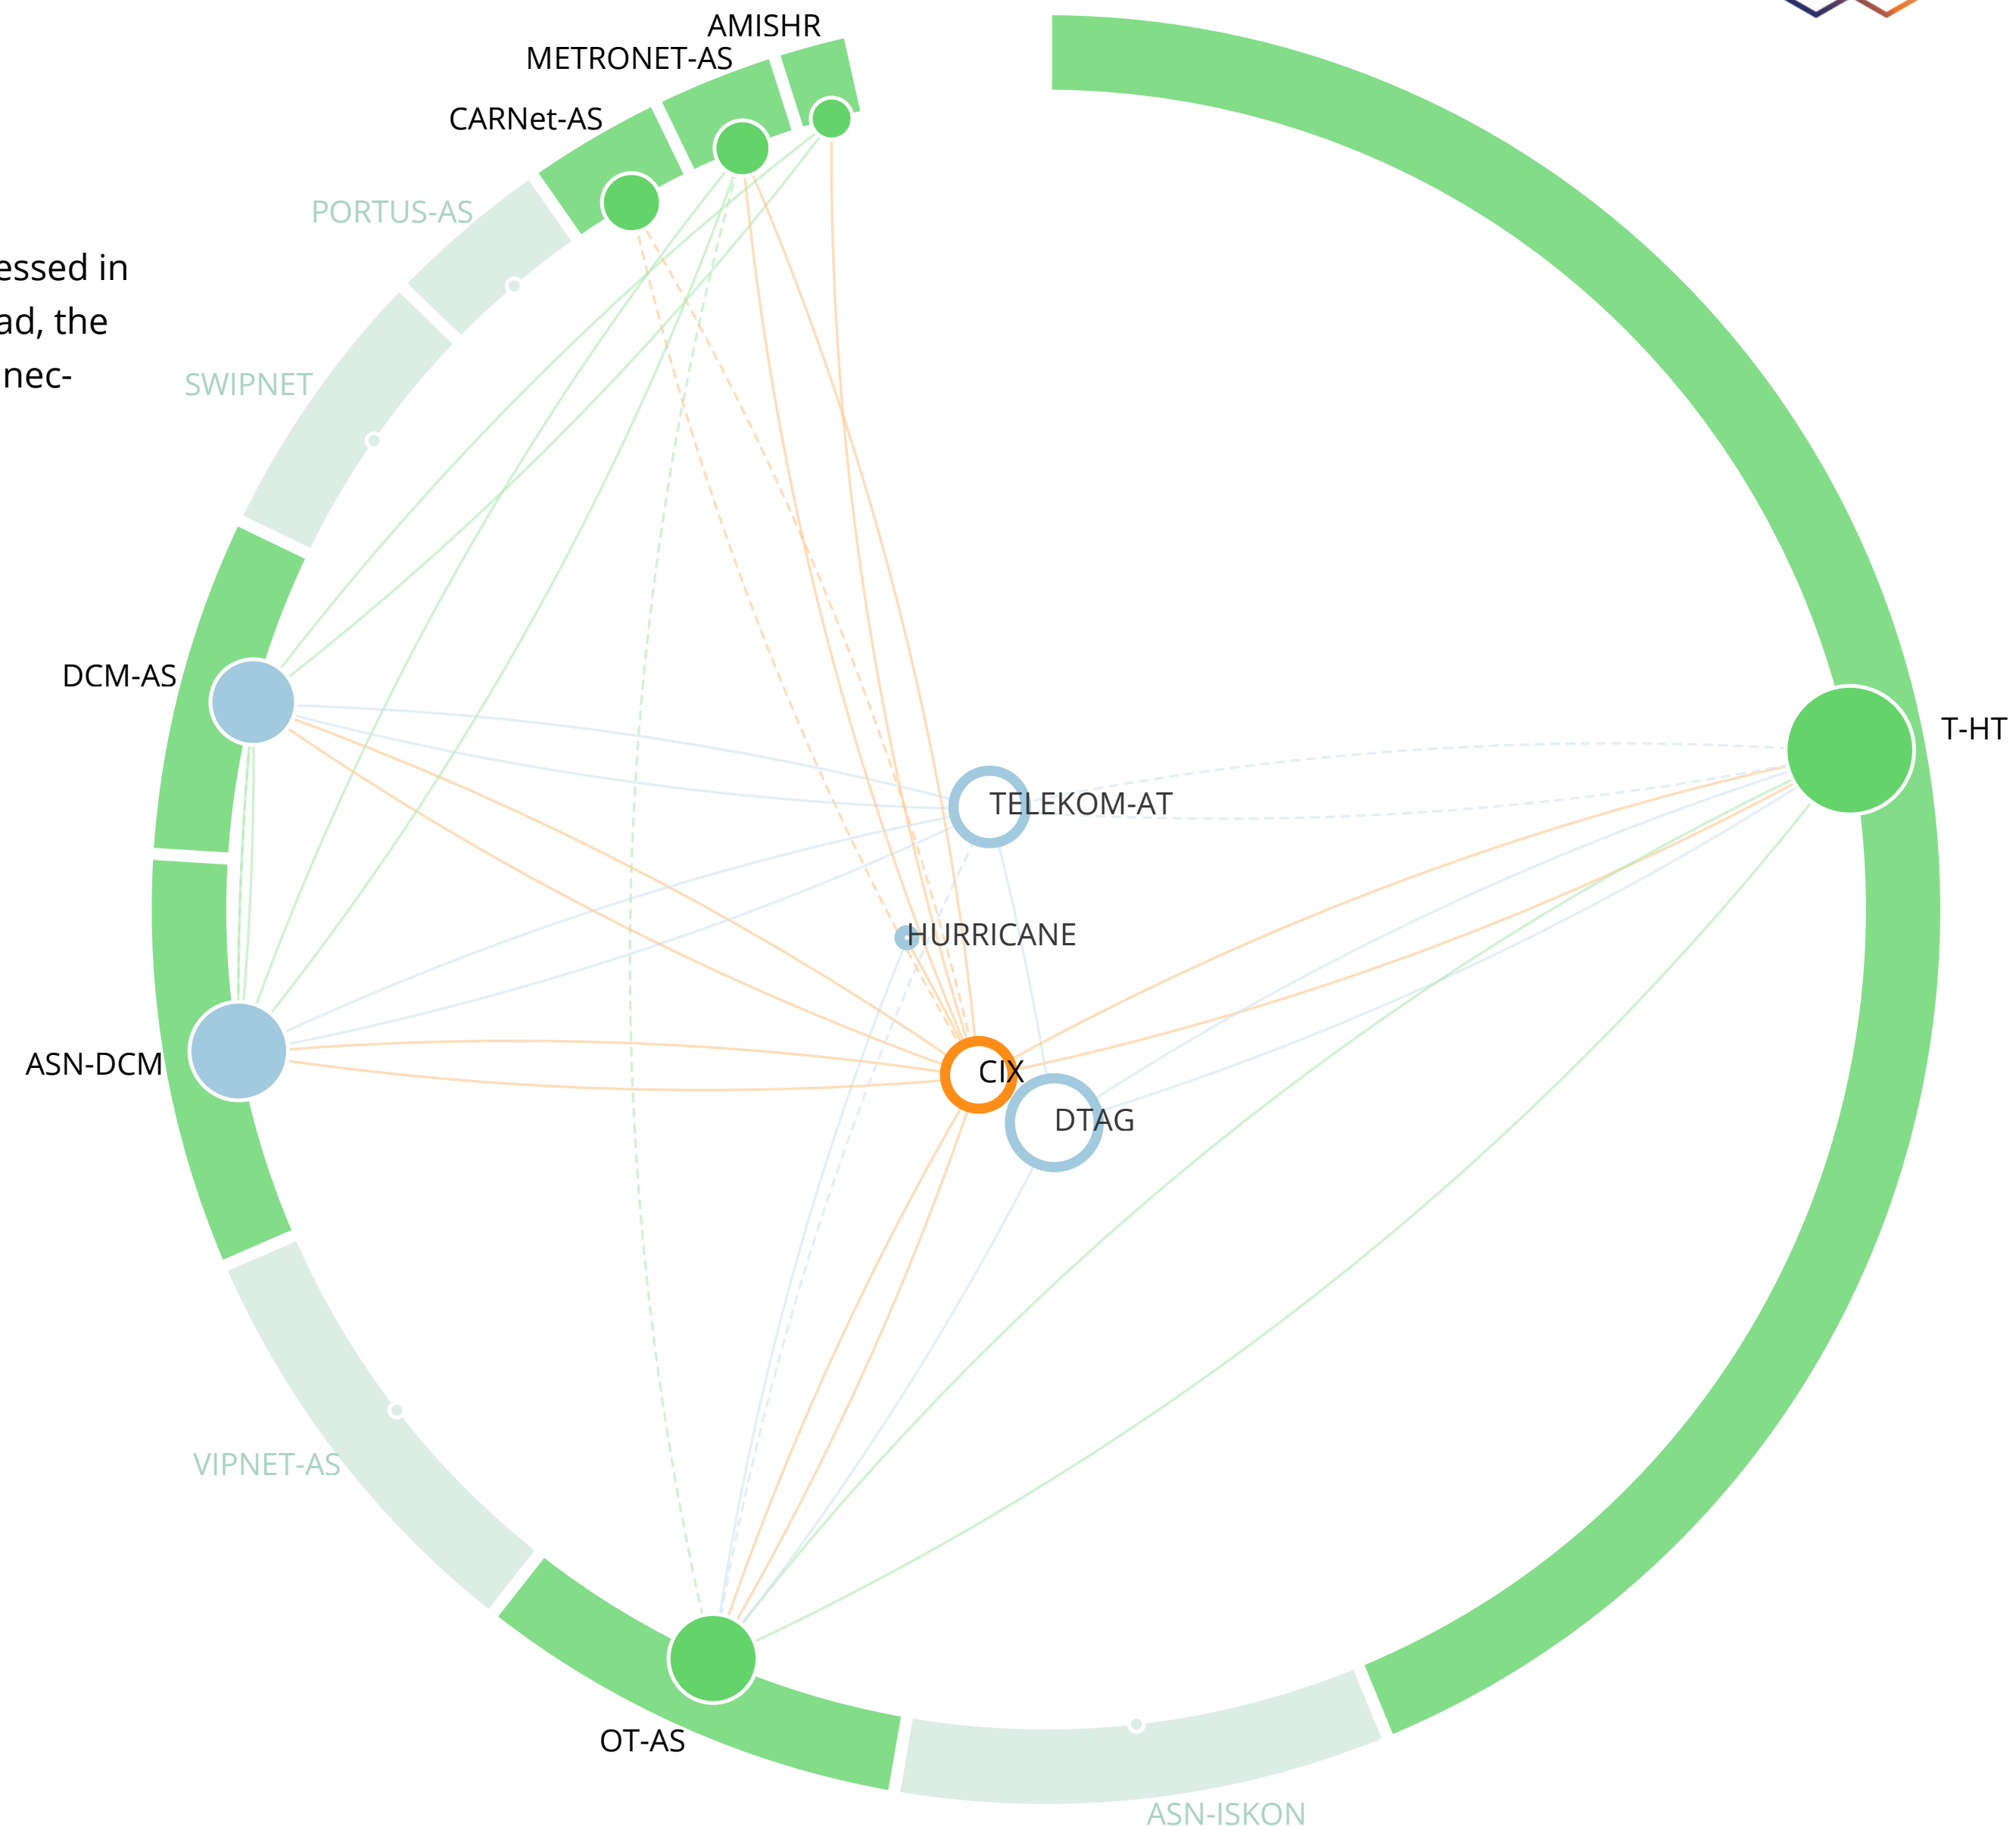
Sketches of the Peer-to-Peer Fabric of a Country



The quality of end-user connections are often expressed in download speeds towards content providers. Instead, the sketches presented here focus on peer-to-peer connections in a country.

The sketches explore the different ways in which end-users are interconnected within the same country: the peer-to-peer fabric. Each sketch represents a snapshot of this fabric at a single given point in time. They try to put a number on the amount of different ways the networks interconnect their users.


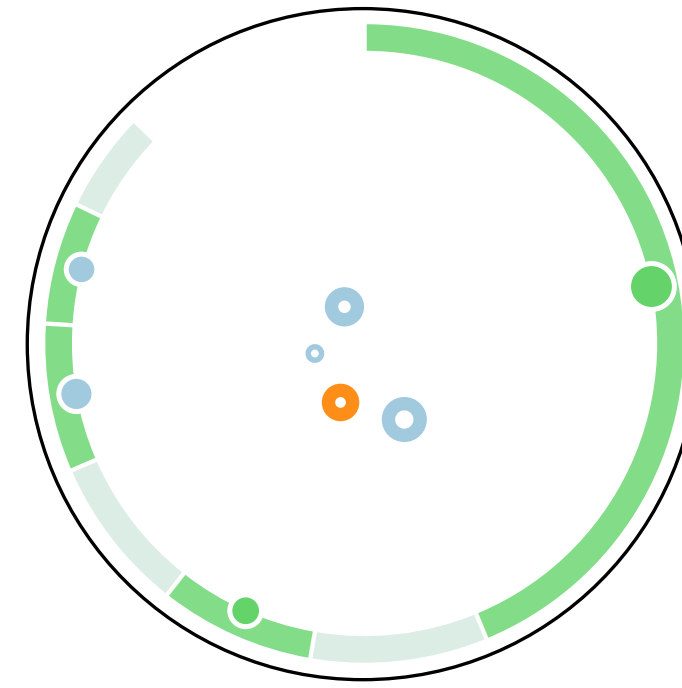
These sketches are created with active measurements from the RIPE Atlas measurement platform, datasets from RIPEstat, AS-to-ORG datasets from CAIDA and a dataset from APNIC that estimates the percentage of end-users in each network.



Sketches of the Peer-to-Peer Fabric of a Country




The full circle represents 100% of the end-users in a country.



Each network that provides connectivity to more than 1% of the end-users is represented by a colored circle segment. The length of the arc of the segment represents the percentage of the end-users in a country.

The darker green denotes an network for which we have peer-to-peer data. The lighter green color denotes networks for which we don't have peer-to-peer data.

The open part of the circle represents the sum of all ASes that provide connectivity to less than 1% of the end-users in a country.




Sketches of the Peer-to-Peer Fabric of a Country




Each ring or circle represents the percentage of the peer-to-peer fabric in a country that passes through this point.

The color of the circle or ring denotes the type of location.




A blue circle on the outer ring represents a network that both serves end-users and provides transit to others end-user networks within the country.



A green circle on the outer ring represents a network that (mainly) serves end-users.

A blue circle in the interior indicates a transit network or an IXP that is external to this country.




An orange circle in the interior indicates an IXP identified with this country.


Sketches of the Peer-to-Peer Fabric of a Country




Orange lines indicate that two end-user networks are connected through an IXP.



Green lines indicate that two end-user networks are directly connected.



Blue lines indicate two end-user networks are connected through a transit network.



Dotted lines of any color indicate that we cannot fully map this path.

Peer-to-Peer Fabric



country

Denmark

snapshot date







1 March 2018

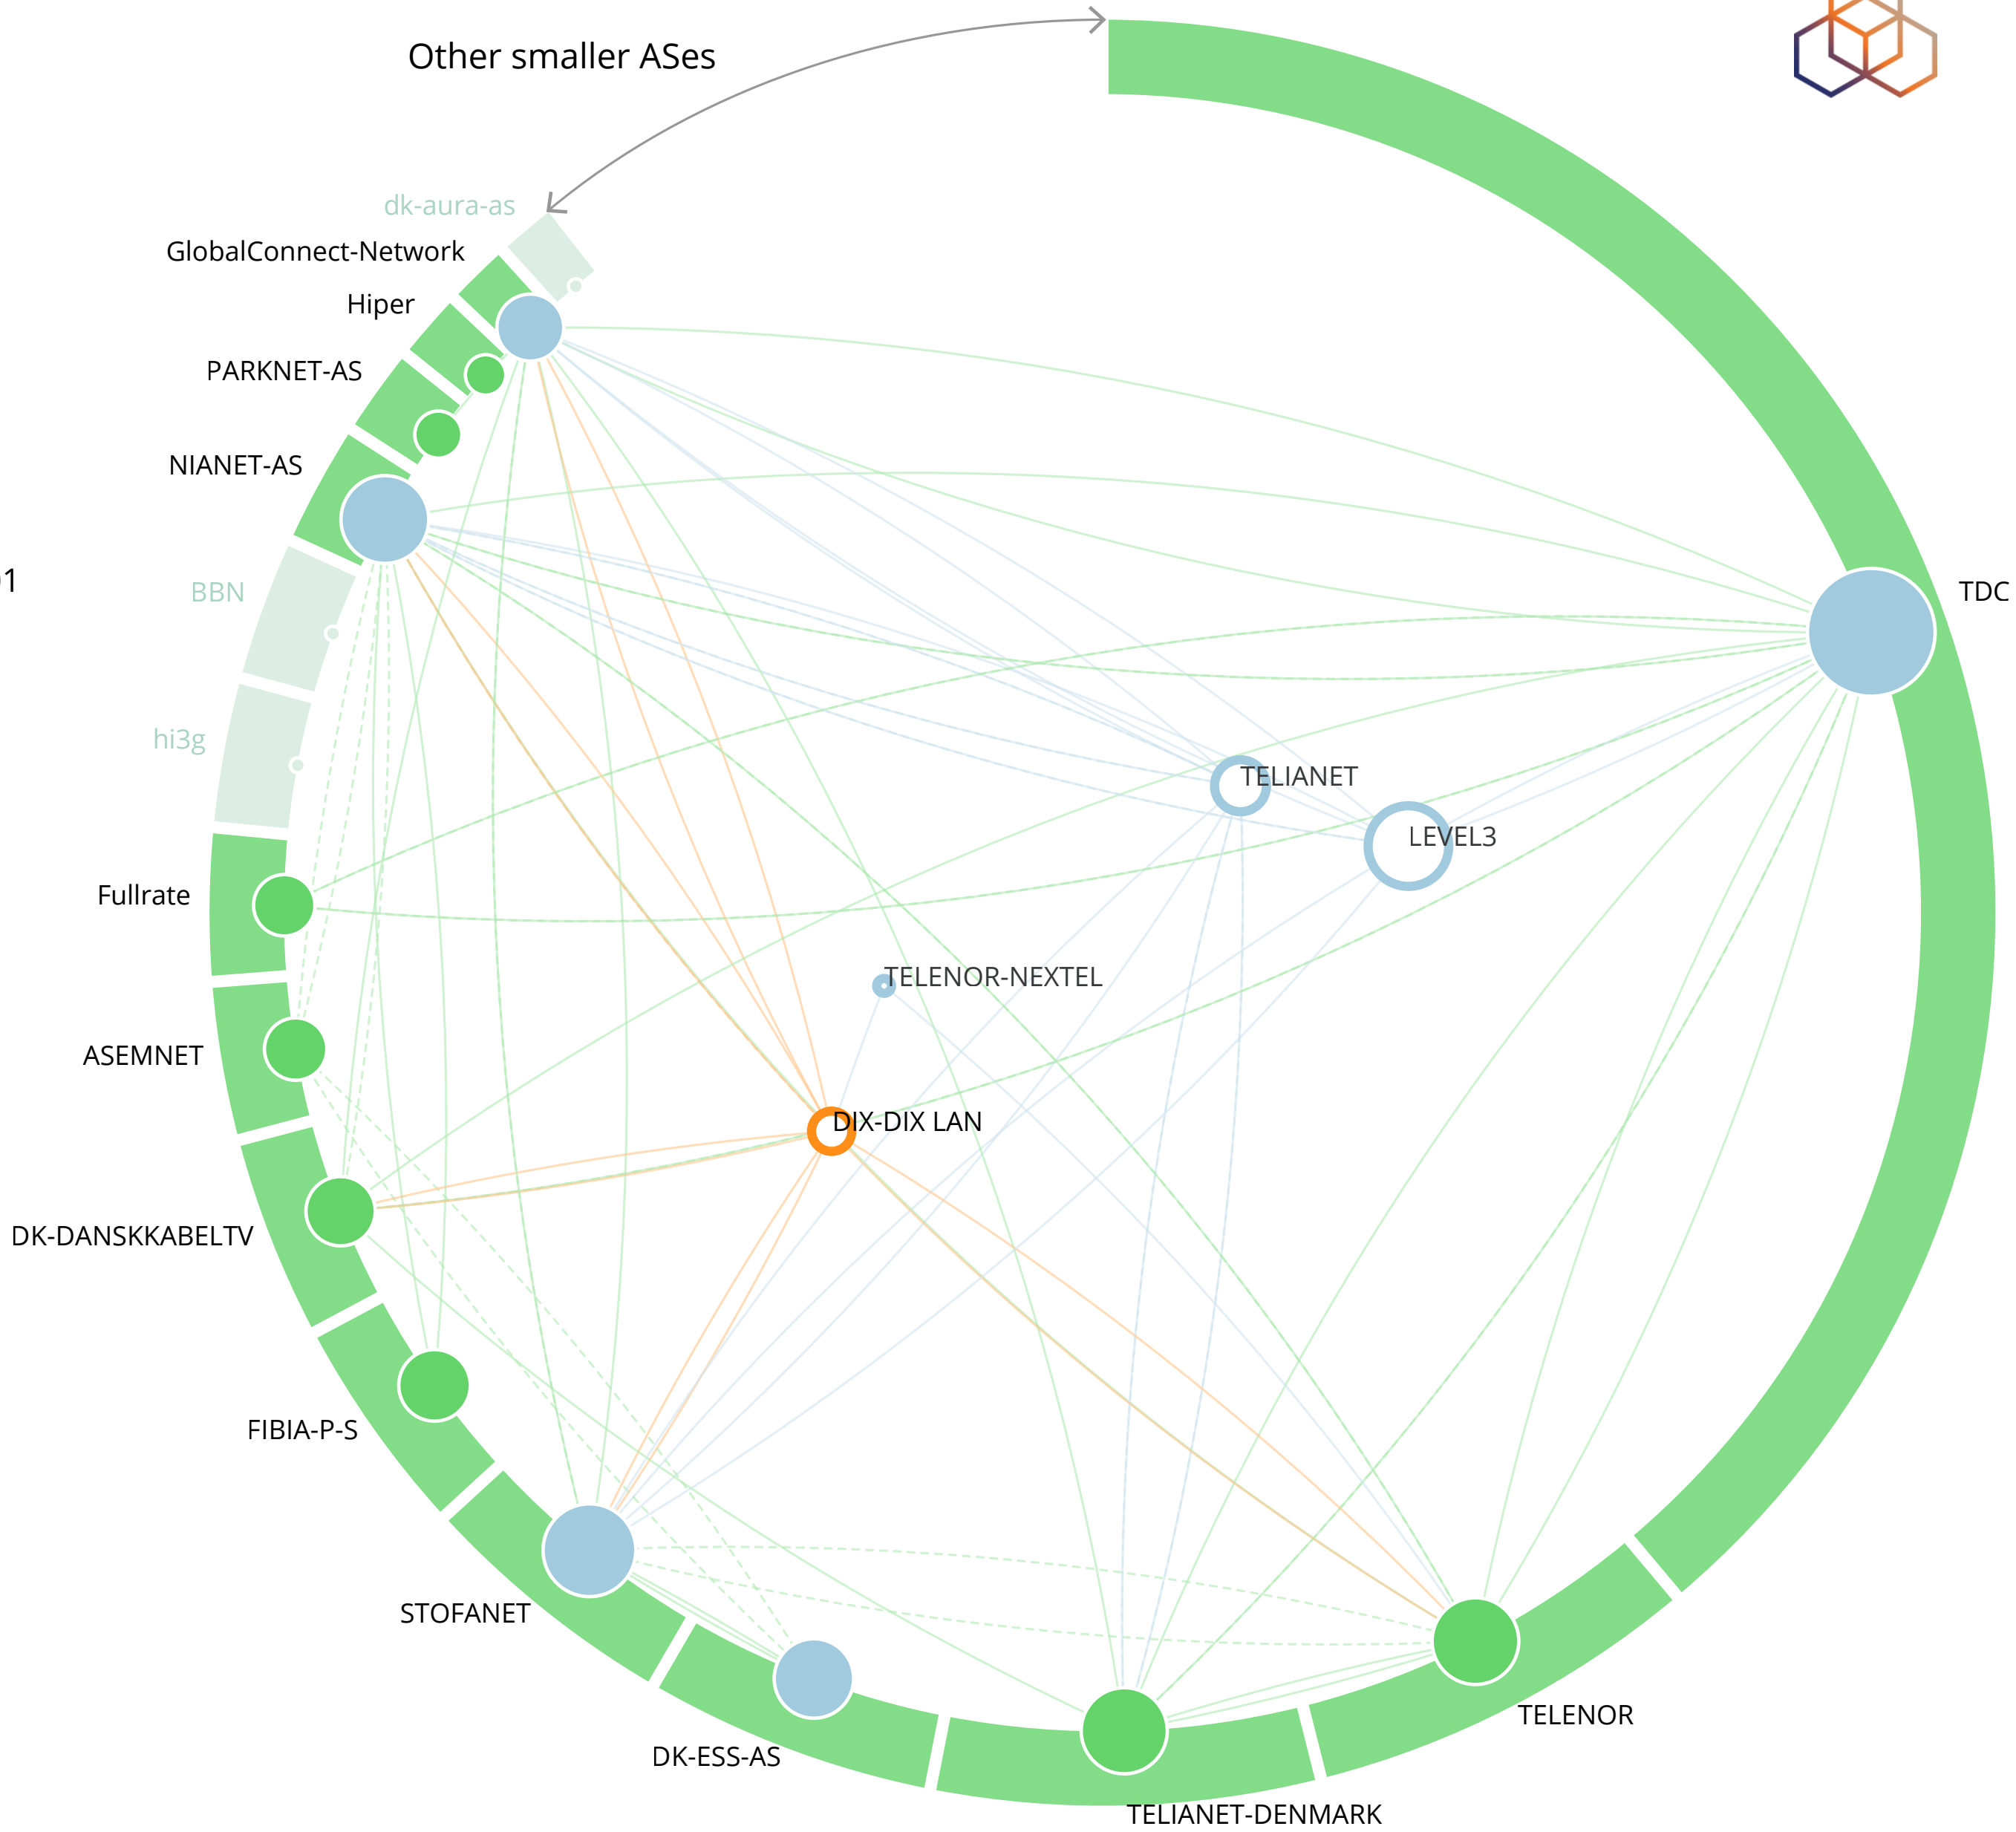
github

<https://github.com/emileaben/ixp-country-jedi/>

url

<http://sg-pub.ripe.net/ixp-country-jedi/dk/2018/03/01>

-  A network that serves end-users
-  A network that serves end-users and provides transit to other end-user networks within the country
-  A transit network or an IXP external to this country
-  An IXP that is identified with this country
-  A sizable end-user network for which we have data
-  A sizable end-user network for which we have no data



Peer-to-Peer Fabric

country

Southern Korea

snapshot date

1 March 2018







github

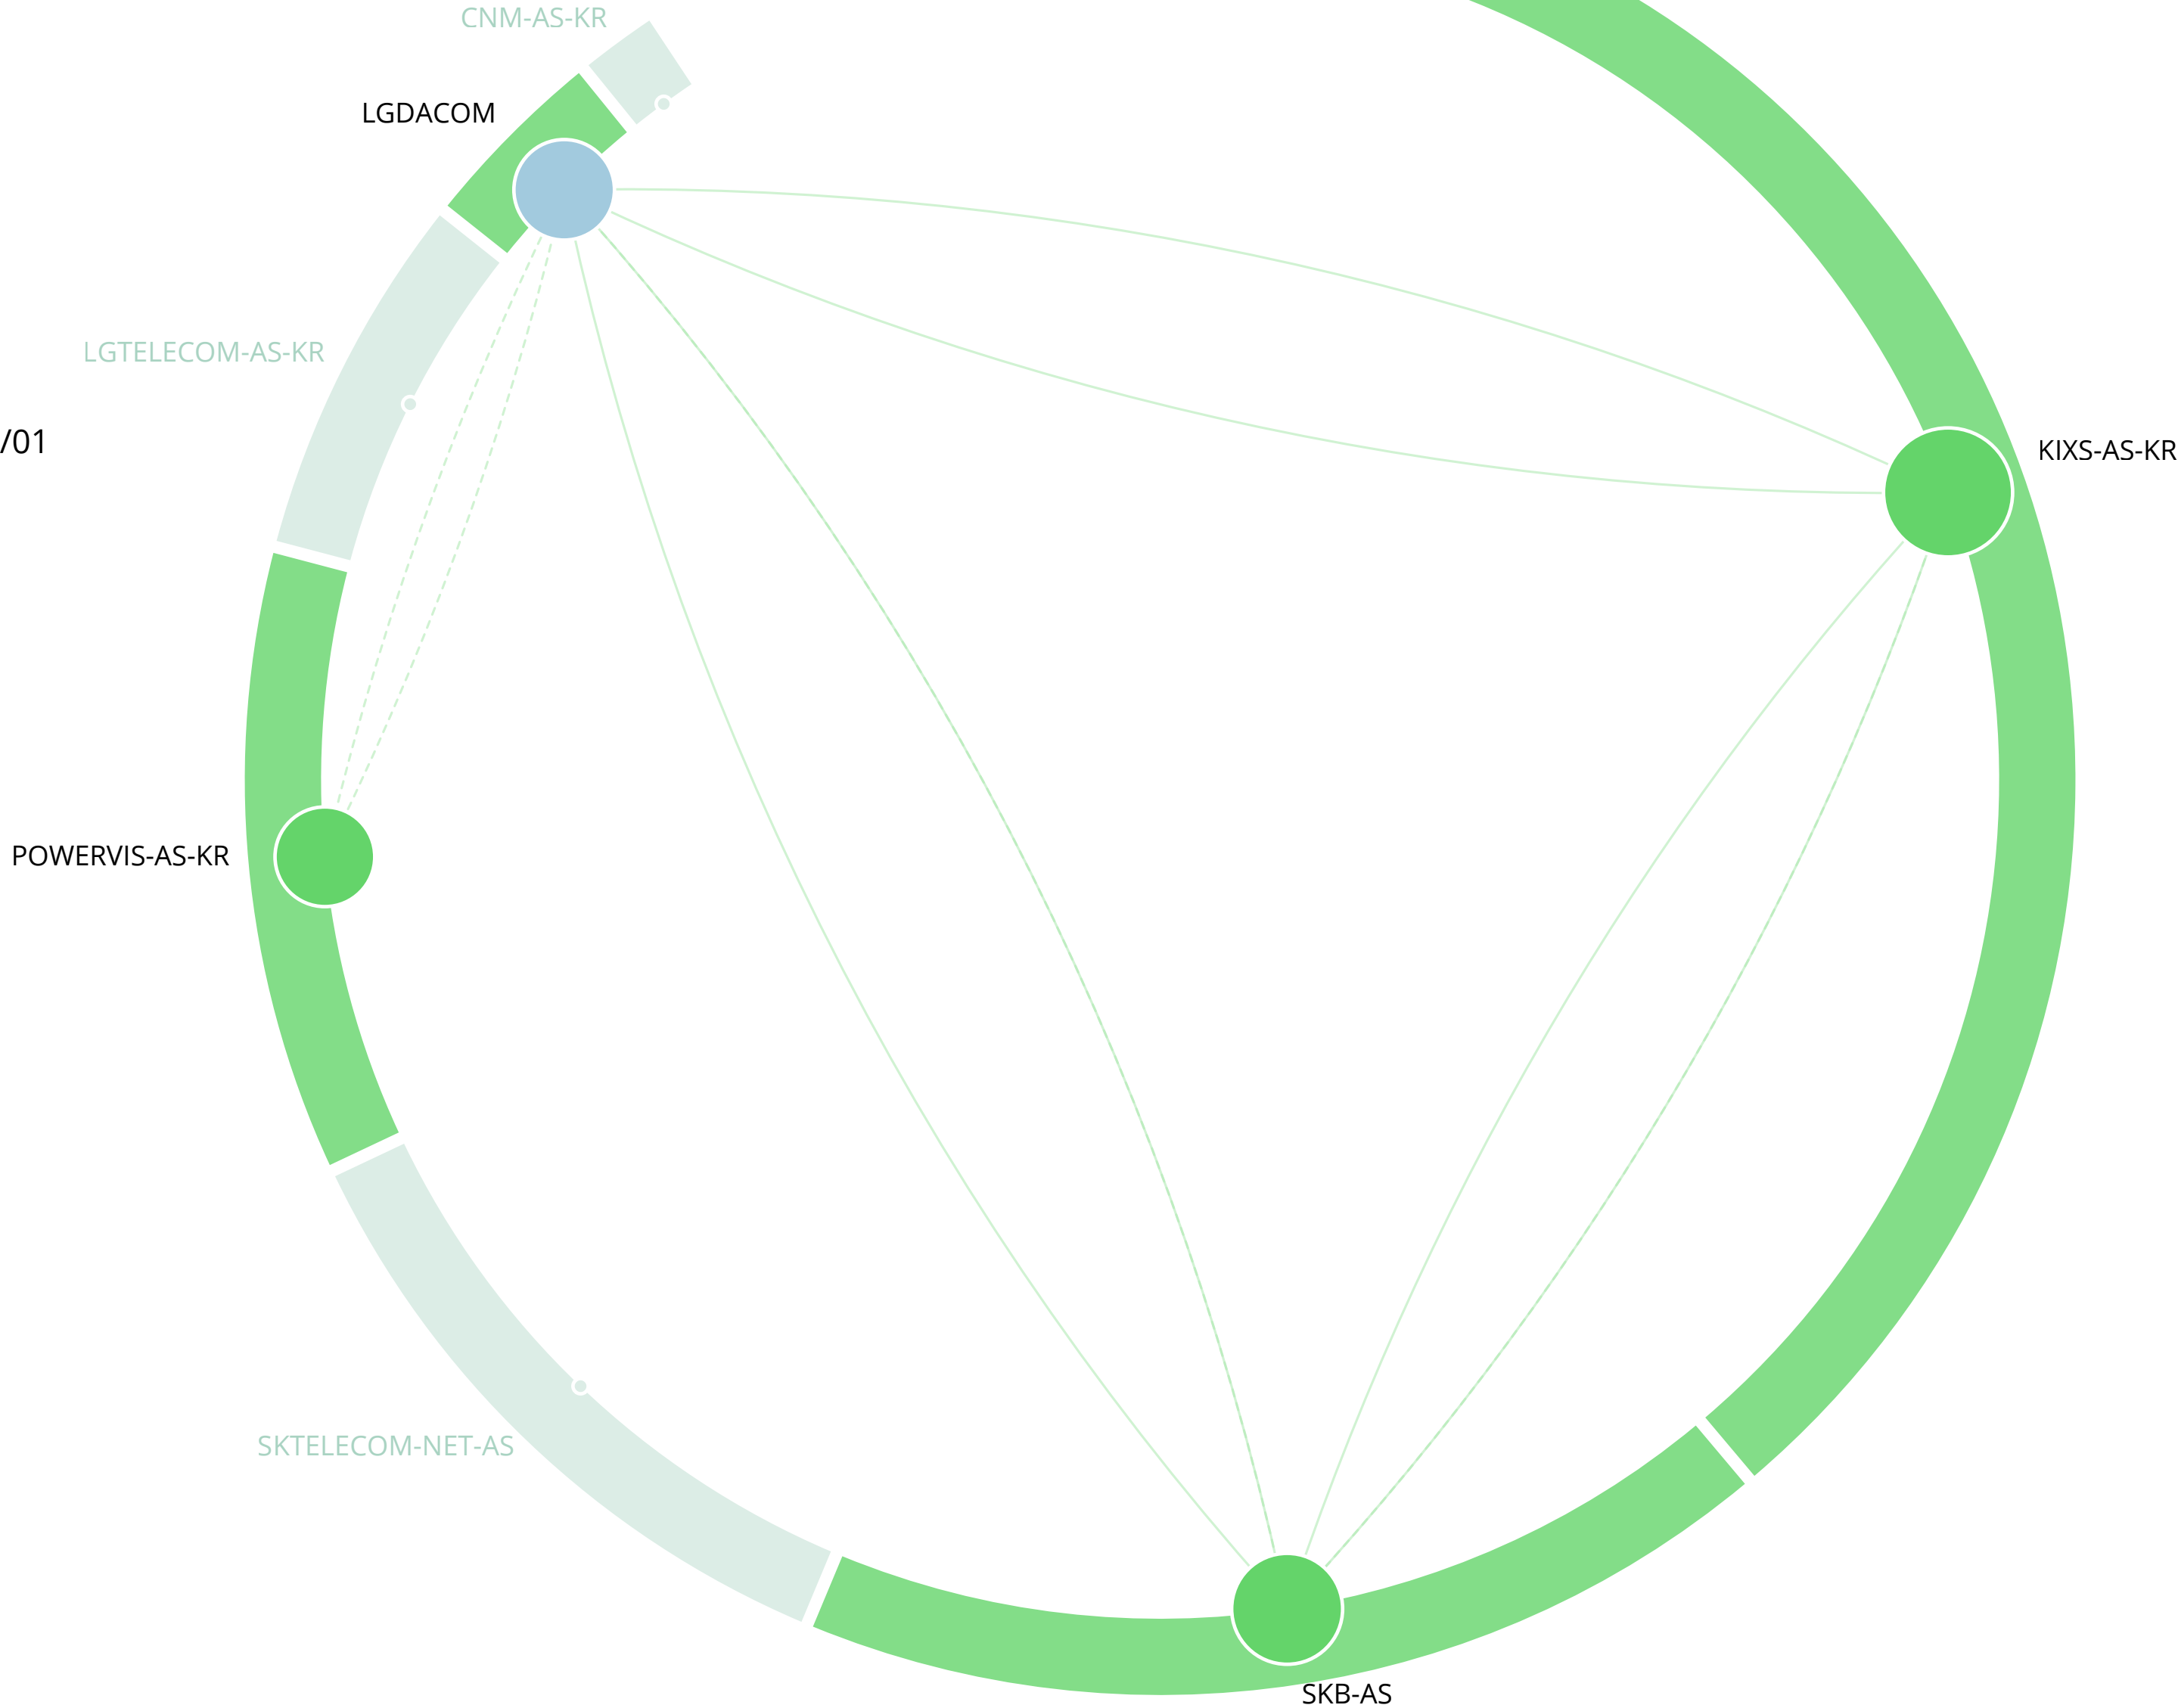
<https://github.com/emileaben/ixp-country-jedi/>

url

<http://sg-pub.ripe.net/ixp-country-jedi/dk/2018/03/01>



-  A network that serves end-users
-  A network that serves end-users and provides transit to other end-user networks within the country
-  A transit network or an IXP external to this country
-  An IXP that is identified with this country
-  A sizable end-user network for which we have data
-  A sizable end-user network for which we have no data



Peer-to-Peer Fabric

country

Ireland

snapshot date

1 March 2018

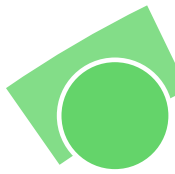
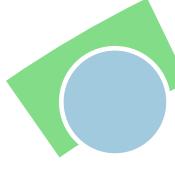


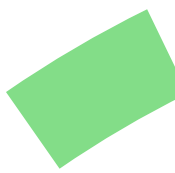

github

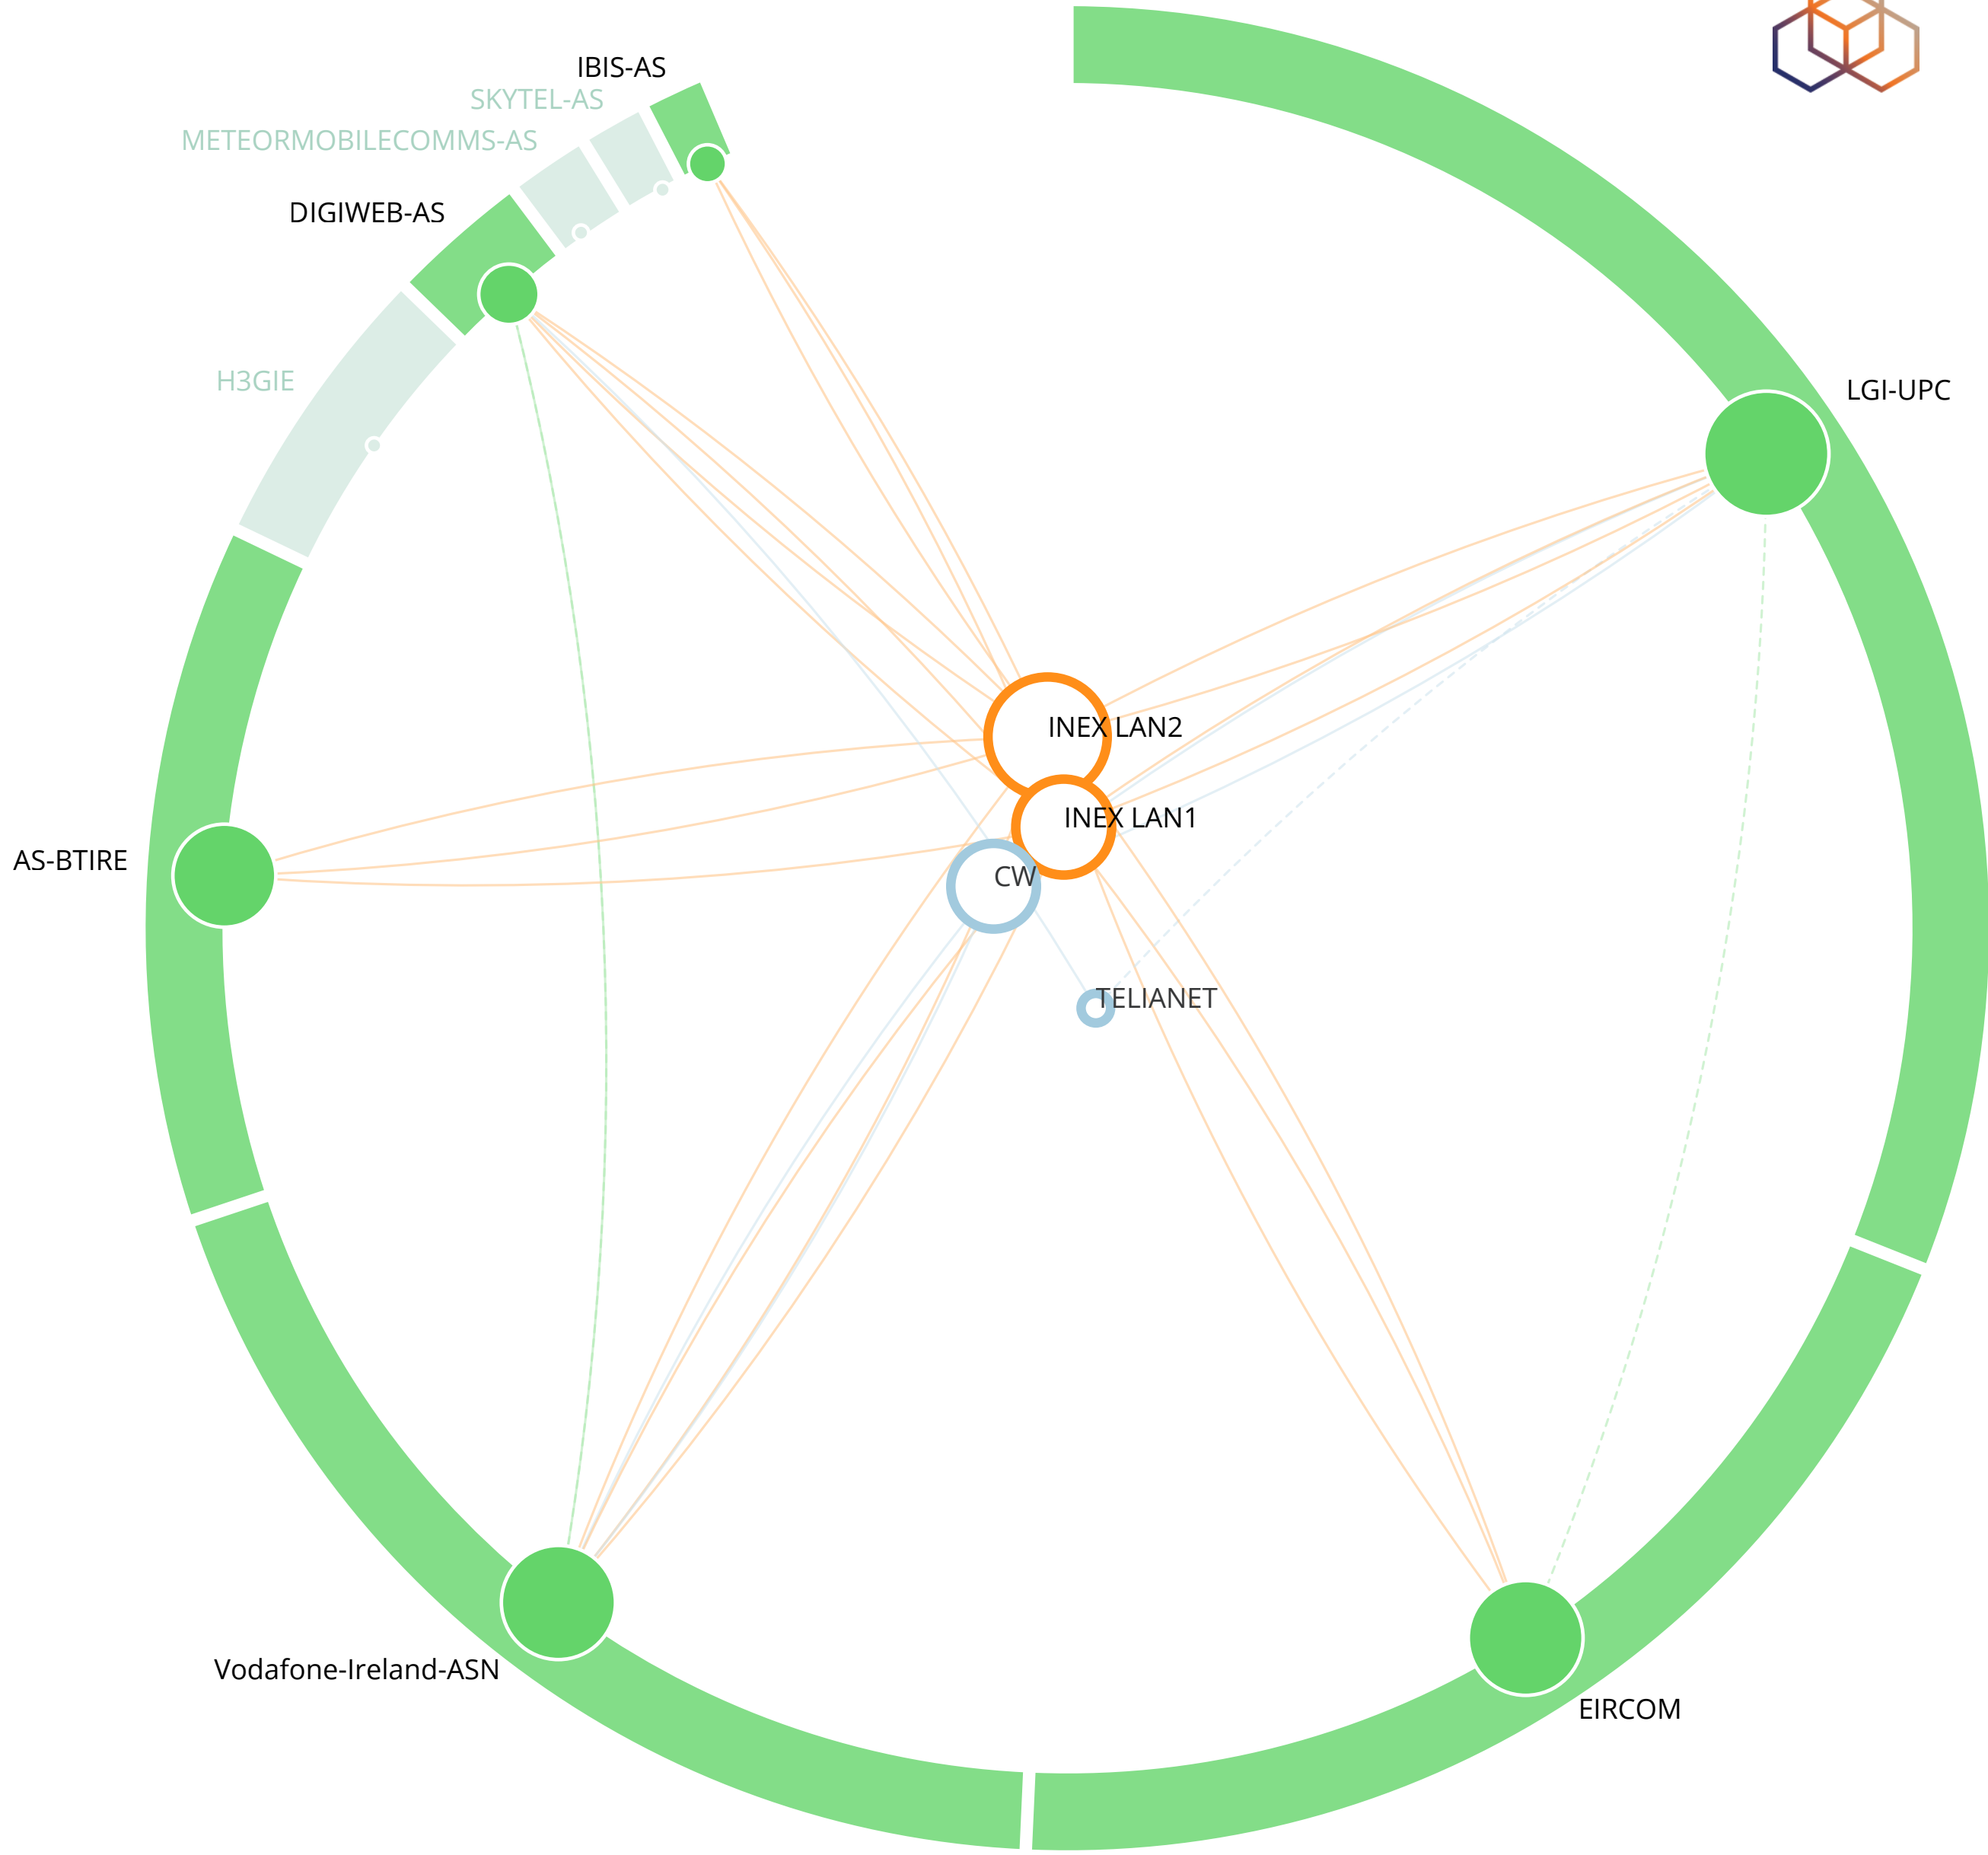
<https://github.com/emileaben/ixp-country-jedi/>

url

<http://sg-pub.ripe.net/ixp-country-jedi/dk/2018/03/01>



-  A network that serves end-users
-  A network that serves end-users and provides transit to other end-user networks within the country
-  A transit network or an IXP external to this country
-  An IXP that is identified with this country
-  A sizable end-user network for which we have data
-  A sizable end-user network for which we have no data



Peer-to-Peer Fabric

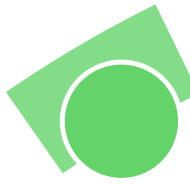
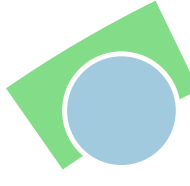

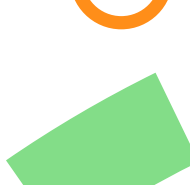




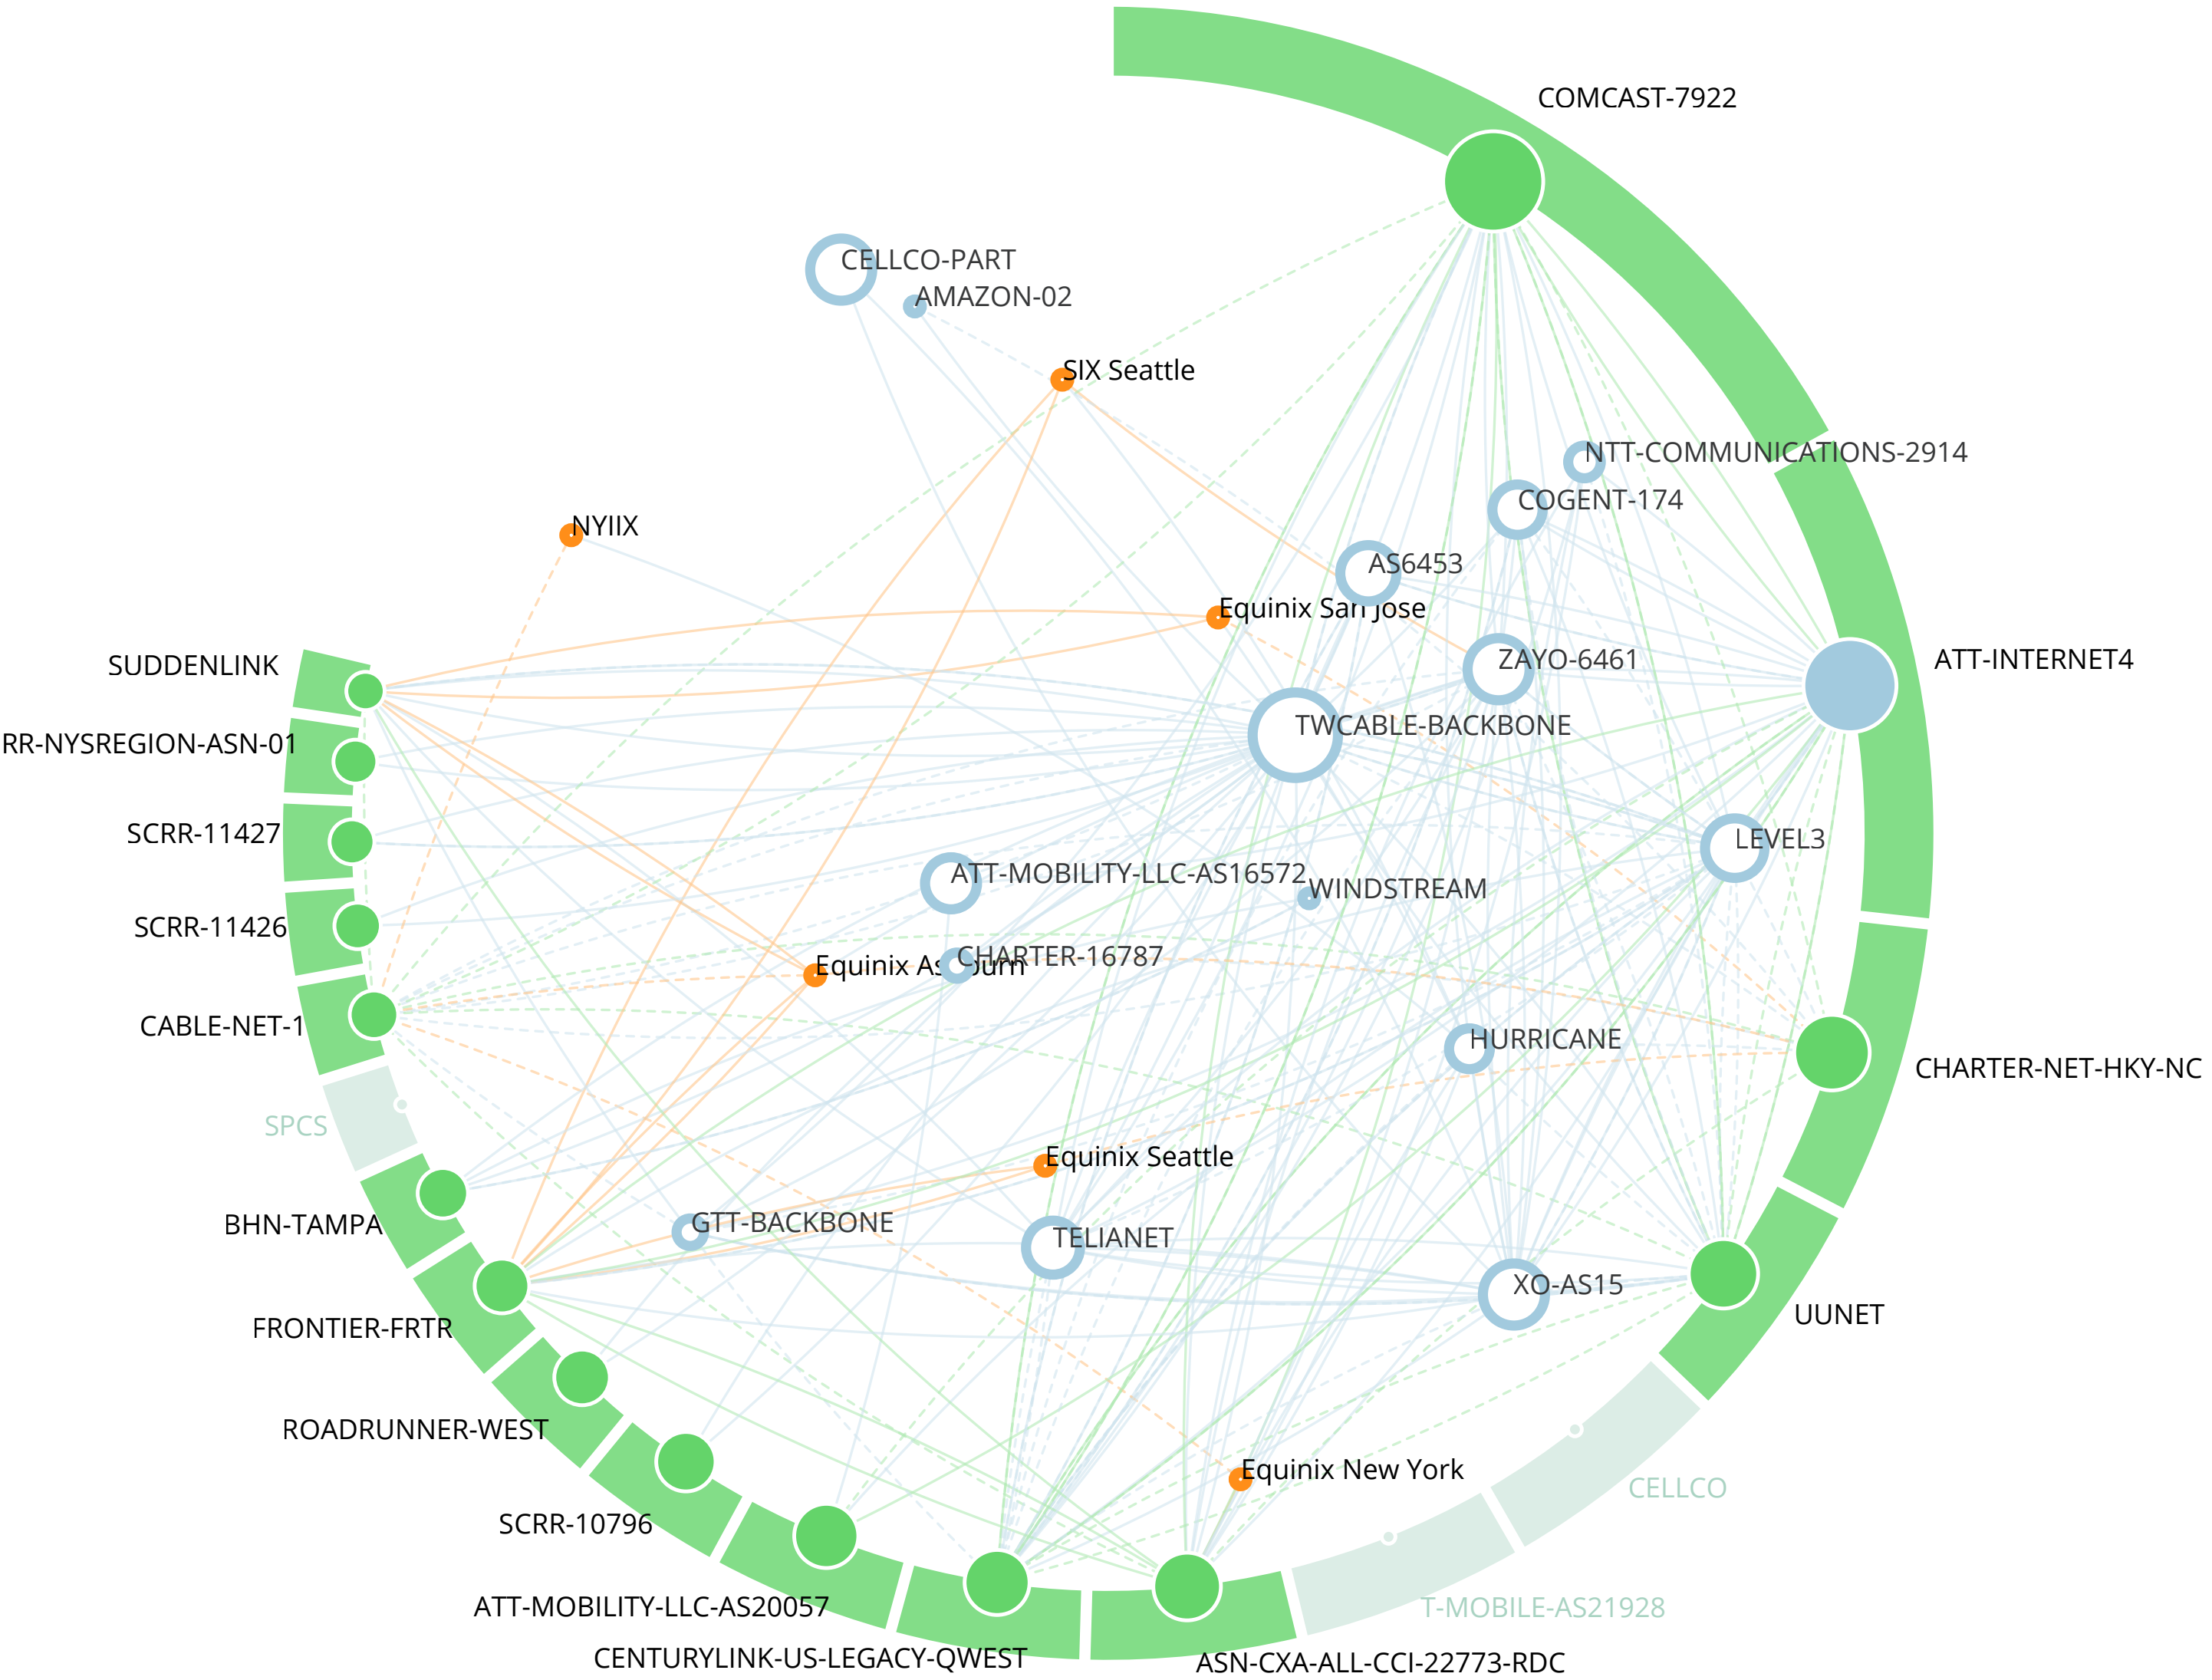
country
USA

snapshot date
1 March 2018

github
<https://github.com/emileaben/ixp-country-jedi/>

url
<http://sg-pub.ripe.net/ixp-country-jedi/dk/2018/03/01>

-  A network that serves end-users
-  A network that serves end-users and provides transit to other end-user networks within the country
-  A transit network or an IXP external to this country
-  An IXP that is identified with this country
-  A sizable end-user network for which we have data
-  A sizable end-user network for which we have no data





Why Geolocation of infra-structure?

Why?



Analyse & Optimise Paths

Why?



Coherence of Geo Information

Geolocation in the RIPE Database



“Please be aware that this information may not be picked up by geolocation providers.”

...

“Please be aware that geolocation information is added by the resource holders in the RIPE Database and the RIPE NCC does not verify this information.”



The problem of Geolocation



Geolocation is hard

Different Research Approaches



- Triangulation a.k.a. trilateration
- Reverse DNS based location inference
- ‘Administrative’ analyses
- Verification/falsification procedures

Commercial Offerings



- Tend to concentrate on end user IP Addresses
- Opaque Methodology
- IPv6 address space largely ignored



Our Integration Attempt

inference engines and crowdsourcing



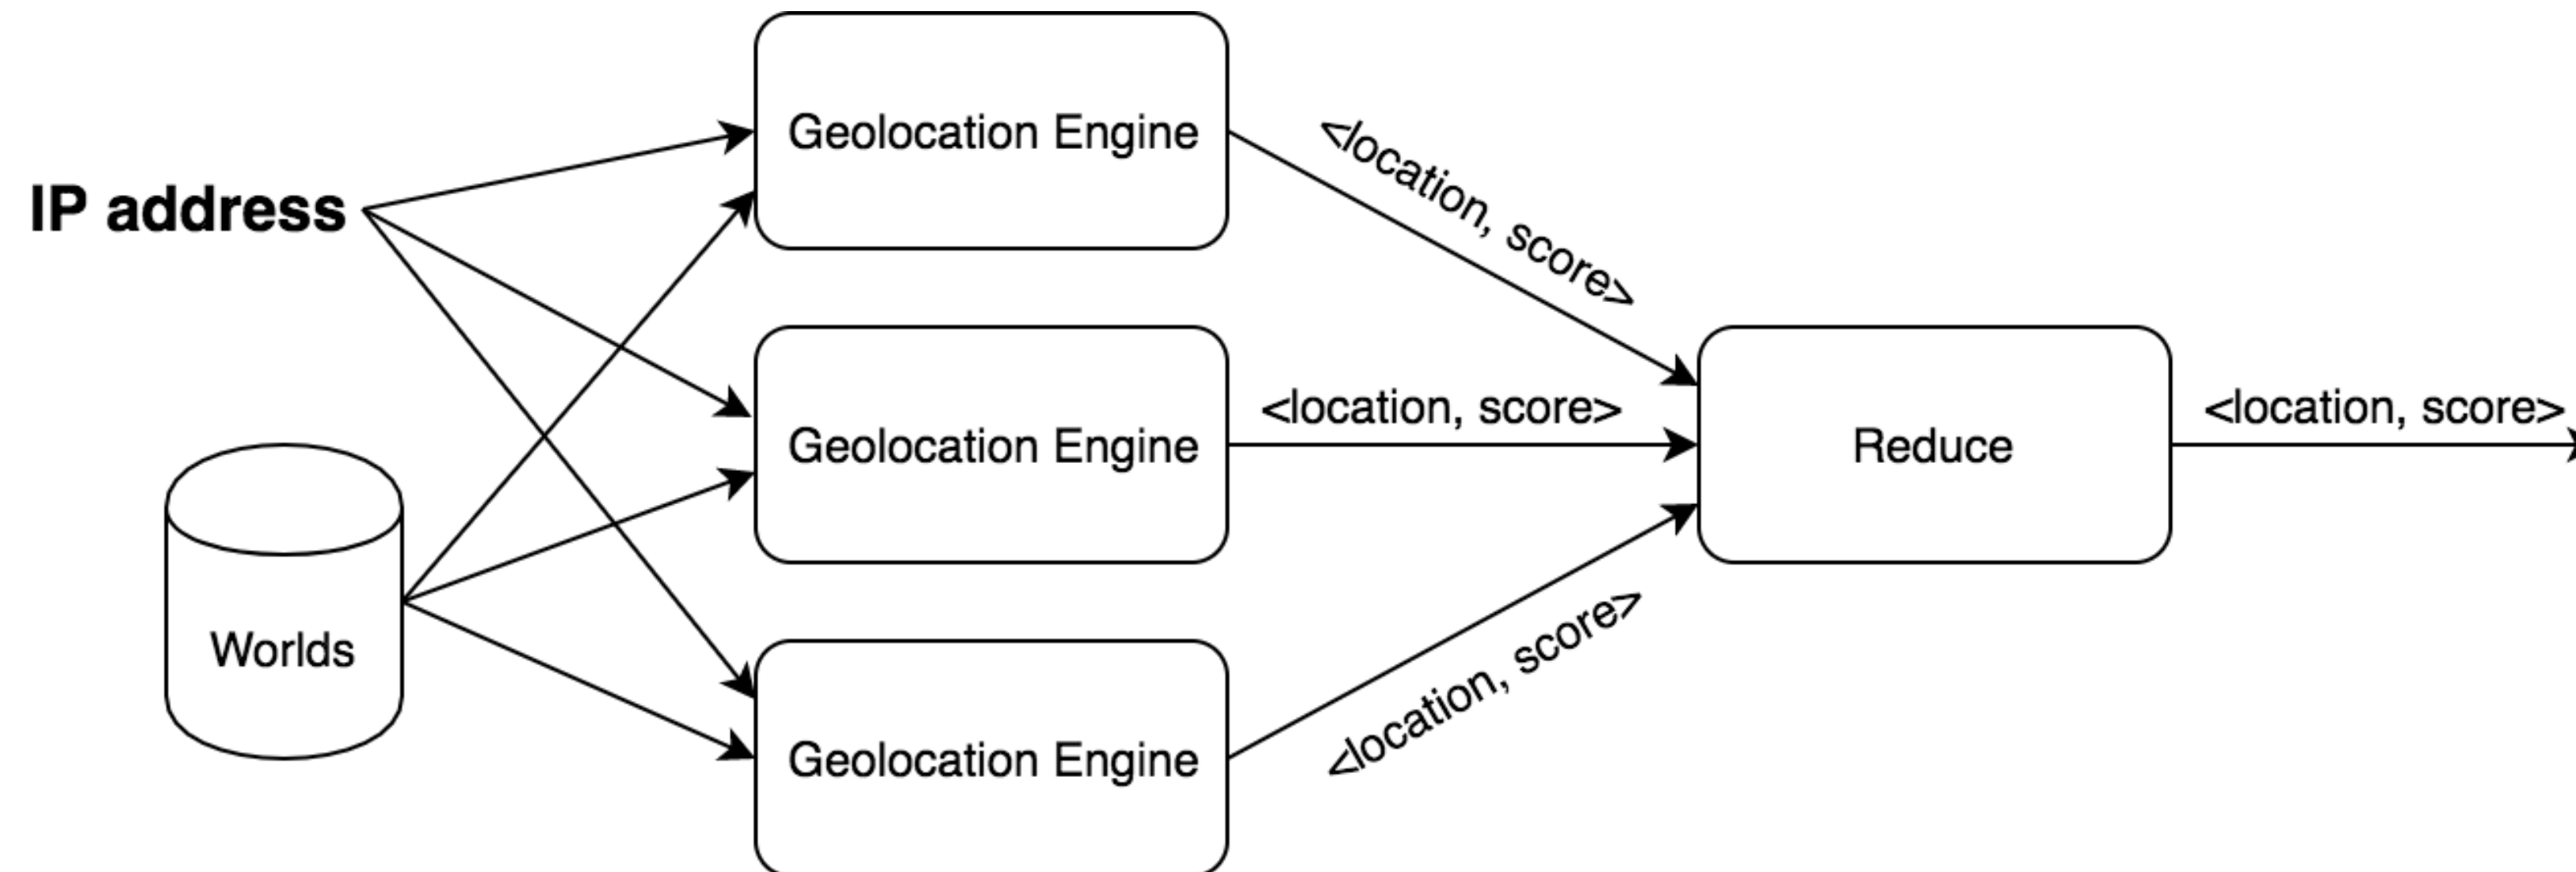
Accumulate Research Efforts as Inference Engines

Inference Engines



- Each of them is applicable only in some cases
- Each of them has a score factor
- Complete transparency about the inference methodology

Inference Engines



geolocation API



<https://openipmap.ripe.net>

```
▼ [
  ▼ {
    "url": "/locate",
    "description": "Geolocation service. It provides geolocation according to a set of passive and active geolocation approaches"
  },
  ▼ {
    "url": "/worlds",
    "description": "Worlds dataset, providing standard geolocation format to all other services"
  },
  ▼ {
    "url": "/crowdsourced",
    "description": "Geolocation service based on crowdsourced information."
  },
  ▼ {
    "url": "/peeringdb",
    "description": "PeeringDB interface for geolocation purposes."
  },
  ▼ {
    "url": "/triangulation",
    "description": "Active geolocation service based on latency triangulation."
  },
  ▼ {
    "url": "/anycast",
    "description": "Anycast geolocation service based on active measurements"
  }
]
```


openipmap.ripe.net/api/locate/ 83.163.50.165/best



```
{
  "location": {
    "score": 145,
    "countryCodeAlpha3": "NLD",
    "countryCodeAlpha2": "NL",
    "cityPopulation": 147590,
    "stateAnsiCode": "07",
    "pointGeometry": "0101000020E61000005C72DC291D8C12401B81785DBF304A40",
    "cityNameAscii": "Haarlem",
    "stateIsoCode": "NL-07",
    "countryName": "Netherlands",
    "stateName": "North Holland",
    "longitude": 4.63683,
    "geonameId": 2755003,
    "latitude": 52.38084,
    "cityName": "Haarlem",
    "type": "city",
    "id": "HAARLEM-NL-07-U173CX8KTBR196ECJF92"
  },
  "meta": {
    "distribution": {
      "version": "17.9.18.1"
    },
    "service": {
      "version": "0.0.1"
    },
    "request": {
      "params": {
        "ip": "83.163.50.165"
      },
      "query": {}
    }
  }
}
```

openipmap.ripe.net/api/locate/ 83.163.50.165/partial



```
{
  "partials": [
    {
      "engine": "probeslocation",
      "description": "Probes location suggestor - based on user setting",
      "scoreFactor": 10,
      "locations": [...] // 1 item
    },
    {
      "engine": "anycastparistech",
      "description": "Anycast engine - Paristech dataset",
      "scoreFactor": 10,
      "locations": []
    },
    {
      "engine": "crowdsourced",
      "description": "Crowdsourced engine",
      "scoreFactor": 9,
      "locations": []
    },
    {
      "engine": "triangulation",
      "description": "Triangulation engine (if empty try in 3 minutes, triangulation requires time)",
      "scoreFactor": 5,
      "locations": [...] // 20 items
    }
  ],
  "meta": {
    "distribution": {
      "version": "17.9.18.1"
    },
    "service": {
      "version": "0.0.1"
    }
  }
}
```

/locate - Active geolocation



If the IP has not been measured yet, a new Ping measurement starts

- Peering DB data and BGP data are used to reduce the locations probed
- Score based on RTT, only RTT <10ms are considered
- PeeringDB facilities and population boost the score
- A list of possible locations will be returned
- We are working on it! (Contributions are welcome!)

let's geolocate 2a07:1480:2:303::1



First step: get a set of ASes that are related to this IP address

AS59626

Prefix Overview (2a07:1480:2:303::1)

Is visible in RIS

This prefix is part of 2a07:1480::/29 announced by

AS59626
CONSCIA - Conscia A/S

Resource	RIR	Country
2a07:1480::/29	RIPE NCC	DK

Show IANA Registry Information

Showing results for 2a07:1480::/29 as of 2018-03-07 08:00:00 UTC

Given resource is not announced but result has been aligned to first-level less-specific (2a07:1480::/29).

[source data](#) [embed code](#) [permalink](#) [info](#)

AS16245

ASN Neighbours (AS59626)

- Left #: 1
- Right #: 0
- Unique #: 1
- Uncertain #: 0

Show entries Search:

ASN	Name	Type	Path Count	IP Version
AS16245	NGDC - Netgroup A/S	left	216	v4 & v6

Showing 1 to 1 of 1 entries

ASN names are valid for 2018-03-07 13:53 UTC

Showing results for AS59626 from 2018-03-07 00:00:00 UTC

Query time has been set to the latest available time (2018-03-07 00:00 UTC)

[source data](#) [embed code](#) [permalink](#) [info](#)

let's geolocate 2a07:1480:2:303::1



Second step: get a set of geographic locations related with these ASes

AS59626

Public Peering Exchange Points		Filter
Exchange ▼ ASN	IPv4 IPv6	Speed RS Peer
Nothing matched your filter You may filter by Exchange, ASN or Speed		
Private Peering Facilities		Filter
Facility ▼ ASN	Country City	
GlobalConnect Copenhagen 59626	Denmark Copenhagen	

AS16245

Public Peering Exchange Points		Filter
Exchange ▼ ASN	IPv4 IPv6	Speed RS Peer
DIX DIX LAN 16245	192.38.7.93 2001:7f8:1f:0:1:6245:93:0	10G <input type="radio"/>
STHIX - Copenhagen 16245	185.1.88.22 2001:7f8:b0::1:6245:1	10G <input type="radio"/>
Private Peering Facilities		Filter
Facility ▼ ASN	Country City	
GlobalConnect Copenhagen 16245	Denmark Copenhagen	
Interxion Copenhagen 16245	Denmark Copenhagen	
Telia Telehouse Copenhagen 16245	Denmark Copenhagen	
UNI-C 16245	Denmark Lyngby	

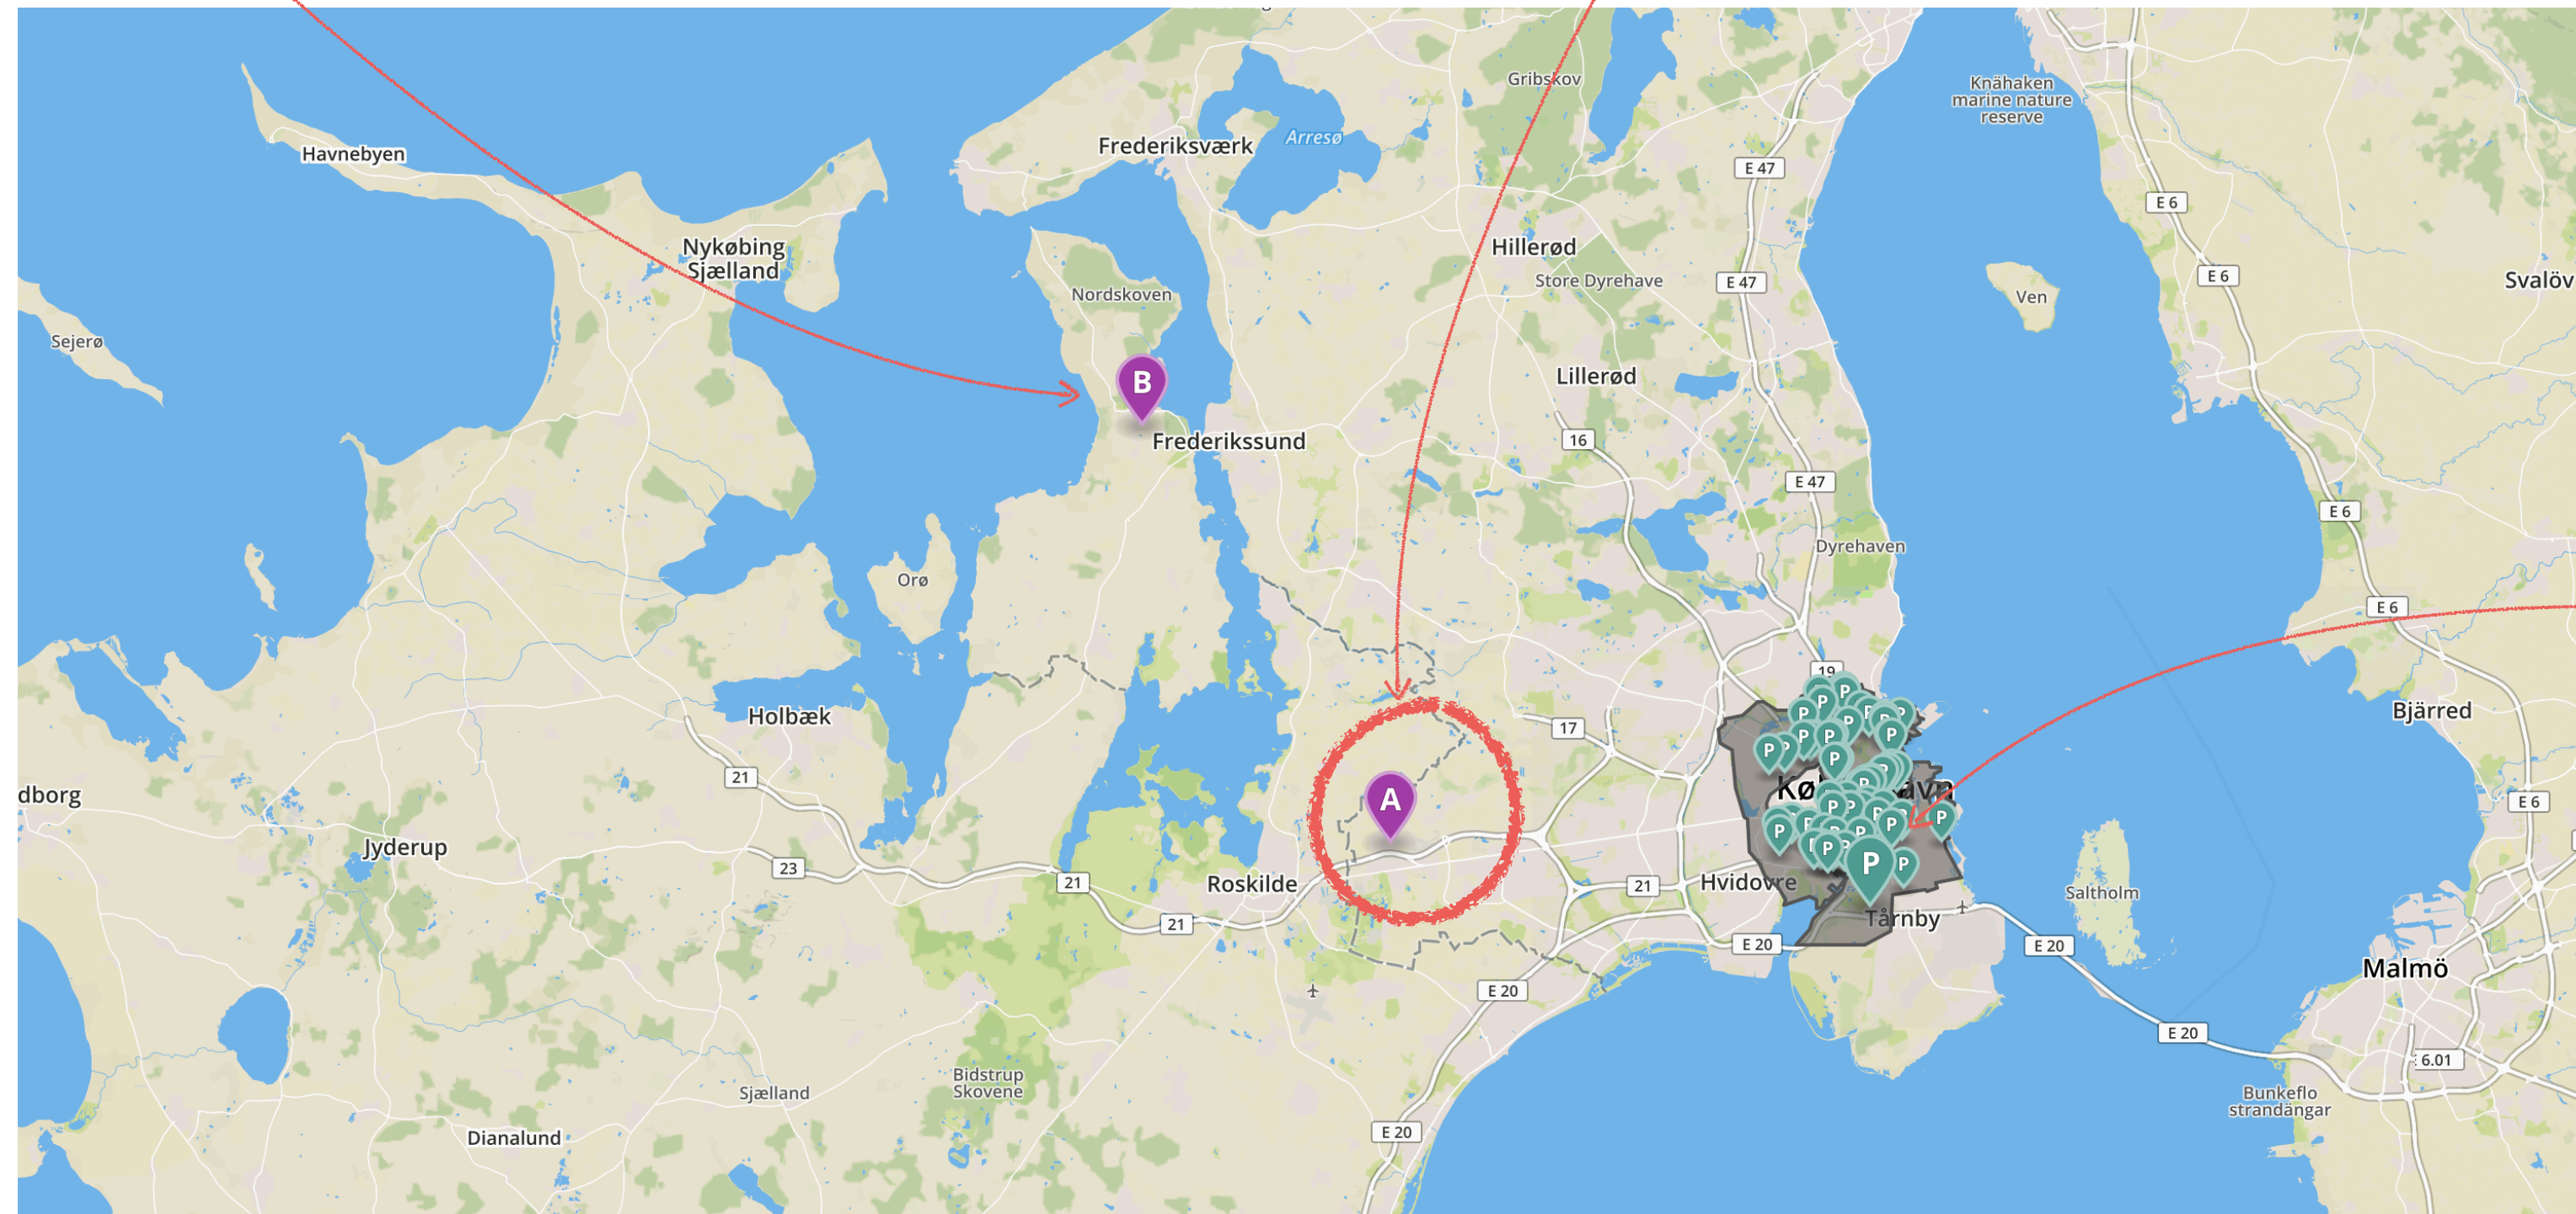
let's geolocate 2a07:1480:2:303::1



Third step: run ping measurements on RIPE Atlas probes within these ASes and these locations.

AS59626 (6,4ms)

AS16245 (2.1ms!)



from location (6,4ms)

let's geolocate 2a07:1480:2:303::1



The screenshot shows the OpenIPmap interface. At the top, the RIPE NCC logo and navigation menu are visible. The main content area features a search bar with the IP address 2a07:1480:2:303::1 entered. Below the search bar, a map of Europe is displayed with a red pin indicating the location of Ballerup, Denmark. A sidebar on the left provides detailed information about the IP location, including the IP address and the specific location: Ballerup, DK-17.

OpenIPmap Geolocating Internet Infrastructure

2a07:1480:2:303::1

About | API reference | Manual

2a07:1480:2:303::1 Ballerup,DK-17 Denmark

IP LOCATION

2a07:1480:2:303::1 Ballerup, DK-17

Home | Sitemap | Contact Us | Service Announcements | Privacy Statement | Legal | Cookies | Copyright Statement | Terms of Service



Crowdsourcing with OpenIPmap

Web UI: put stuff on a map



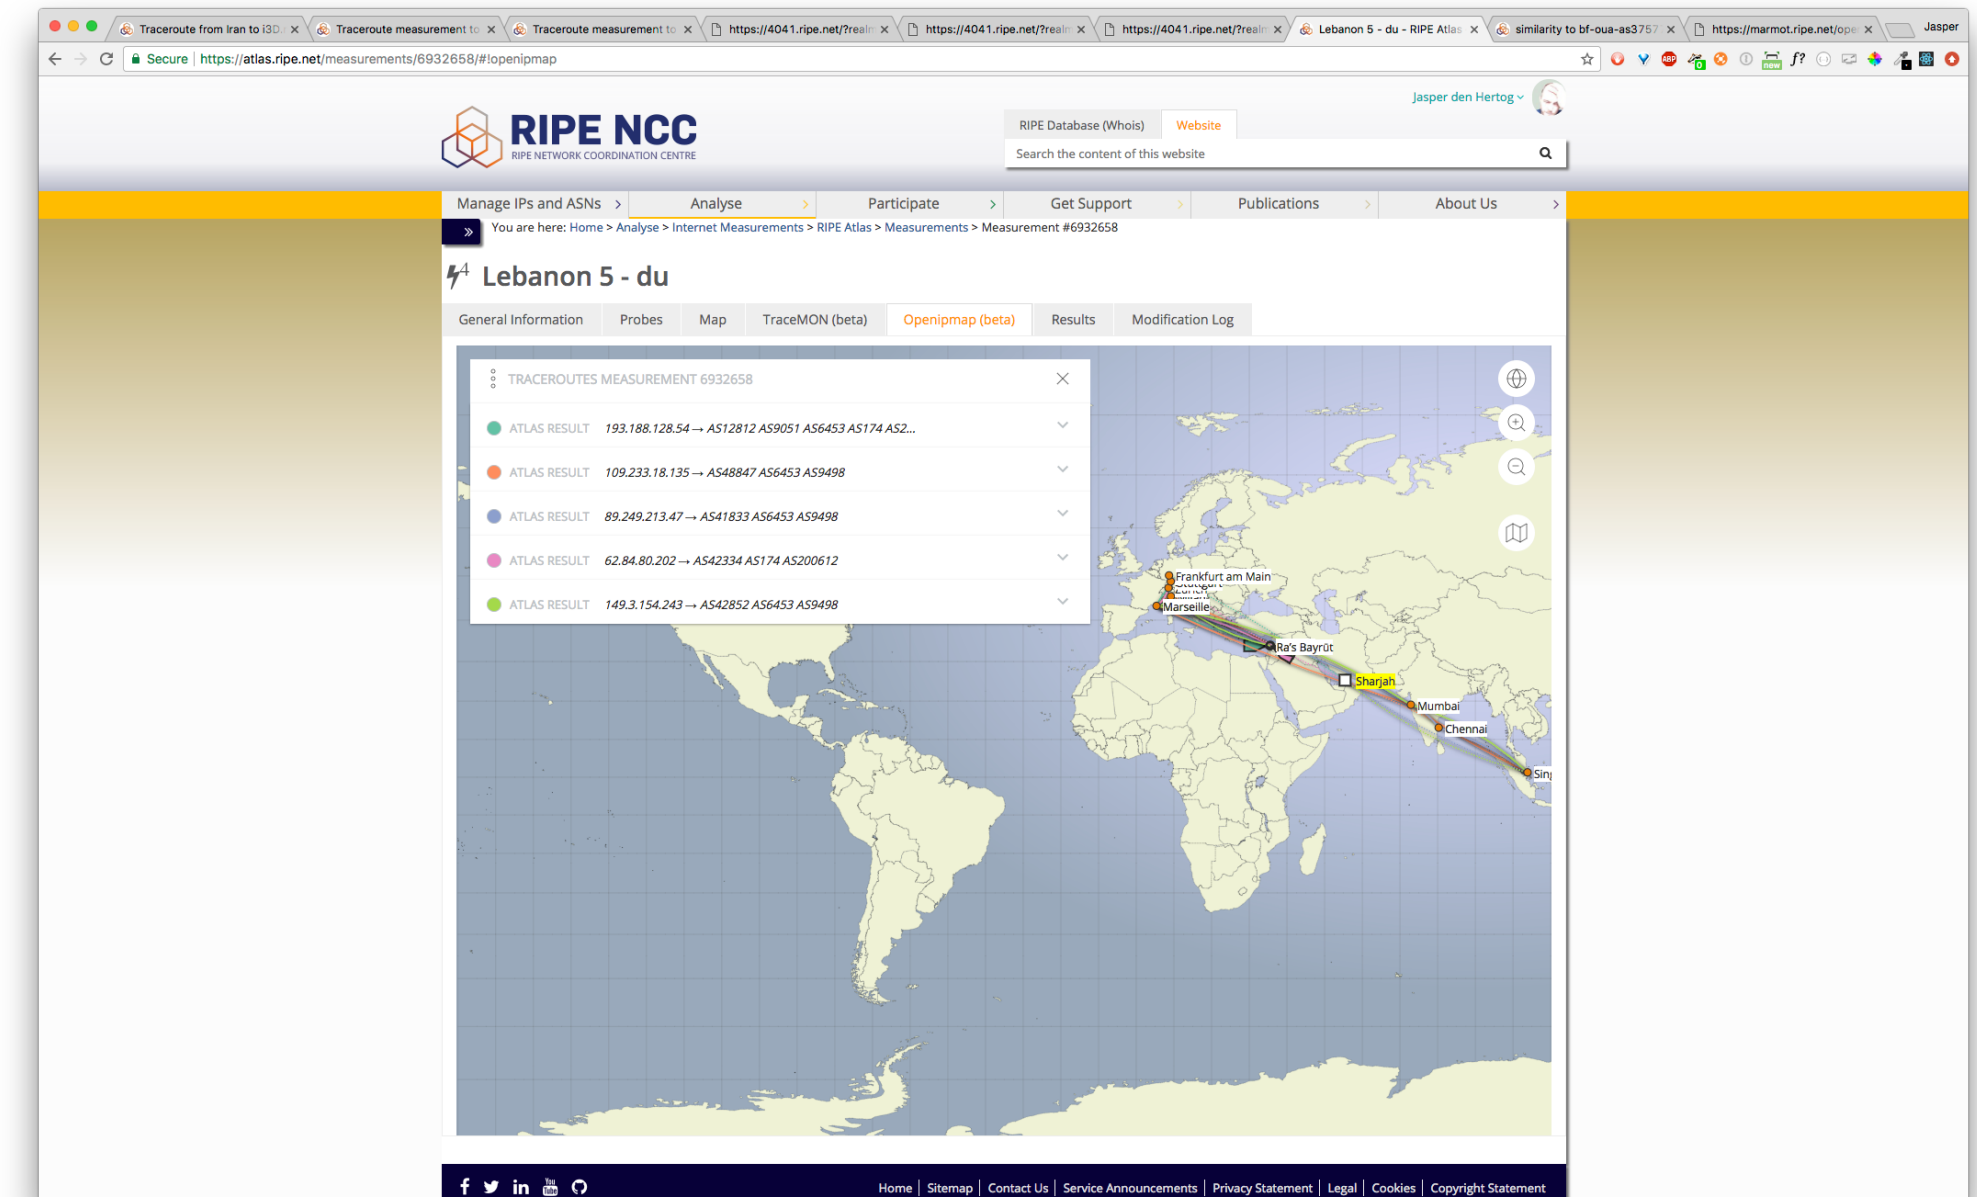
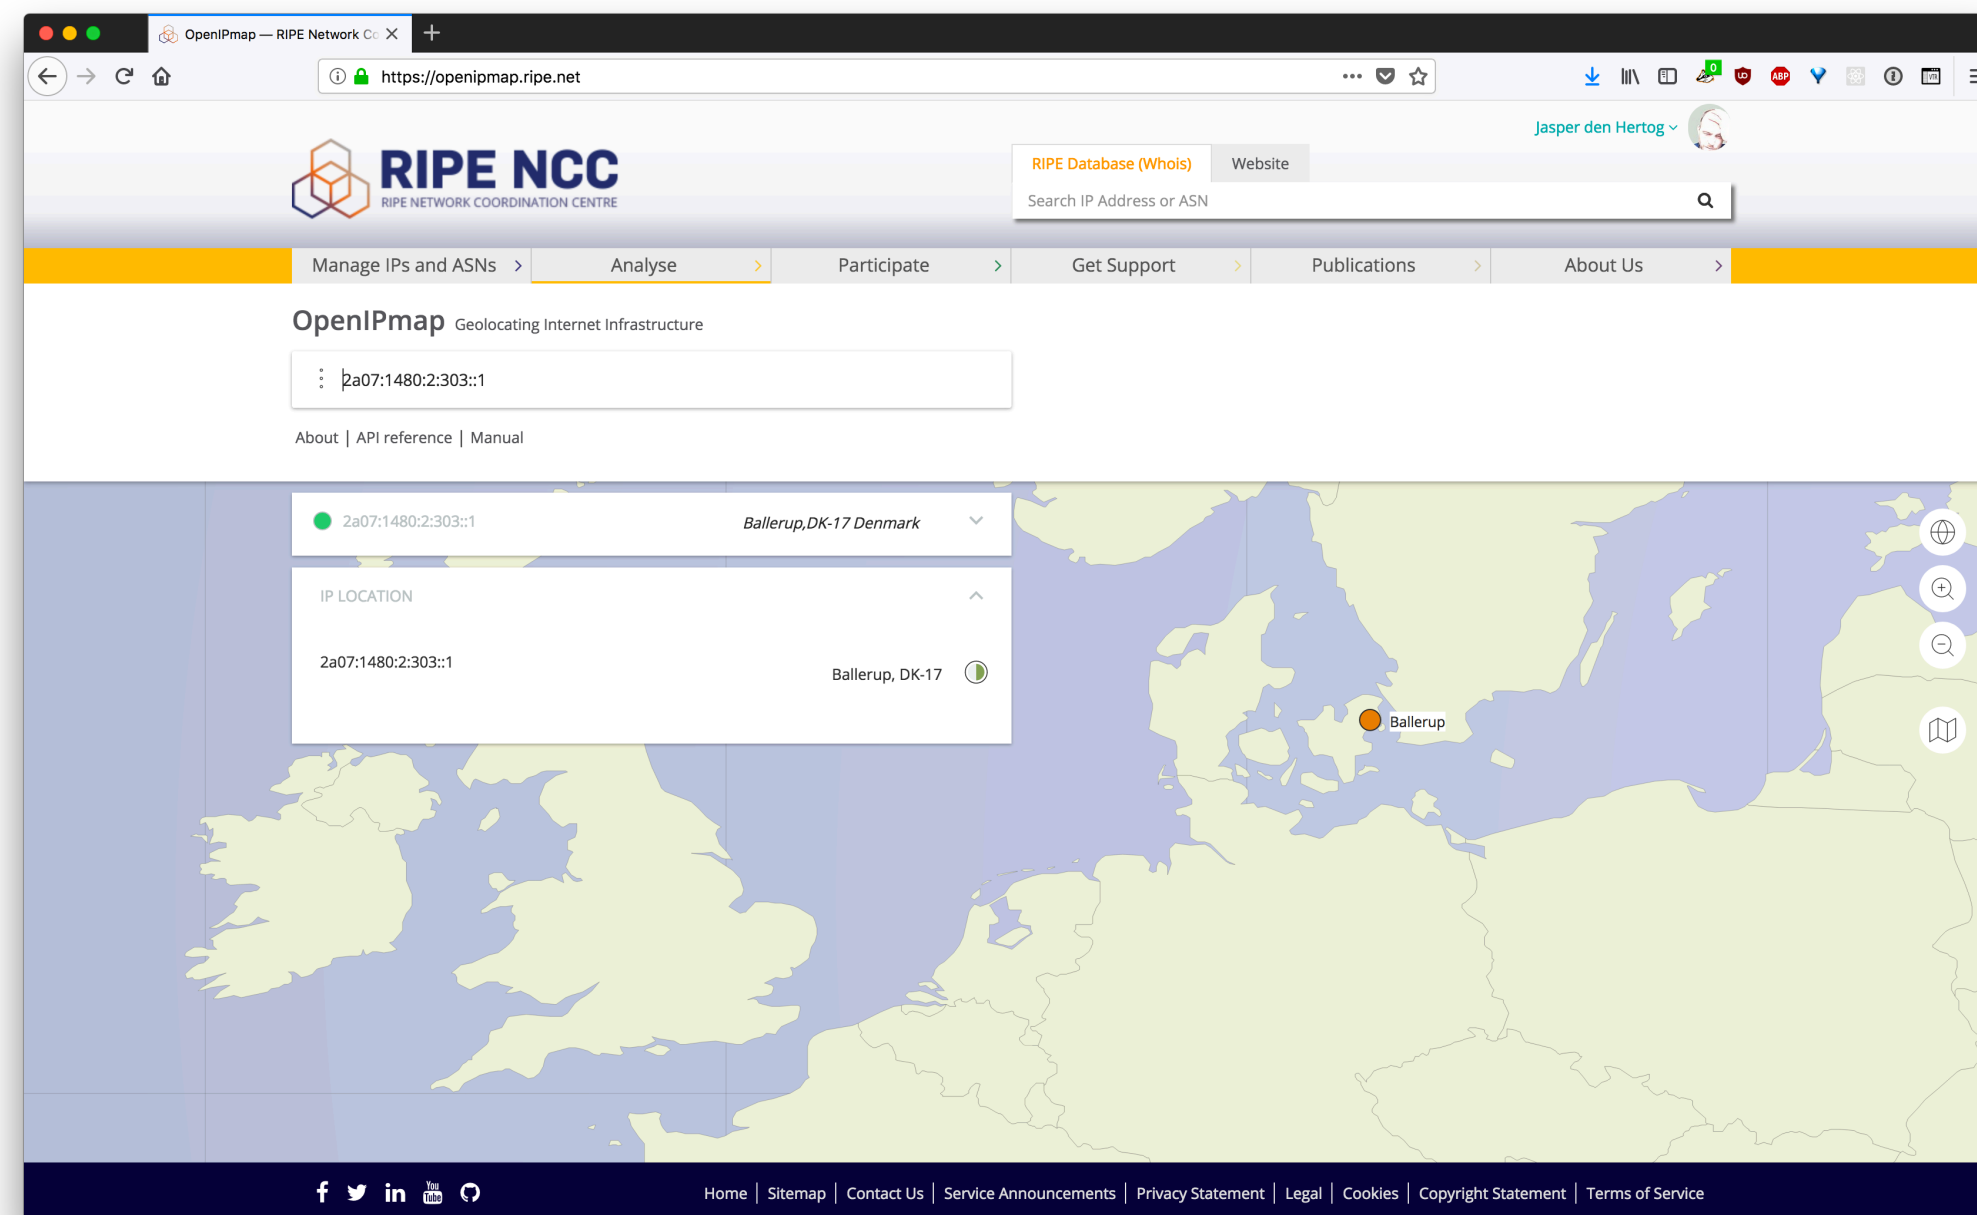
- Make sense of RTTs in one glimpse
- Understand network relationships among countries
- Verify geographical optimisation policies
- Crowdsource IP addresses to geolocation

...actually we have two interfaces



<https://openipmap.ripe.net>

<https://atlas.ripe.net/measurements/<TRACEROUTEMSM>>



The first one we already saw...



The screenshot shows the OpenIPmap interface. At the top, the RIPE NCC logo and navigation menu are visible. The main content area features a search bar with the IP address `2a07:1480:2:303::1` entered. Below the search bar, a map of Europe is displayed with a red pin marking the location of Ballerup, Denmark. A sidebar on the left provides details for the IP location, including the IP address and the location name.

OpenIPmap Geolocating Internet Infrastructure

2a07:1480:2:303::1

About | API reference | Manual

2a07:1480:2:303::1 Ballerup,DK-17 Denmark

IP LOCATION

2a07:1480:2:303::1 Ballerup, DK-17

Home | Sitemap | Contact Us | Service Announcements | Privacy Statement | Legal | Cookies | Copyright Statement | Terms of Service

The second shows traceroutes



<https://atlas.ripe.net/measurements/11585456/#!openipmap>

The screenshot displays the RIPE Atlas OpenIPMap interface. The browser address bar shows the URL <https://atlas.ripe.net/measurements/11585456/#!openipmap>. The interface includes a navigation menu with tabs for General Information, Probes, Map, TraceMON, Openipmap (beta), and Results. The main content area is divided into two panels. The left panel, titled 'TRACEROUTES MEASUREMENT 11585456', lists six ATLAS RESULT entries, each with a colored dot and a truncated IP address range. Below this is a 'SUMMARY' section showing a tree view of the traceroute path. The path starts at AS16245 (IP 2001:1448:380:102:c24a:ff:fe...) and branches into several paths, including one through AS59626 and another through an AS unknown. The right panel shows a map of Europe with several colored lines representing traceroute paths between various locations, including Arhus, Hvidovreagenz, and Stockholm. A 'REMOVE FROM MAP' button is visible below the summary section.

TRACEROUTES MEASUREMENT 11585456

- ATLAS RESULT 2001:470:28:8b2:6666:b3ff:feb0:ef80 → AS6939 AS16245 A...
- ATLAS RESULT 2a02:980:2510:ba00:220:4aff:fec8:25be → AS39554 AS3292 AS162...
- ATLAS RESULT 2a01:4f0:2:666::2 → AS28717 AS48374 AS16245 AS5...
- ATLAS RESULT 2a03:5440:100:42::2 → AS42541 AS16245
- ATLAS RESULT 2001:470:28:8b2:6666:b3ff:feb0:e87c → AS6939 AS1835 AS1624...
- ATLAS RESULT 2001:1448:380:102:c24a:ff:fecc:75b8 → AS16245 AS5...

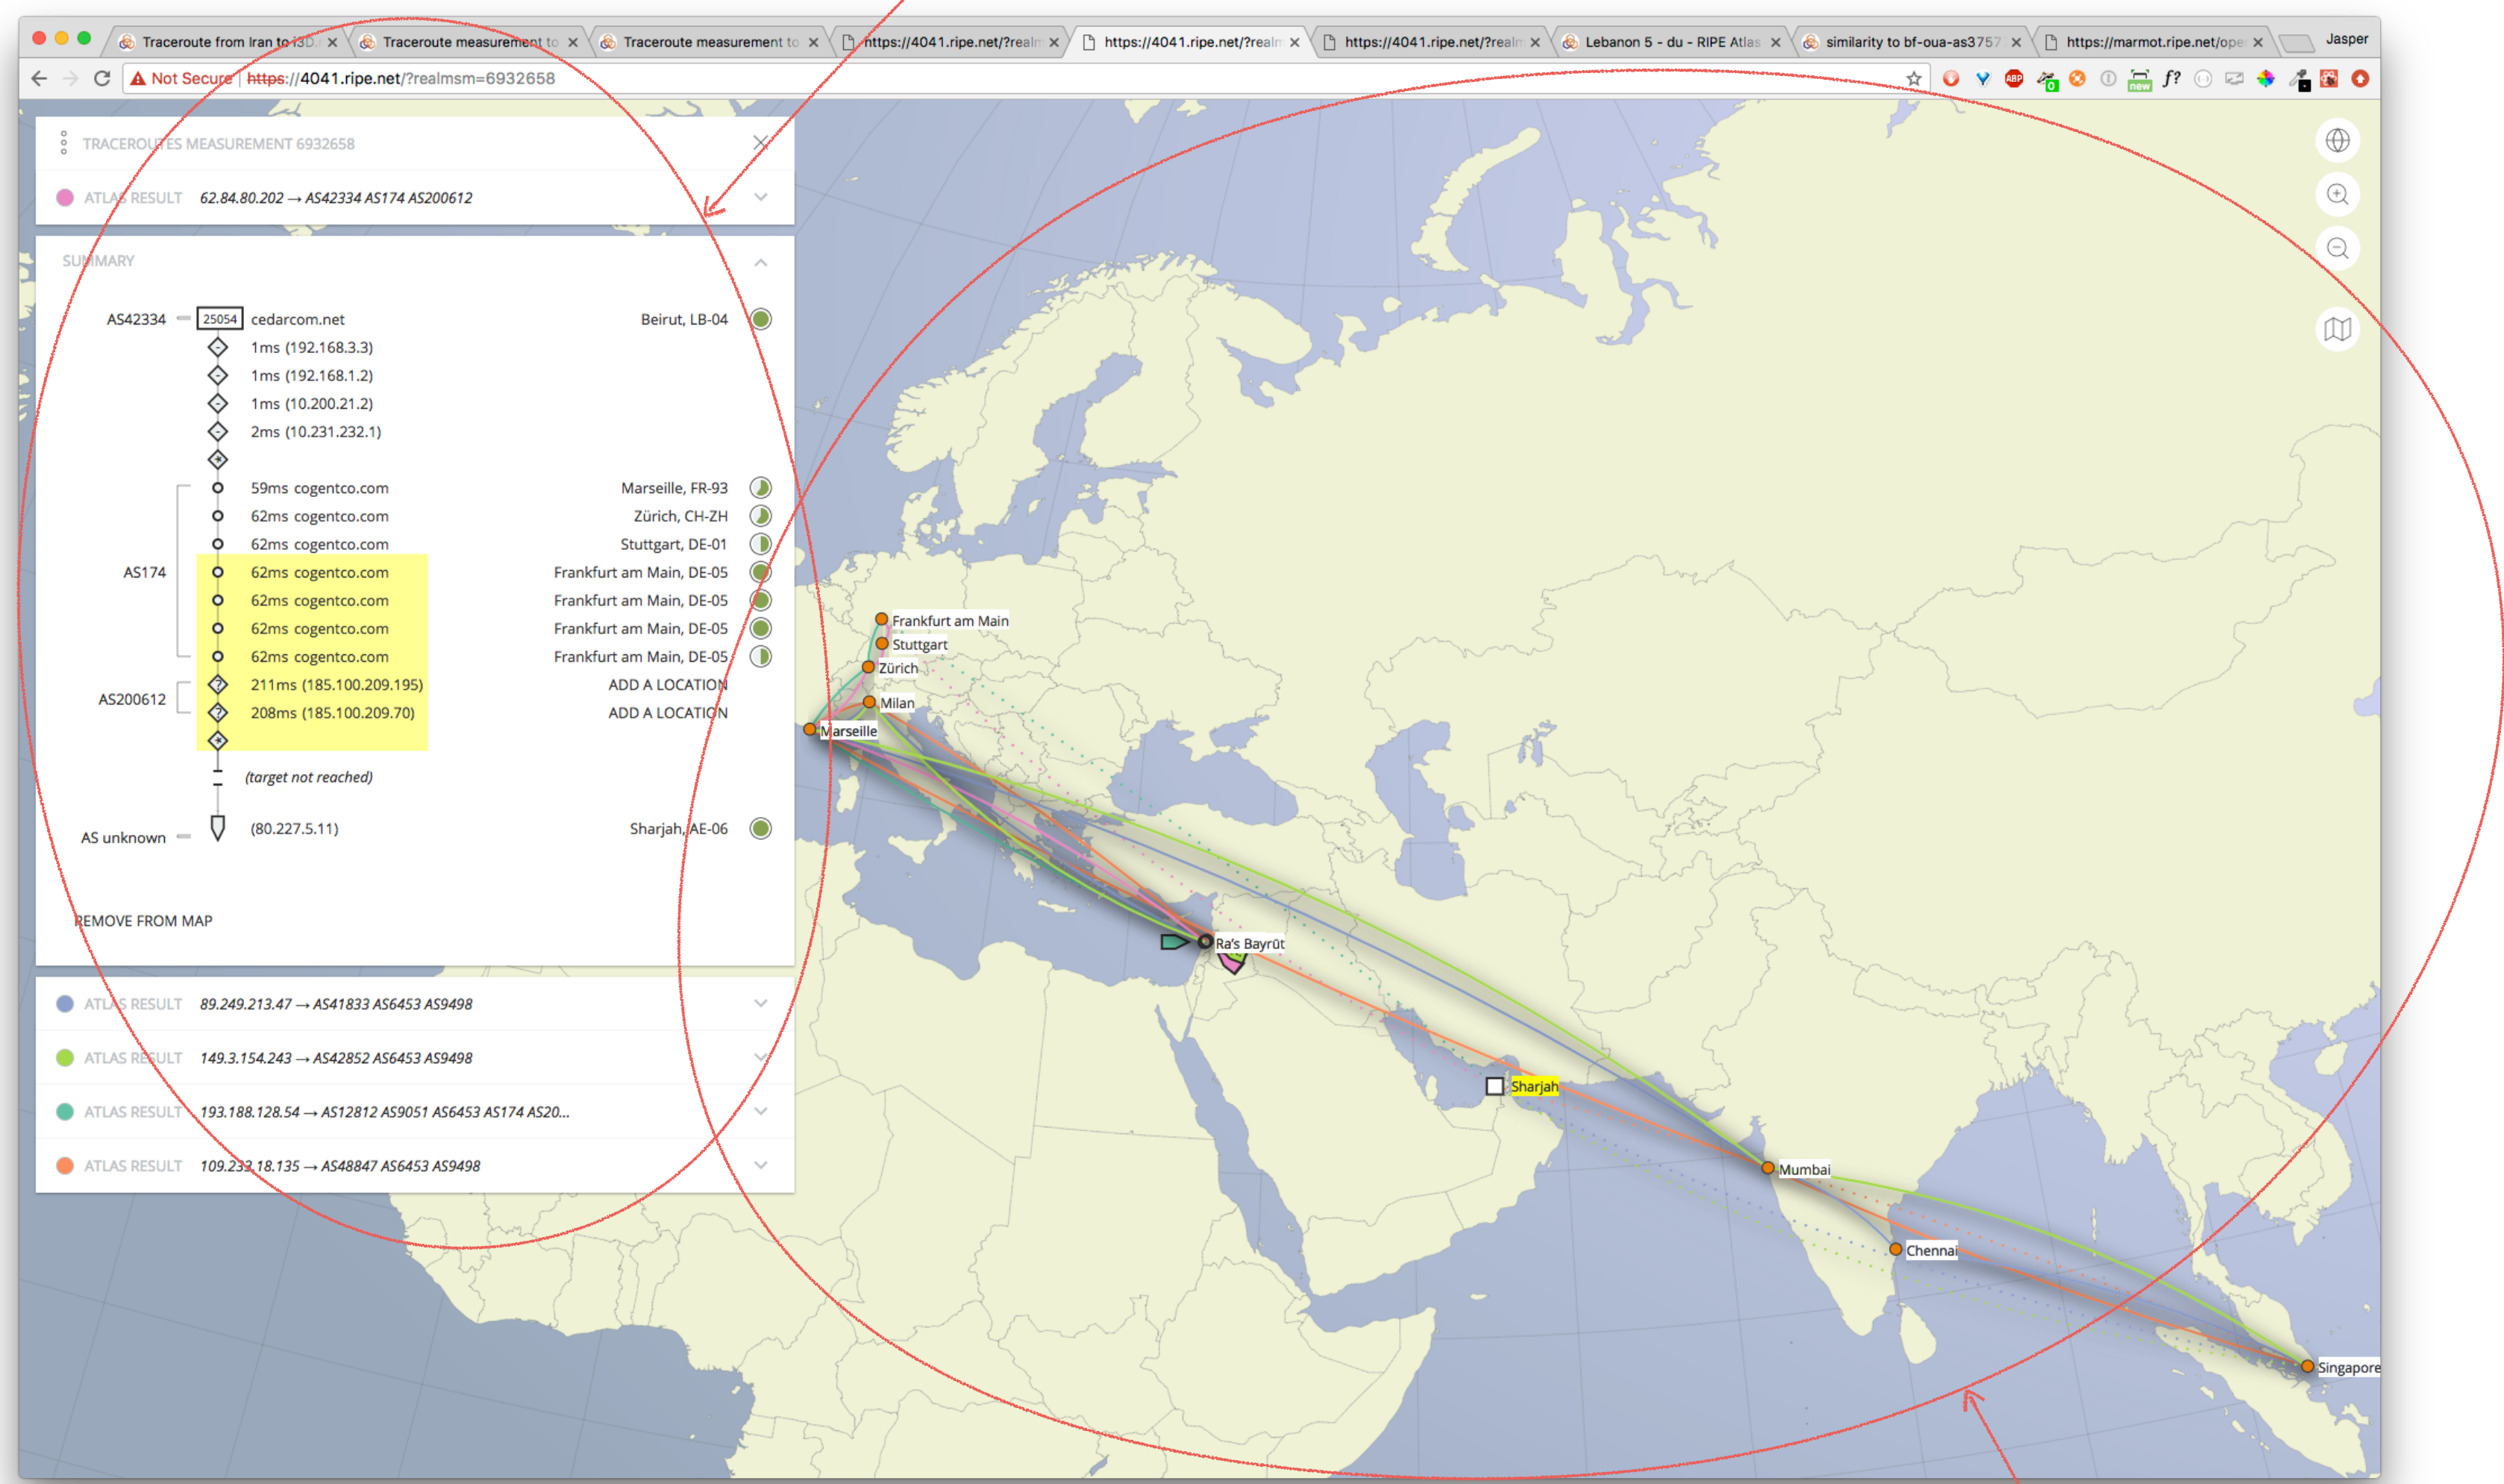
SUMMARY

- 18910 (2001:1448:380:102:c24a:ff:fe...) Taastrup, DK-17
- AS16245
 - 2ms (2001:1448:380:102::1) Roskilde, DK-20
 - 2ms (2001:1448:0:a402::1) Ballerup, DK-17
 - 3ms ngdc.net Ballerup, DK-17
 - 3ms ngdc.net Copenhagen, DK-17
- AS59626 2ms (2a07:1480:2:303::1) Ballerup, DK-17
- AS unknown (2a07:1480:2:303::1) Ballerup, DK-17

REMOVE FROM MAP

- ATLAS RESULT Pfd45:de5c:ec80:0:c24a:ff:fea0:17ea-0.03561073631270...
- ATLAS RESULT 2a05:f6c7:748:0:eade:27ff:fe69:dd4e → AS203953 AS42525 AS162...

elements



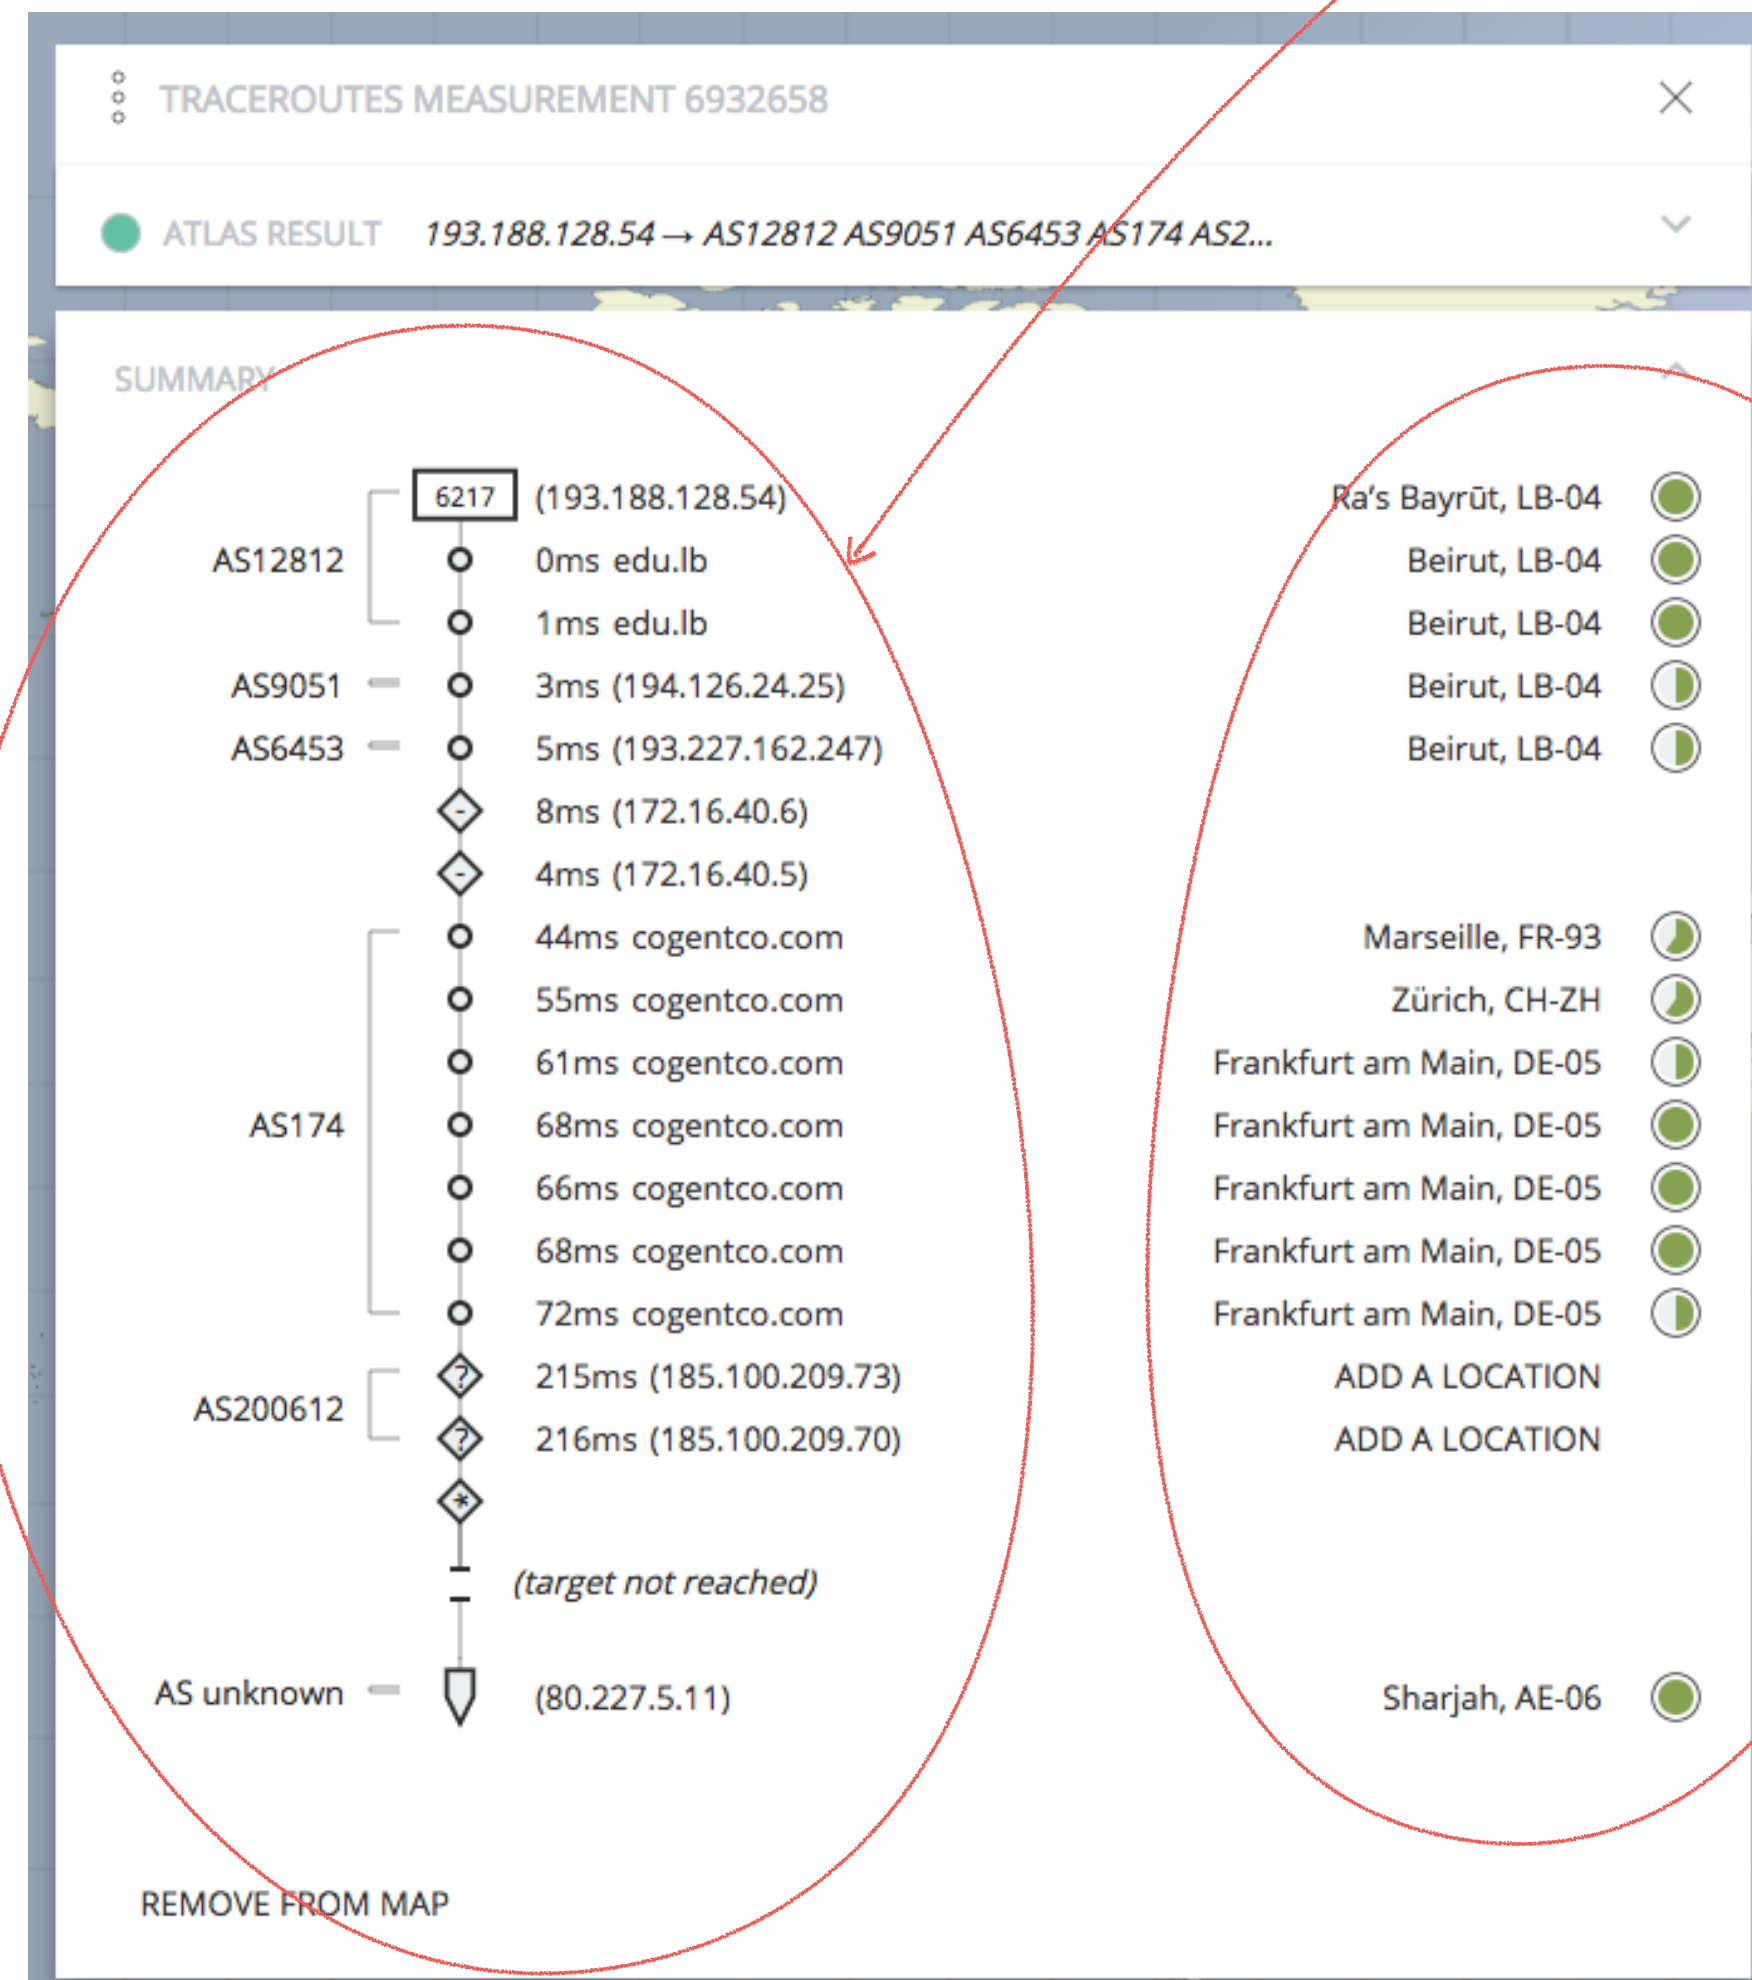
sidebar

map

sidebar

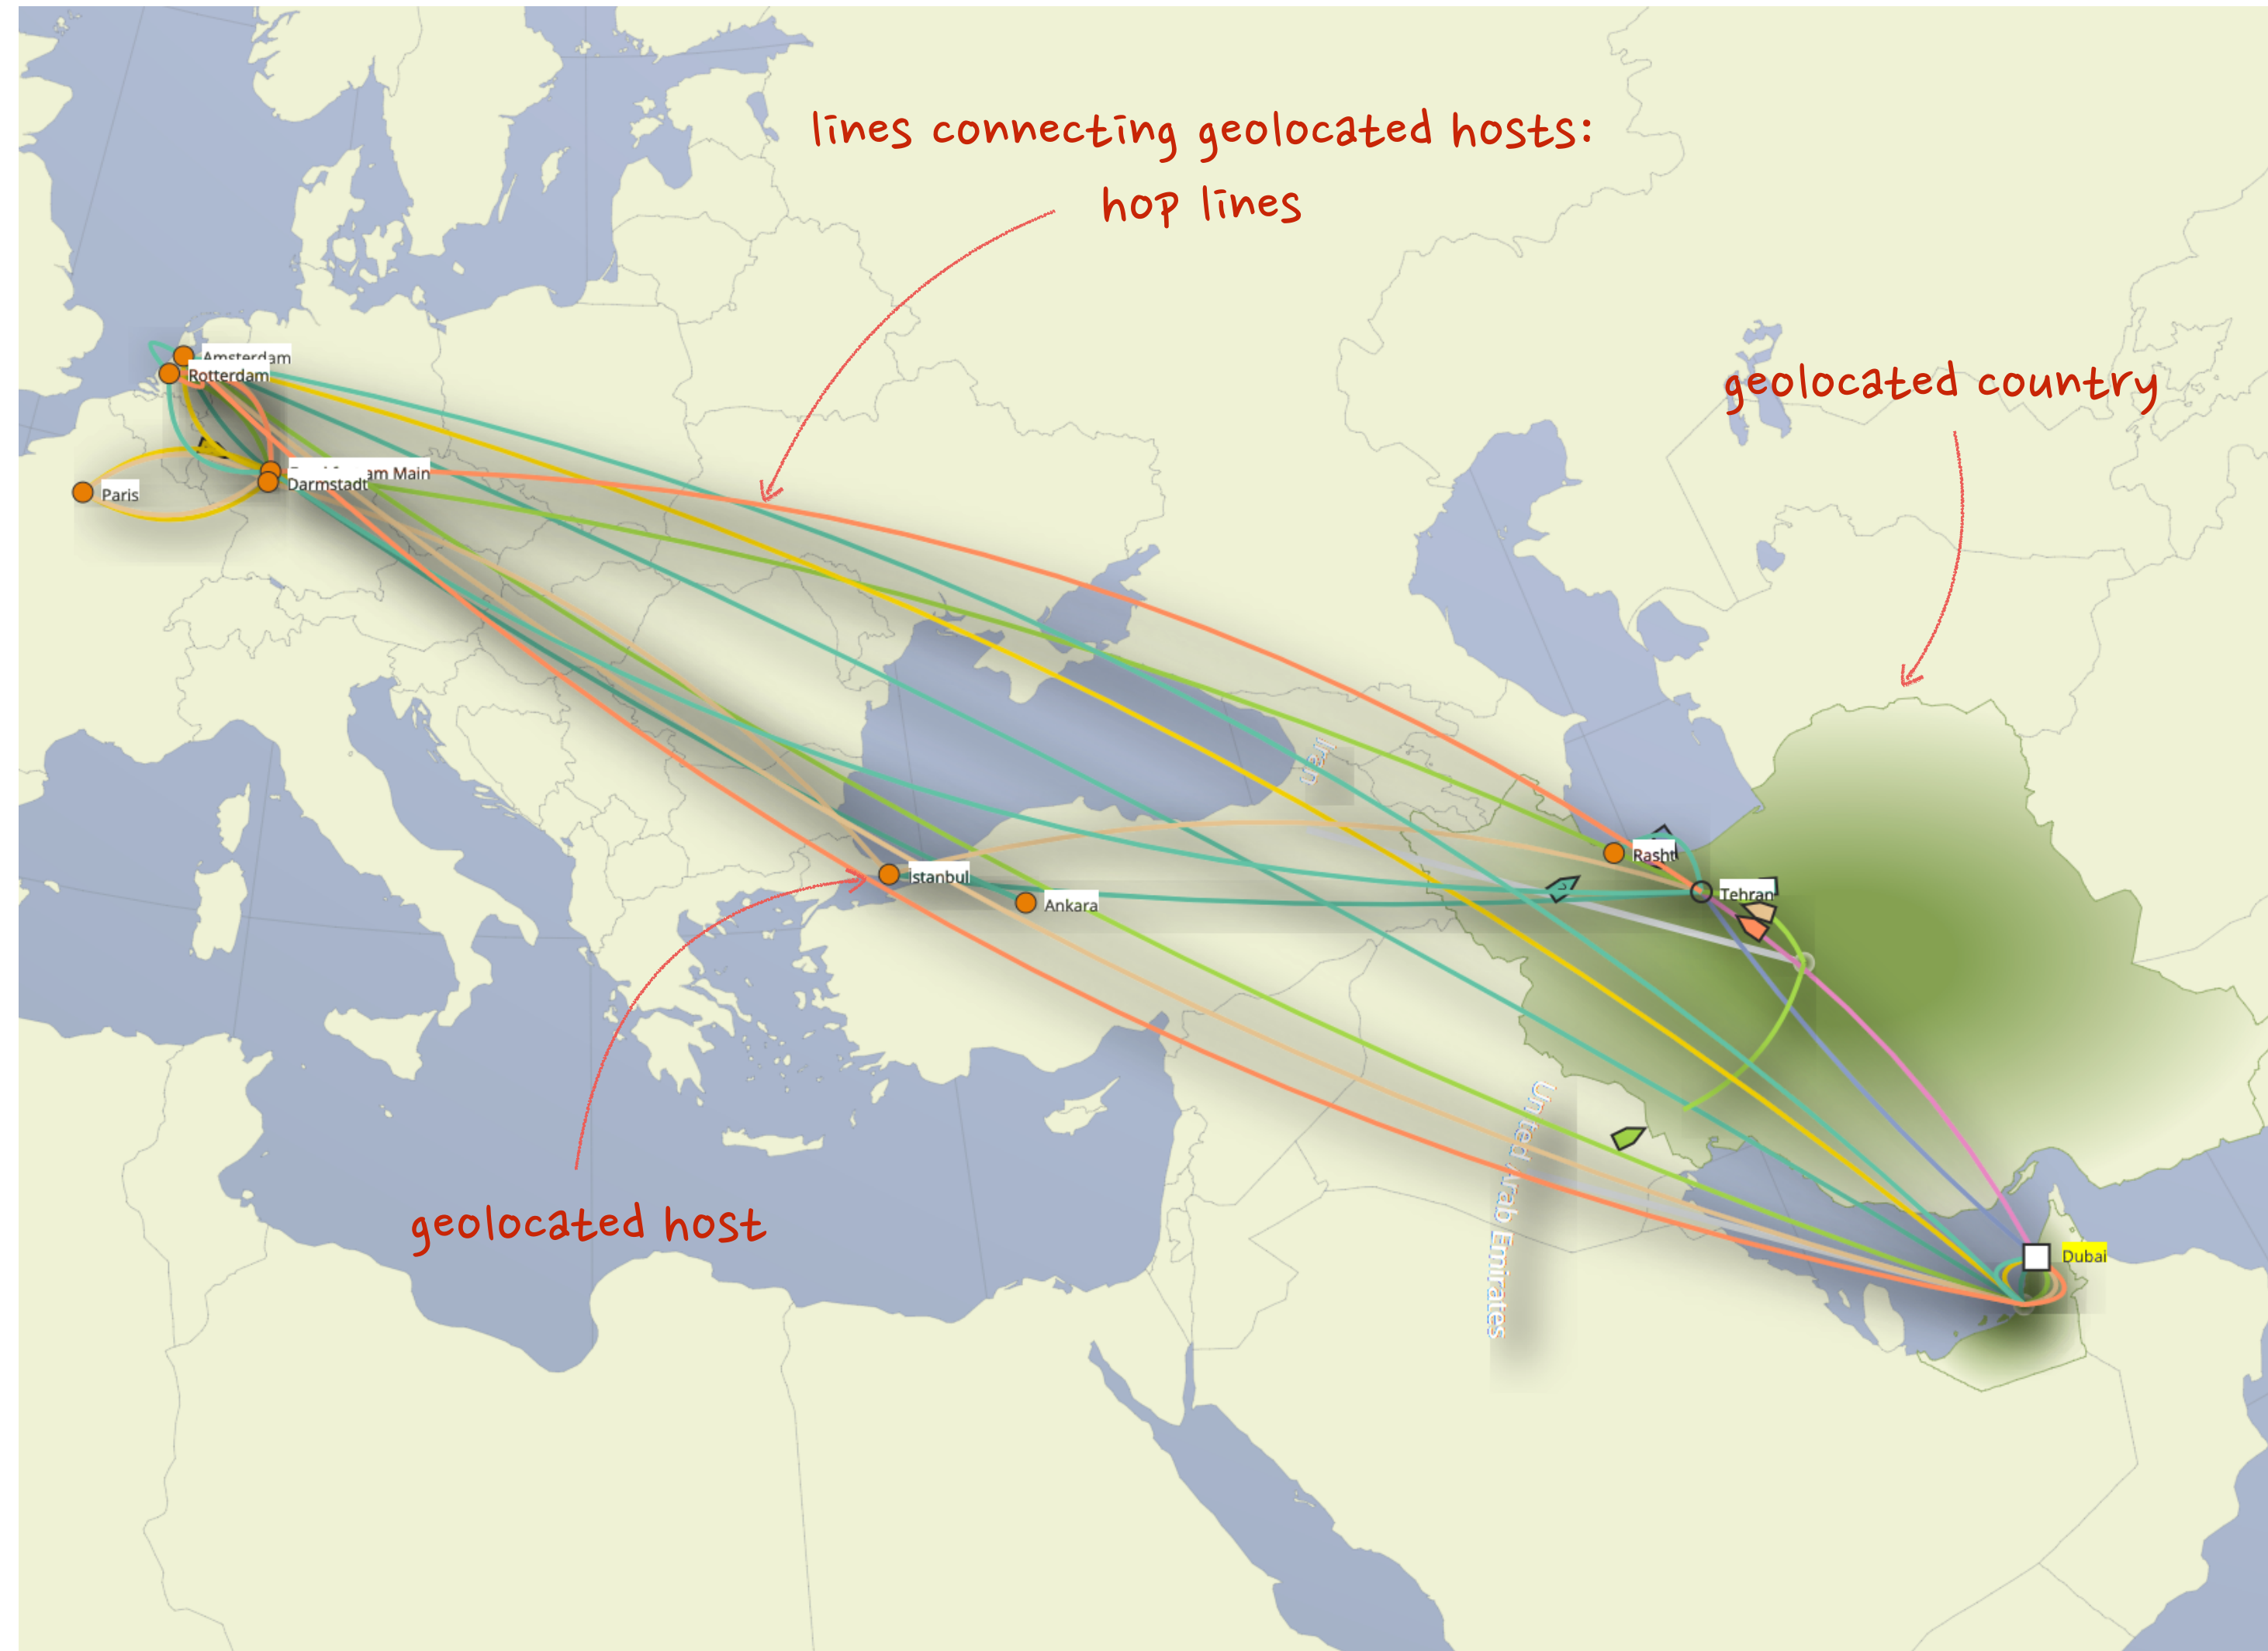


hops with enriched info



Geographical information

map



interaction sidebar and map



The screenshot displays a web interface with two traceroute measurement results in the sidebar and a map of South America. The top result is for 201.52.240.184, and the bottom result is for 189.100.75.210. The sidebar shows hop details with IP addresses and round-trip times. The map shows a network of hops connecting various locations in South America, with a specific hop line highlighted in green. Red annotations include a circle around a hop in the first result and arrows pointing to hop lines on the map.

Traceroute 1: 201.52.240.184

AS	IP	RTT
AS28573	16730 (201.52.240.184)	
	1ms (192.168.130.1)	1ms
	1ms (192.168.131.1)	1ms
	2ms (192.168.0.1)	2ms
	9ms (10.17.0.1)	9ms
	13ms (201.6.45.1)	13ms
AS28573	14ms com.br	14ms
	15ms net.br	15ms
	15ms net.br	15ms
AS4230	17ms net.br	17ms
	14ms net.br	14ms
	12ms net.br	12ms
AS54113	14ms (151.101.56.19)	14ms
AS unknown	(151.101.56.193)	

Traceroute 2: 189.100.75.210

AS	IP	RTT
AS28573	10349 (189.100.75.210)	
	0ms (192.168.1.1)	0ms
	1ms (192.168.0.1)	1ms
	28ms com.br	28ms
AS28573	10ms com.br	10ms
	11ms (201.6.40.45)	11ms
	12ms com.br	12ms
	13ms net.br	13ms
	15ms net.br	15ms
AS4230	14ms net.br	14ms
	13ms net.br	13ms
	12ms net.br	12ms
AS54113	13ms (151.101.20.193)	13ms
AS unknown	(151.101.20.193)	

Map Locations: Knoxville, Atlanta, Houston, West Palm Beach, Fortaleza, Belo Horizonte, Rio de Janeiro, Barueri, Passo Fundo, Lajeado, Curitiba Alegre, São Paulo, BR-27, Brazil, Barueri, BR-27, Fortaleza, BR-06, Syracuse, US-NY.

Annotations:
- "shows traceroute and highlights hops along this line" (red arrow pointing to the green hop line on the map)
- "double-click hop line" (red arrow pointing to a hop line on the map)

to each traceroute its own line



data doubt: scores for locations



score indicator in sidebar

Perth, AU-08

LOCATION SCORE

The likelihood of this being the correct location is very low. You should not consider this an accurate estimate.

DISMISS

São Paulo, BR-27

LOCATION SCORE

OpenIPMap has calculated that the likelihood of this being the actual location is around 50%.

DISMISS

São Paulo, BR-27

LOCATION SCORE

There is a high likelihood that this is the actual location.

DISMISS

crowdsource: add locations



SUMMARY

AS28573 — 14635 (189.63.153.221)

- 1ms (192.168.1.1)
- 11ms (10.35.128.1)

AS28573

- 9ms com.br
- 13ms com.br
- 11ms com.br

AS4230

- 20ms (200.167.43.9)
- 32ms (200.244.214.13)

AS54113

- 36ms net.br
- 34ms net.br
- 35ms net.br

AS unknown

- 32ms (151.101.56.193)

REMOVE FROM MAP

click here and...

Porto Alegre, BR-23

Lajeado, BR-23

Passo Fundo, BR-23

ADD A LOCATION

EDIT HOST LOCATION

HOST
200.244.214.13
REVERSE DNS NOT AVAILABLE

CURRENT LOCATION

City

Country

CANCEL CONFIRM

ATLAS RESULT 177.32.129.64 → AS28573 AS4230 AS54113

ATLAS RESULT 189.33.89.42 → AS28573 AS4230 AS54113

ATLAS RESULT 201.52.240.184 → AS28573 AS4230 AS54113

ATLAS RESULT 191.180.97.223 → AS28573 AS4230 AS54113

ATLAS RESULT 189.101.11.138 → AS28573 AS4230 AS3356 AS2914 AS54113

ATLAS RESULT 187.180.71.87 → AS28573 AS4230 AS54113

start typing a city or...

a country

confirm and boost location score



SUMMARY

AS28573 — 14635 (189.63.153.221)

- 1ms (192.168.1.1)
- 11ms (10.35.128.1)

AS28573

- 9ms com.br
- 13ms com.br
- 11ms com.br

AS4230

- 20ms (200.167.43.9)
- 32ms (200.244.214.13)
- 36ms net.br
- 34ms net.br
- 35ms net.br

AS54113 — 32ms (151.101.56.193)

AS unknown — (151.101.56.193)

REMOVE FROM MAP

Porto Alegre, BR-23

Lajeado, BR-23

Passo Fundo, BR-23

ADD A LOCATION

ADD A LOCATION

EDIT HOST LOCATION

HOST

200.230.252.110
ebt-B1421-tcore01.spoph.embratel.net.br

CURRENT LOCATION

City

São Paulo

Country

Brazil

CANCEL CONFIRM

ATLAS RESULT 177.32.129.64 → AS28573 AS4230 AS54113

ATLAS RESULT 189.33.89.42 → AS28573 AS4230 AS54113

ATLAS RESULT 201.52.240.184 → AS28573 AS4230 AS54113

ATLAS RESULT 191.180.97.223 → AS28573 AS4230 AS54113

ATLAS RESULT 189.101.11.138 → AS28573 AS4230 AS3356 AS2914 AS54113

ATLAS RESULT 187.180.71.87 → AS28573 AS4230 AS54113

boost the score for this location

Change and improve locations



The image shows two sequential screenshots of a mobile application's 'EDIT HOST LOCATION' dialog box. The dialog box contains fields for 'HOST' (201.6.35.193, c90623c1.virtua.com.br), 'CURRENT LOCATION' (City: São Paulo, Country: Brazil), and a list of 'cities that match your input' (São Paulo, Brazil BR-27). The 'SUBMIT' button is visible at the bottom right.

start typing and...

...choose a new or more specific location and...

now you can submit!

Change and improve



TRACEROUTES MEASUREMENT 9354519

ATLAS RESULT 189.63.153.221 → AS28573 AS4230 AS54113

SUMMARY

AS28573	14635	(189.63.153.221)
	◇	1ms (192.168.1.1)
	◇	11ms (10.35.128.1)
AS28573	○	9ms com.br
	○	13ms com.br
	○	11ms com.br
	○	20ms (200.167.43.9)
	○	32ms (200.244.214.13)
AS4230	○	36ms net.br
	○	34ms net.br
	○	35ms net.br
AS54113	○	32ms (151.101.56.193)
AS unknown	▽	(151.101.56.193)

REMOVE FROM MAP

ATLAS RESULT 201.52.240.184 → AS28573 AS4230 AS54113

SUMMARY

AS28573	16730	(201.52.240.184)
	◇	1ms (192.168.130.1)
	◇	1ms (192.168.131.1)
	◇	2ms (192.168.0.1)
	◇	9ms (10.17.0.1)
AS28573	○	13ms (201.6.36.1)
	○	13ms com.br
	○	14ms com.br
	○	15ms net.br
	○	15ms net.br
AS4230	○	17ms net.br
	○	14ms net.br
	○	12ms net.br
AS54113	○	14ms (151.101.56.193)
AS unknown	▽	(151.101.56.193)

REMOVE FROM MAP

ADD A LOCATION

- São Paulo, BR-27
- Belo Horizonte, BR-15
- São Paulo, BR-27
- São Paulo, BR-27
- São Paulo, BR-27
- São Paulo, BR-27
- São Paulo, BR-27
- São Paulo, BR-27
- Houston, US-TX
- Houston, US-TX

instant changes in map and...

all other hops that have this host

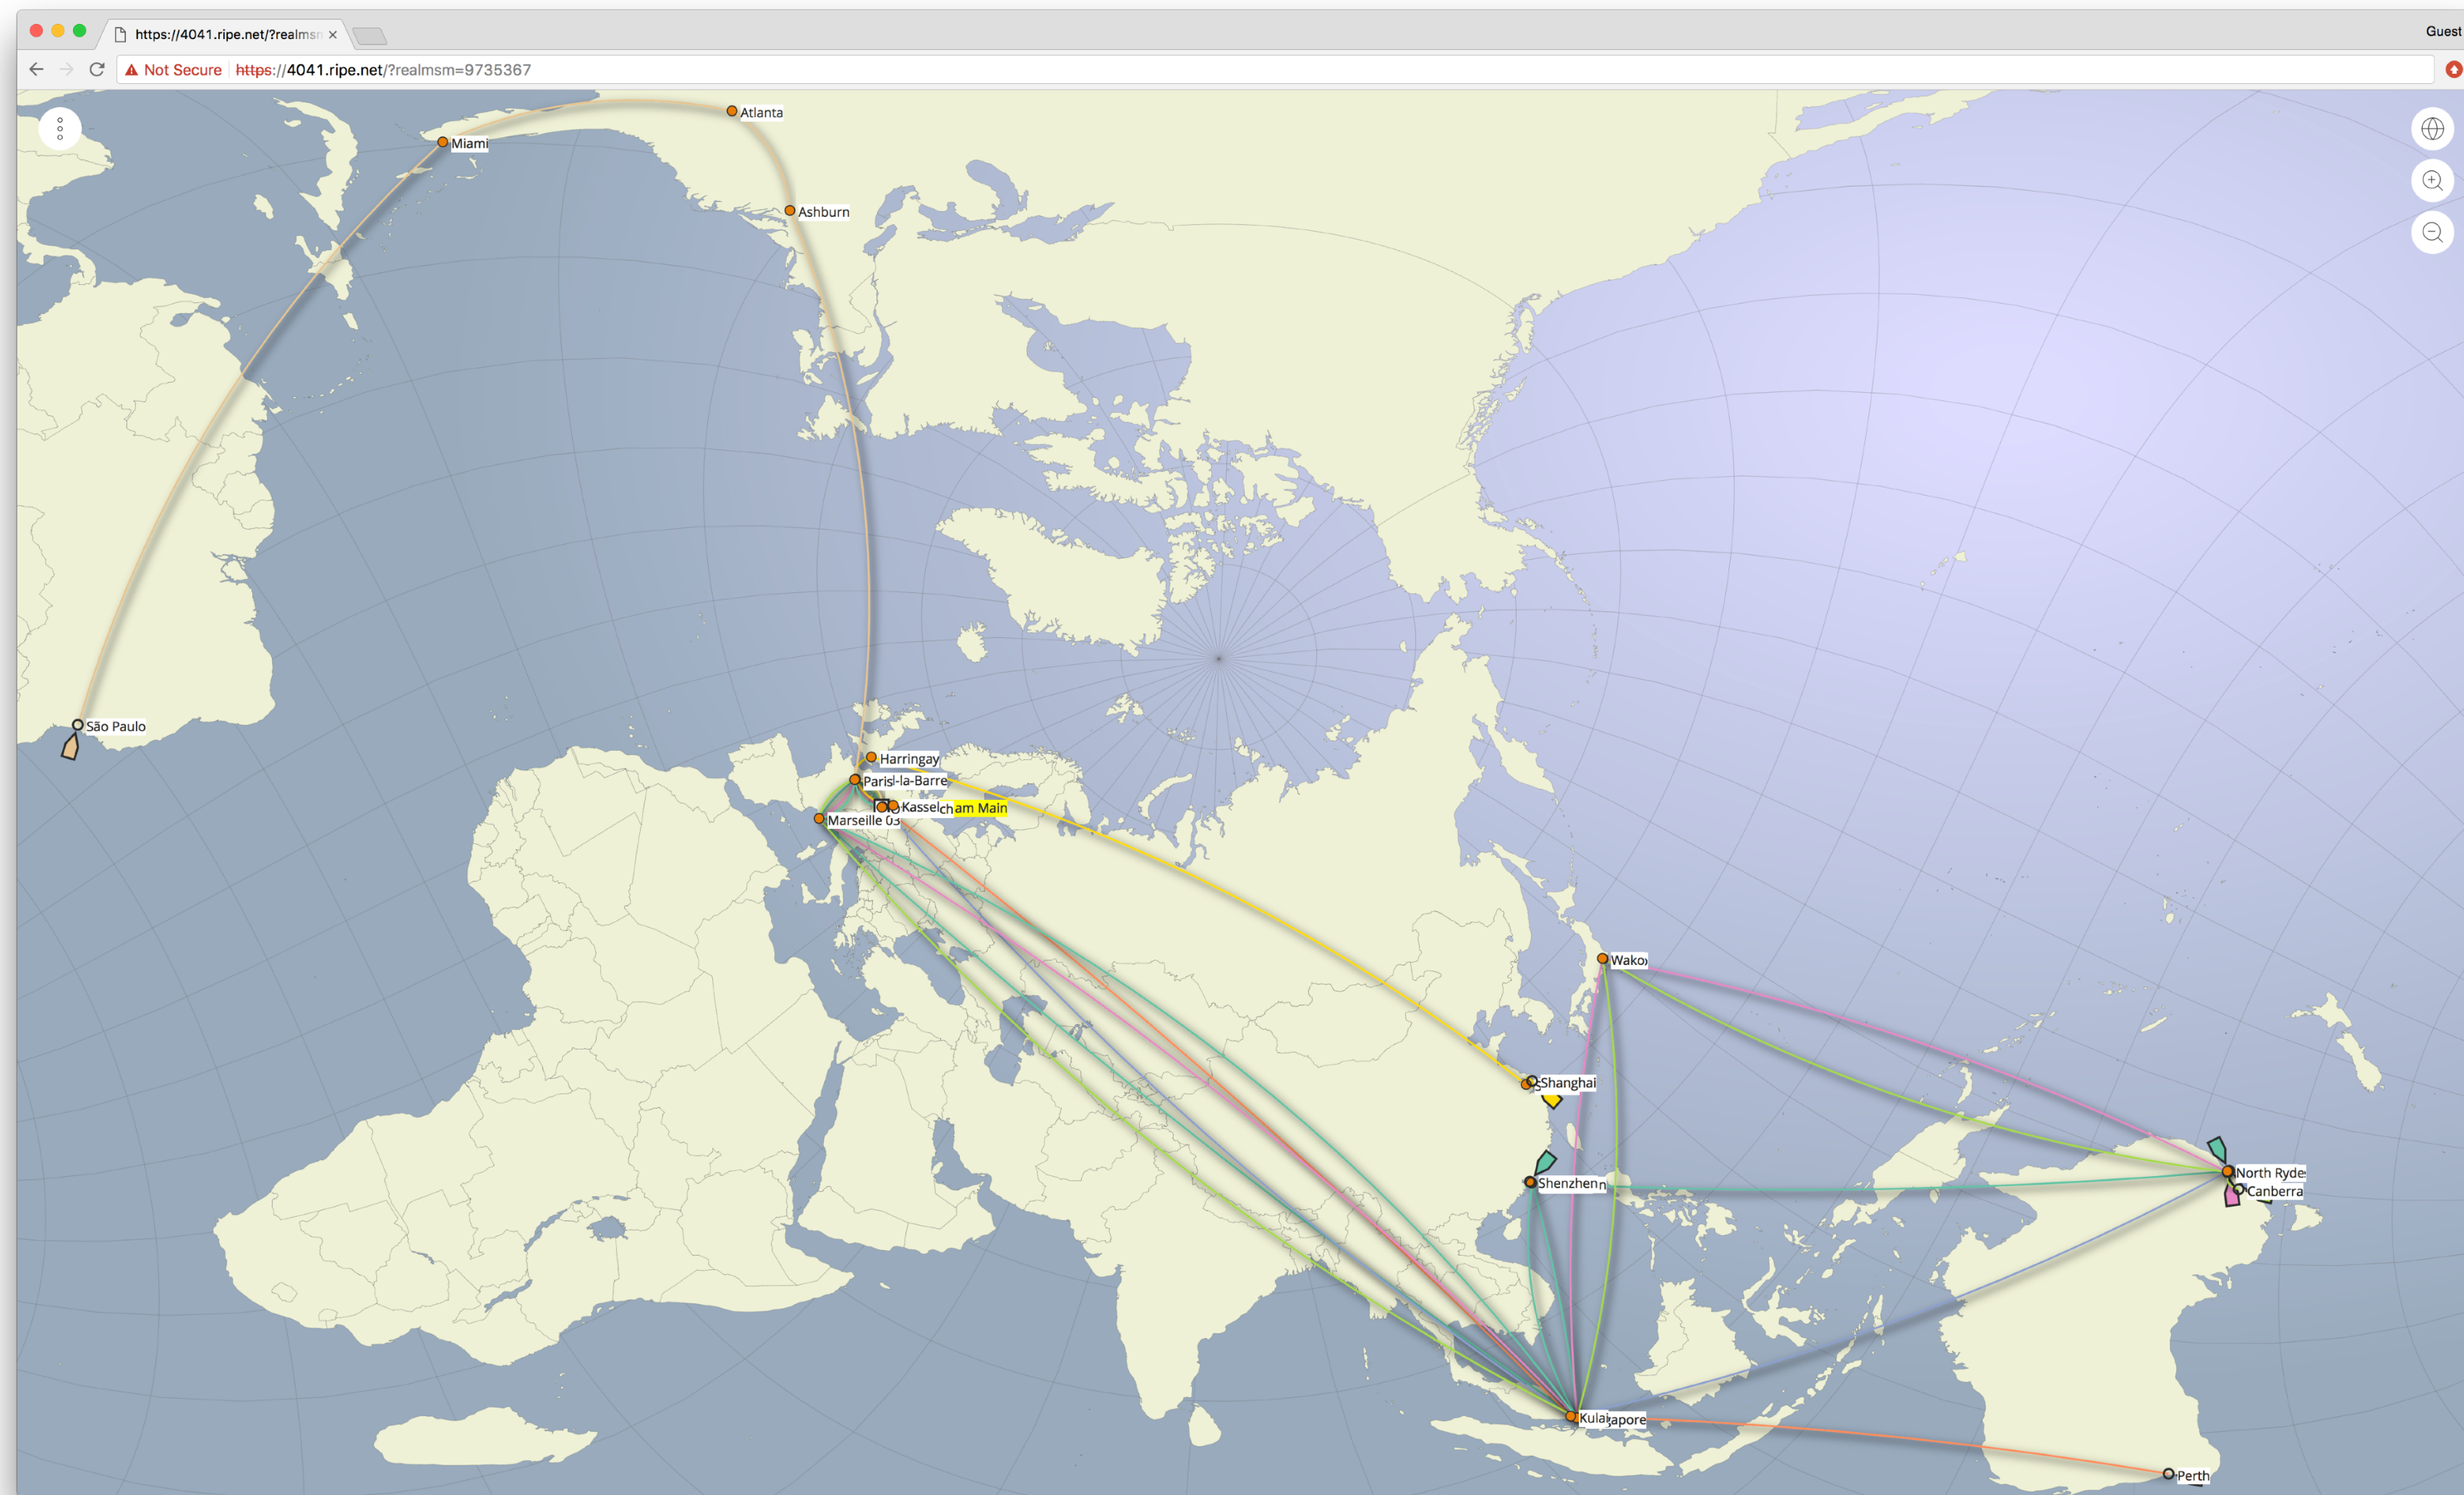
Future Work



New Inference Engines

- Increase research collaborations
- Integrate RIR data
- Reverse DNS engine
- Attempt to separate infra-structure IP addresses from end-user IP addresses
- Automated IP address discovery (IPv6!)

<https://openipmap.ripe.net>



jdenhertog@ripe.net
[@density215](mailto:jdenhertog@density215.com)