

Characterizing International Routing Detours


Anant Shah, Romain Fontugne, Christos Papadopoulos



Work supported by NSF #CNS1305404, DHS #D15PC00205,
Cable Labs and the Australian Government.

Introduction

Geographic Routing Anomalies




Detour between 2 floors of same building in Singapore
Lack of peering between NTT & Tinet



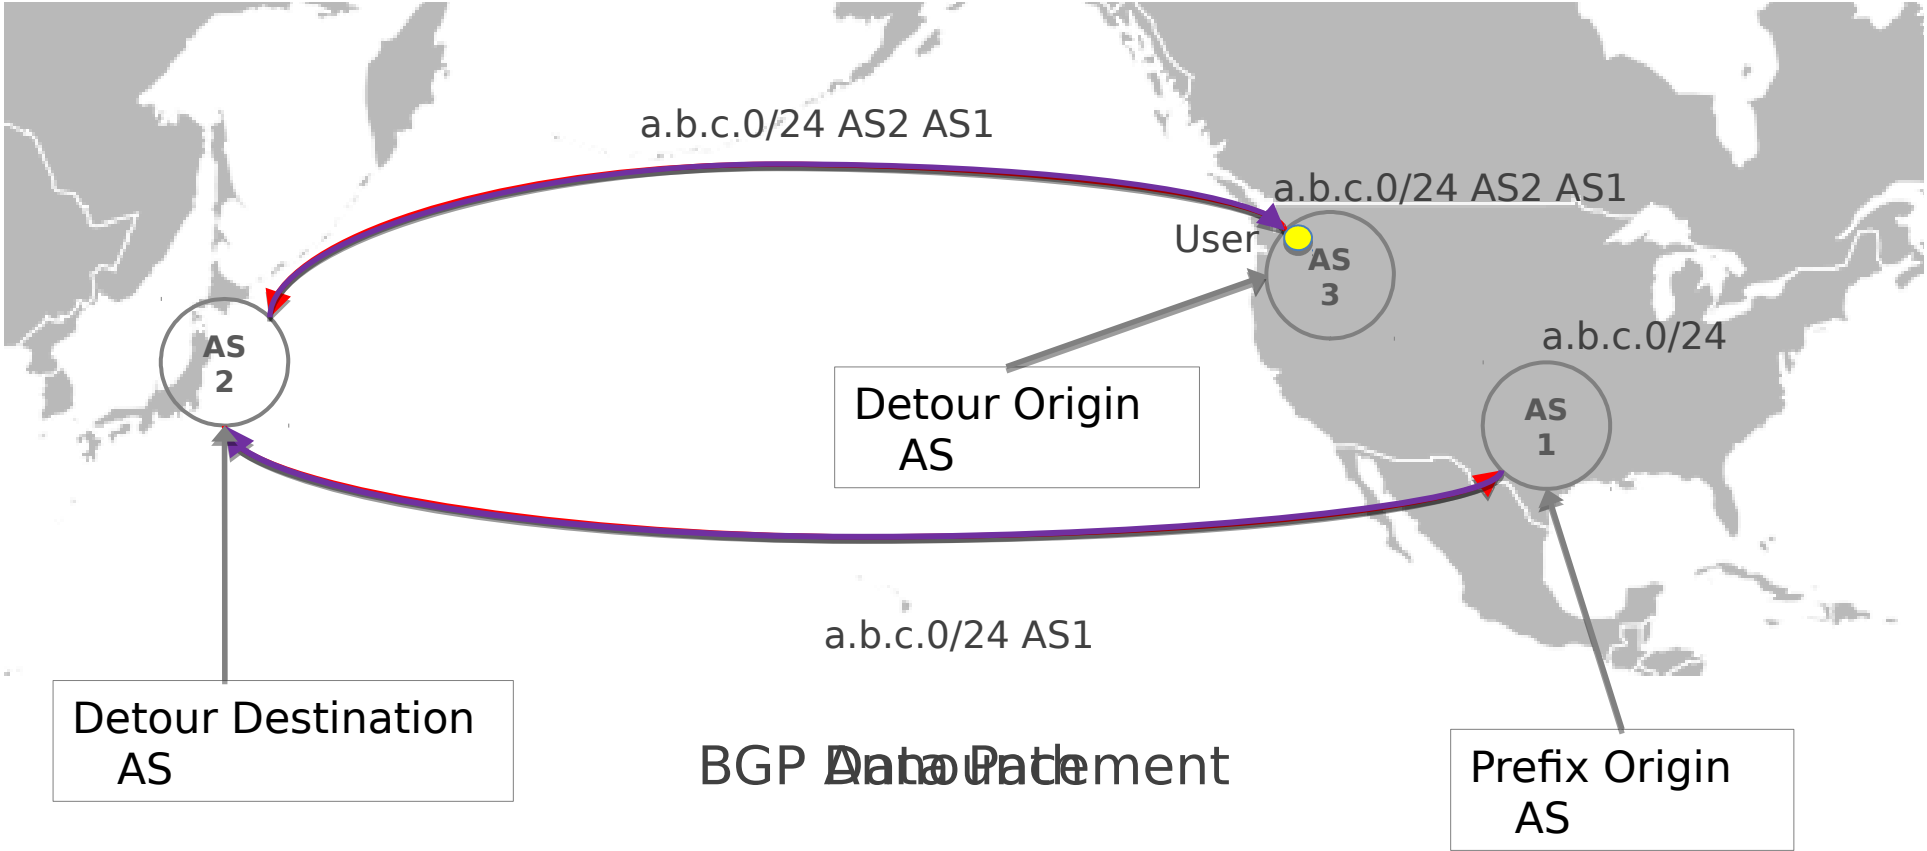
- It has been observed that in some cases routing paths do not take the geographic route as expected

International Routing Detour

A path that originates in one country,
cross international boundaries
and
returns back to origin country



Detour Example



Why do Ops care?

- Determine regulatory compliance
 - State authorities are getting conscious about how national traffic is routed
 - ISPs will need tools to show traffic is being kept local
- Detect network problems
 - A failure in network might lead to such anomalous routes
- Access traffic sniffing potential
 - Malicious ASs can announce bogus routes that lead to detours
 - If a detour occurs, knowing the detour destination AS can help evaluate legitimacy
- Locate areas of sparse network presence
 - Lack of infrastructure can force traffic to be detoured
 - Evaluating the impact on security and latency can motivate further deployment


What's Missing?

- Existing free solutions to detect such events work **post mortem**
- Systematic characterization at **global scale** not publicly available
 - Which countries are more affected?
 - Who causes such events most?
- Datasets and tools to detect anomalous routes are
 - **Unavailable** to public or
 - **Uncorrelated** (Not designed to overlap)

Goals

- Create a fast (near real-time) methodology to detect detours
- Provide further characterization
 - Impact/visibility
 - Duration
 - Stability
- Make the analysis publicly available

Operations Use



- Debug cause of large latencies
- Share info on mailing lists
- Meet regulatory compliance
- Find areas where more infrastructure can be deployed
- Short-lived detours can indicate attacks or misconfigured BGP 'fall back' routes

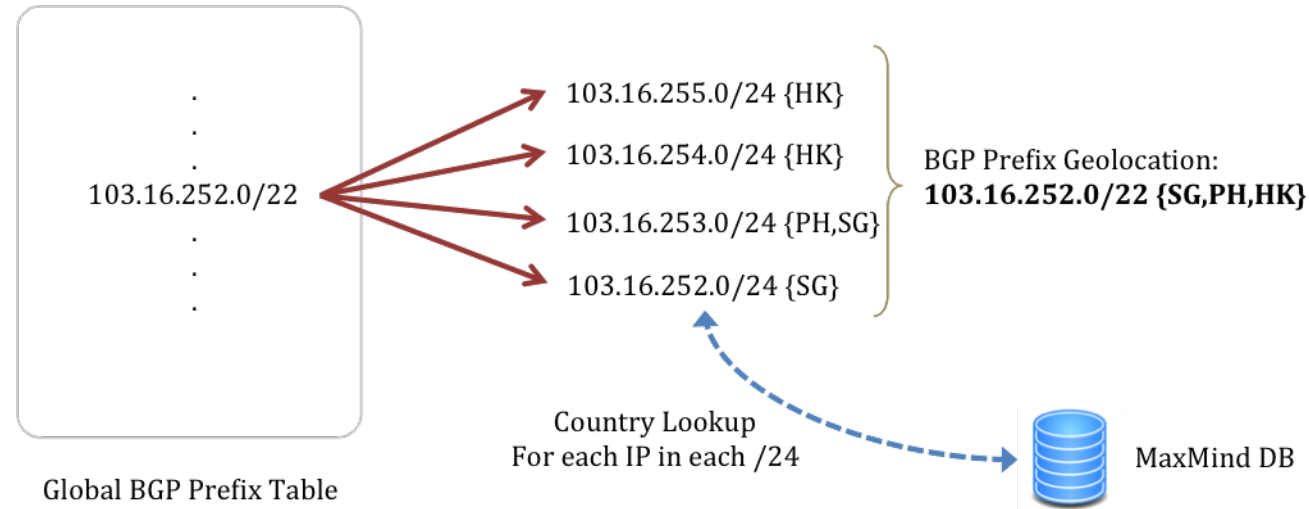
Using Control plane

- Using control plane for inferring routing paths is more scalable
 - Less data
 - Each router provides a global routing view
 - No extra traffic is generated
- Routers forward data based on their knowledge of routes learnt from BGP as 'AS paths'
- To detect **detours in control plane** we need to first create a mapping from AS to country

Definition of AS Geolocation

- Geolocation of an AS is defined as the **presence** of an AS in a country
- An AS has presence in country C if:
 - It announces prefixes that geolocate to C, or
 - Has presence at an IXP located in C, or
 - Has infrastructure IPs that geolocate to C

Step 1: Geolocate All BGP Prefixes




Important: we are interested in **country-level geolocation**

- To geolocate a BGP prefix we first geolocate all constituent /24s using Maxmind by looking up all IP addresses
- Prefix country geolocation is the set of resulting countries

Step 2: Add IXP and Peering Presence

- We parse 300+ IXP websites to gather participant lists
- Add peering mappings from PeeringDB
- Packet Clearing House IXP participant datasets

Where Do ASes Geolocate?



- ASs that geolocate to a single country: 88%
- Possible reason:
 - Organizations use different ASNs in different countries
 - Most ASs are small institutions

Detour Detection Methodology

1

Generate country level path

- Map each AS in the AS path to a country set
- Select paths that start and end in same country
- Eg. {US} {US,CA} {US}

2

Filter paths with Peering

- Detours might not occur if Detour Origin AS and Detour Return AS are peers
- Discard detour if it could be avoided by peering
- We do this at the cost of false negatives

3

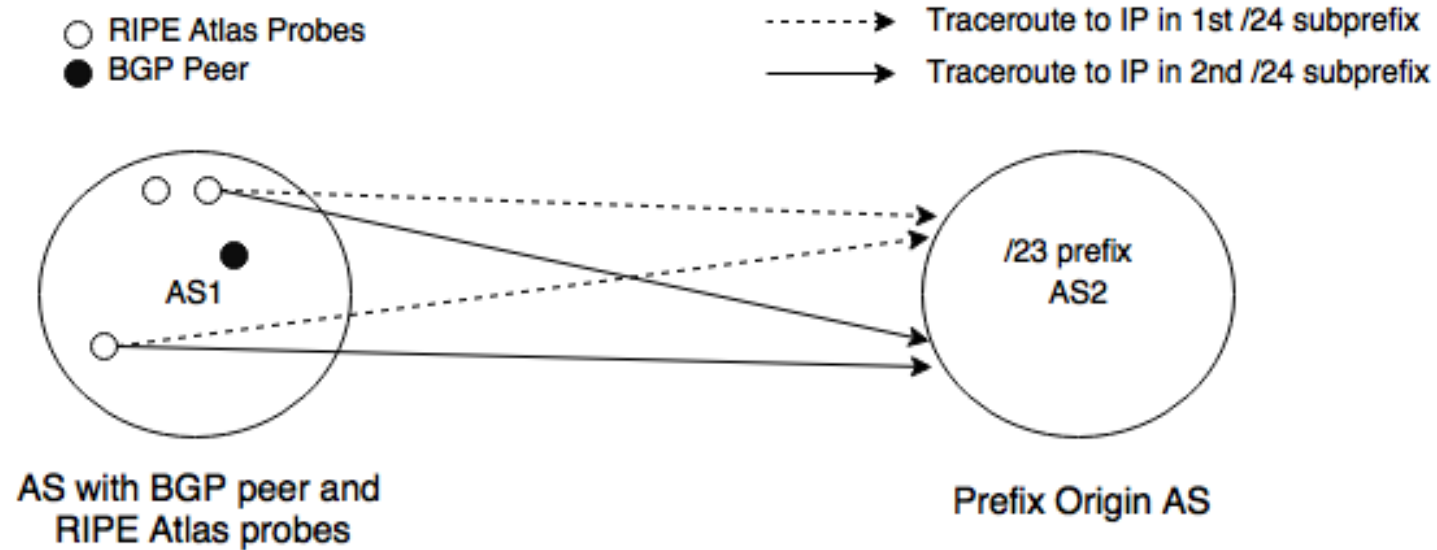
Keep track of the duration

- Helps in characterization process

Validation using RIPE Atlas

- When a detour is detected (control plane), run corresponding traceroute (data plane) using RIPE Atlas probes
 - From same country and same AS
- Check if the traceroute and detour see similar AS path
- Validate if same detour is seen on both planes

Data Plane Detour Validation



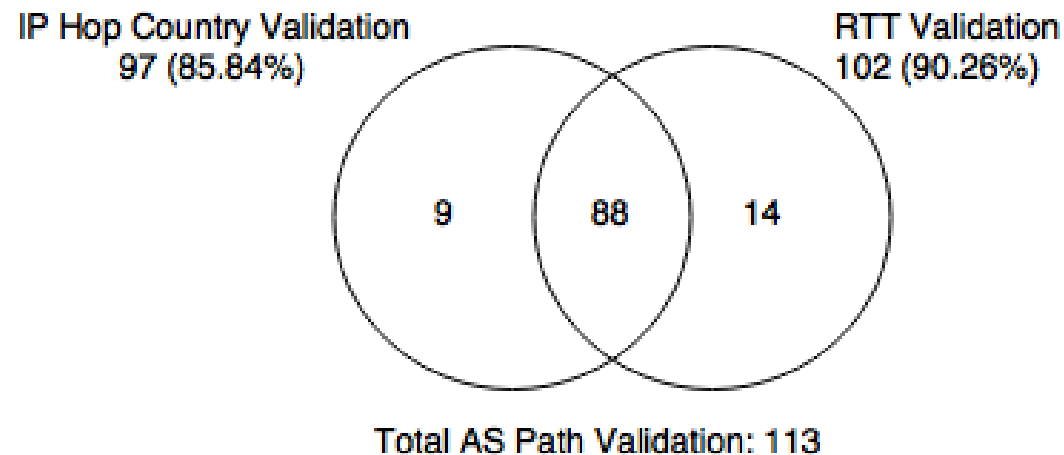
- Probes are in the same AS as BGP peers
- Target IPs belong to different constituent /24 prefixes of a BGP prefix

Detour Validation Tests

- Country-based:
 - Geolocate each IP in the traceroute and check if the expected detour is seen
- RTT-based:
 - Detect an order of magnitude jump in RTTs of consecutive hops observed in traceroute

Validation Results

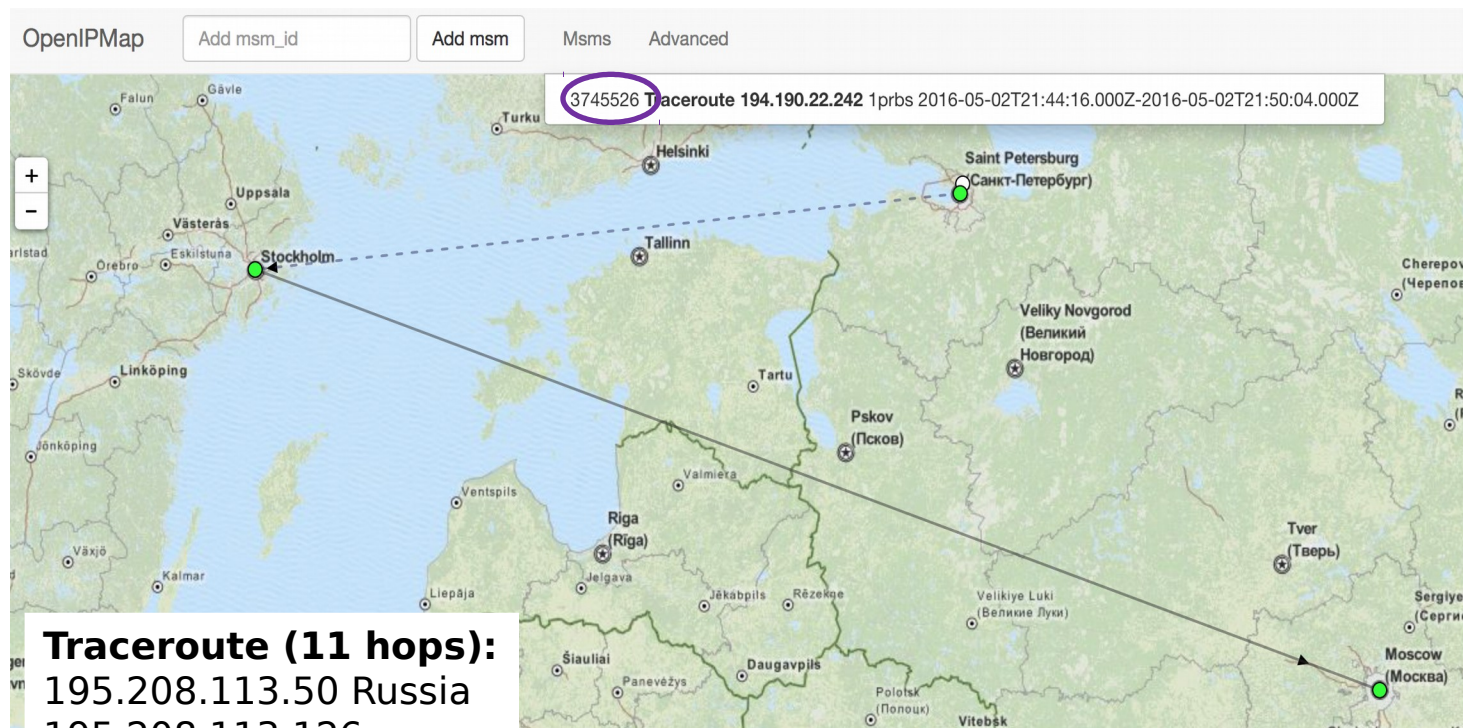
- Ran live detour detection on May 2nd, 2016 for 12 hours
- Accuracy of 85.8% and 90% respectively with the 2 methods



Visualization using OpenIPMap

Most stable detour on May 2nd, 2016

- Step1: Detected detour in control plane
- Step2: Launched traceroute to prefix 194.190.22/24 from AS3277 using RIPE Atlas
 - Within 30min window of detection
- Easy visualization with OpenIPMap possible



·
·
109.105.102.45 Sweden
·
79.104.235.190 Russia
·
194.190.22.242 Russia

BGP Path

BGP Path: 3277 3267 2603 3216

3277
(RUSnet
)
{**RU**}

3267
(SIIT&T)
{'NL',
RU}

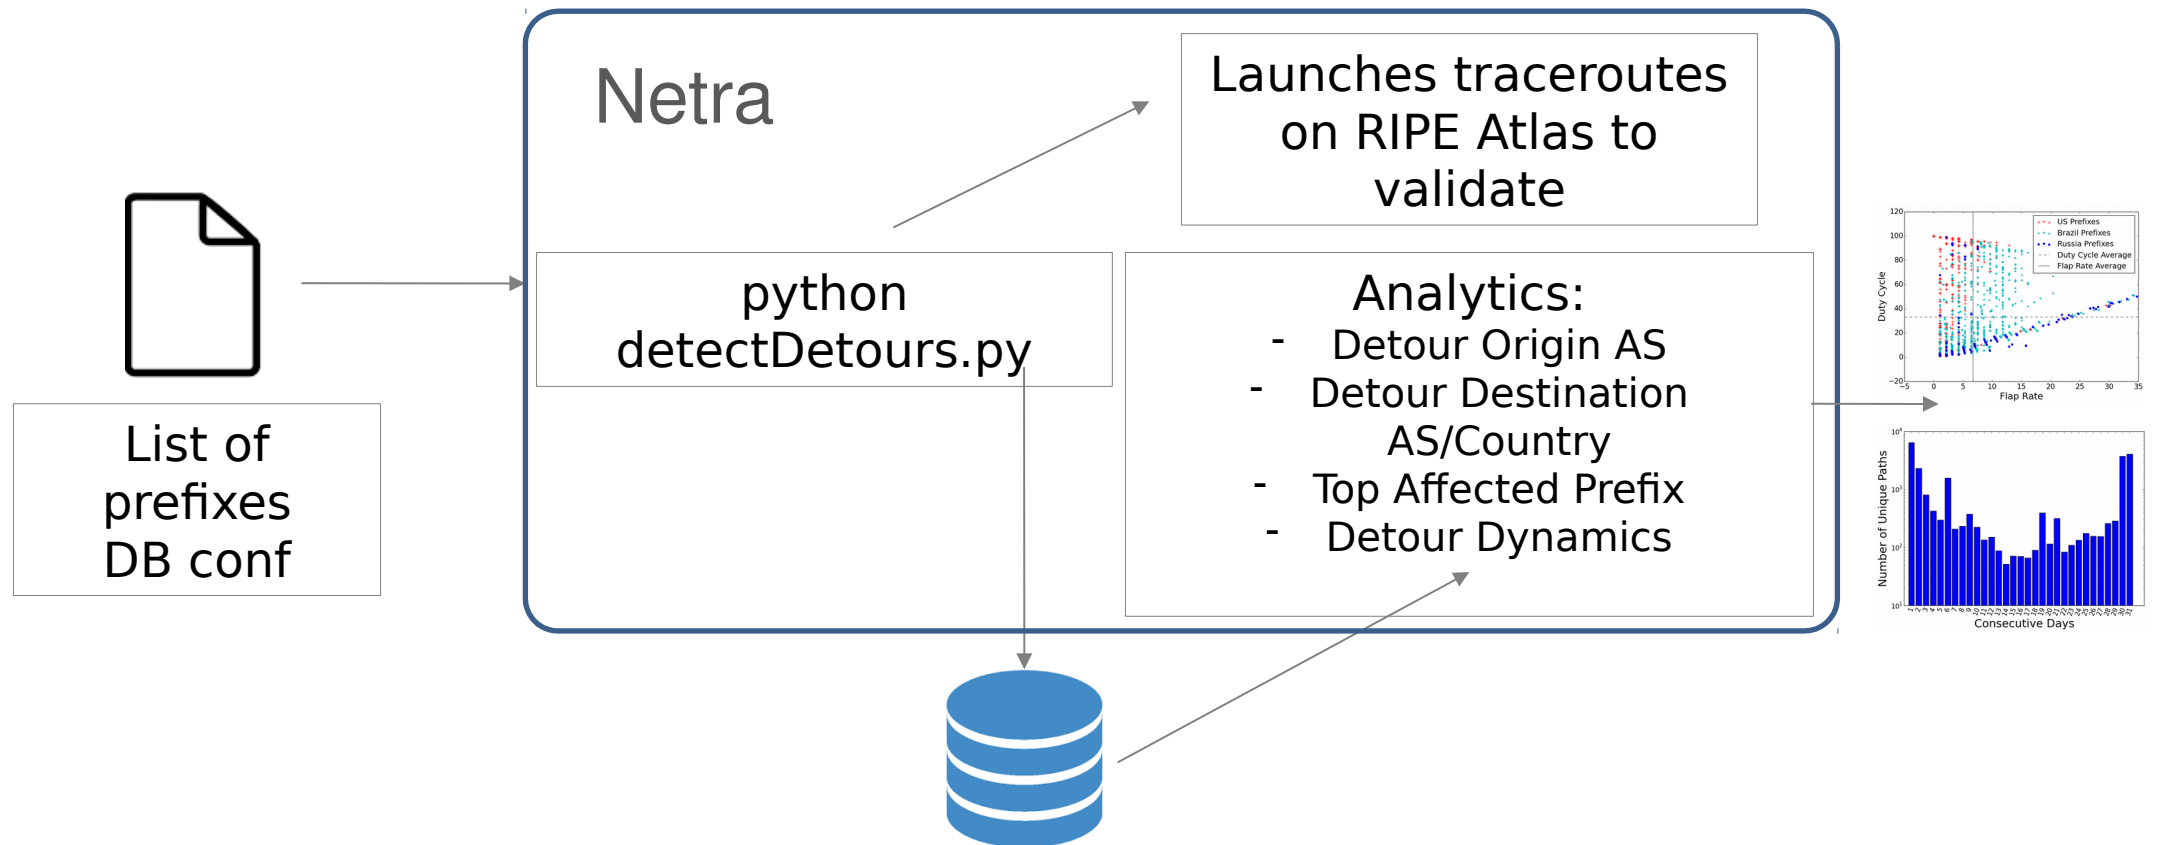
50002 2603
(NORDUnet
)
{'DE', 'IS',
'US', 'DK',
'NO', 'CH',
SE, 'NL',
'GB'}

3216
(Vimplecom)
{'DE', 'CN',
'US', 'EU', 'SE',
'PL', 'GB', 'NL',
RU}

50002
(Renaissance
Insurance)
{**RU**}

Netra

- Use Netra¹, Launch analysis for your prefixes:



¹github.com/akshah/netra

Results on historical BGP data

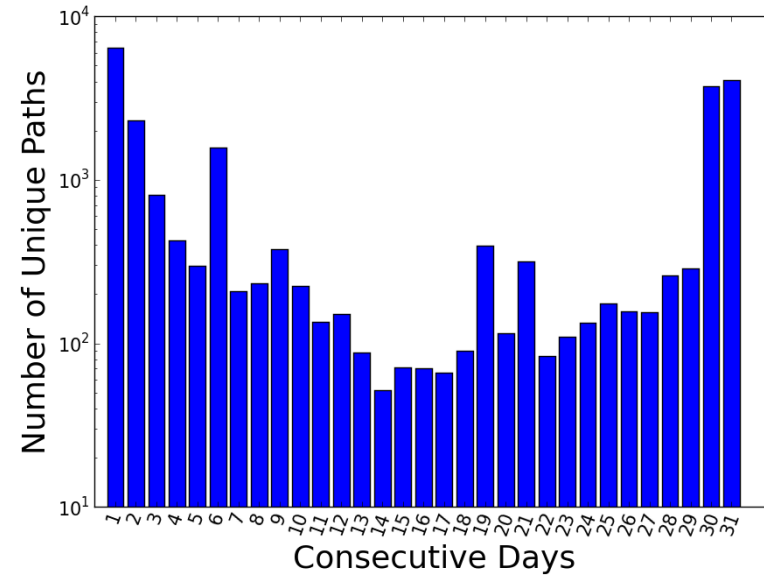
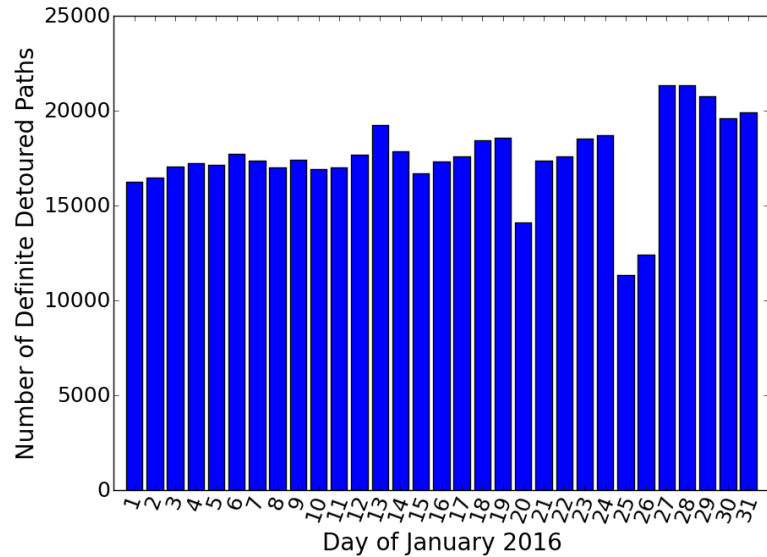
- Datasets: 1 month BGP data from January 2016
 - From 416 peers, spanning 30 countries

	January 2016
Total RIB Entries	14,366,653,046
Detoured Entries	544,484
Number of Unique Detoured Entries	18,995

Characterization Metrics

Metric		Type of Detour
Duration (How long, number of continuous hours, each detour lasts)		1. Transient 2. Persistent
Detour Dynamics	Duty Cycle (Percentage of time detour was active)	1. Recurring 2. Non-recurring
	Flap Rate (Rate of detour appearance)	1. Stable 2. Unstable


Persistence of Detours



- Detours are seen throughout the month

- Most detours either last for couple of days or persist throughout the month

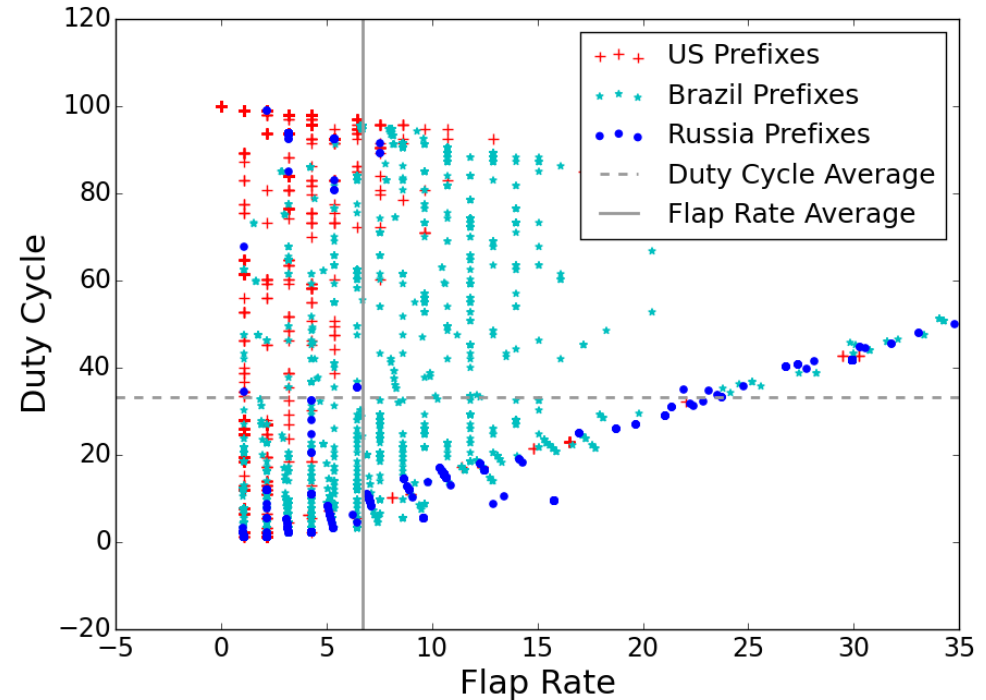
How long do Detours last?



- CDF: Number of hours detours lasted
- More than 90% of detours lasted less than 72 hours
 - Transient Detours

Detour Dynamics

- Dynamics of detours seen in different countries are different
- US, Brazil and Russia accounted for more than 90% of the detours
- US detours are more stable than Brazilian and Russian



Flash Detours

- Detours that appeared in only one RIB throughout the dataset
 - Subset of transient detours
- Some interesting prefixes that suffered flash detours:

Prefix Affected	Owner	Detour Destination
170.61.199.0/24	Mellon Bank, US	28513 (Uninet, MX)
192.230.0.0/20	Washington State Department of Information Services, US	7660(Asia Pacific Advanced Network, JP)
212.11.152.0/21	Moscow Mayor Office, RU	2603(NORDUnet, NO)
208.79.7.0/24	Security Equipment Inc, US	53185(William Roberto Zago, BR)
161.151.72.0/21	The Prudential Insurance Company of America, US	2510(Infoweb Fujitsu, JP)

New Services

New Service: BGPMon Archive

- Web-based archive with time-based BGP data retrieval
- Contains **all data** from RouteViews and Colorado State University collectors
- BGP update messages & RIBs, in MRT, JSON and protobuf format
- Enables continuous pull of data with option to receive only new updates since the previous request
- Works now, try it: <http://bgpmon.io/archive/help>

One way to use the Archive

```
from bgpDataEngine.bgpDataEngine import  
bgpDataEngine
```

```
if __name__ == '__main__':
```

```
    #Setup
```

```
    bde=bgpDataEngine()
```

```
    bde.accessToBGPMonArchive = True
```

```
    bde.accessToRVArchive = False
```


```
    bde.accessToRipeArchive = False
```

```
    #Fetch entire month
```

```
    bde.getMonth('updates','2016','03',load2db=False,\n                 collectors=['route-views.jinx','rrc00','bgpmon'])
```

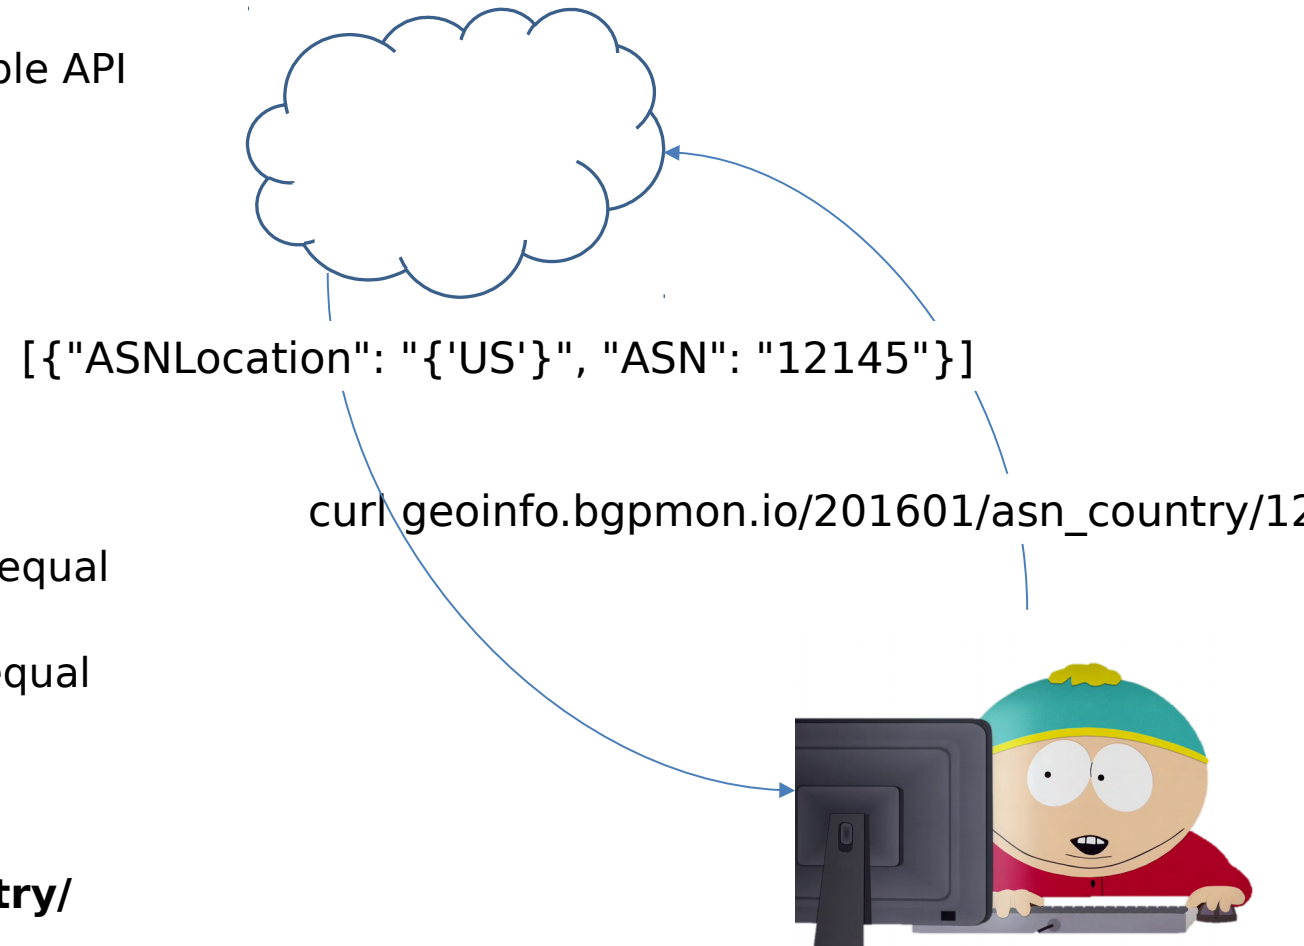
```
    #Fetch range of few days
```

```
    bde.getRange('ribs','20140701','20140705')
```



ASMap: BGPmon GeoInfo API

- Our geolocation data is publically available via simple API @
 - Works now, try it: <http://geoinfo.bgpmon.io>
- Type of queries API supports:
 - AS geolocation
 - BGP prefix geolocation
 - Prefixes announced from given AS
 - Prefixes announced from given country
 - BGP prefixes that geolocate to [more/less than/equal to] X number of countries
 - /24 prefixes that geolocate to [more/less than/equal to] X number of countries
- Example:
 - **`curl geoinfo.bgpmon.io/[MONTH]/asn_country/[ASN]`**

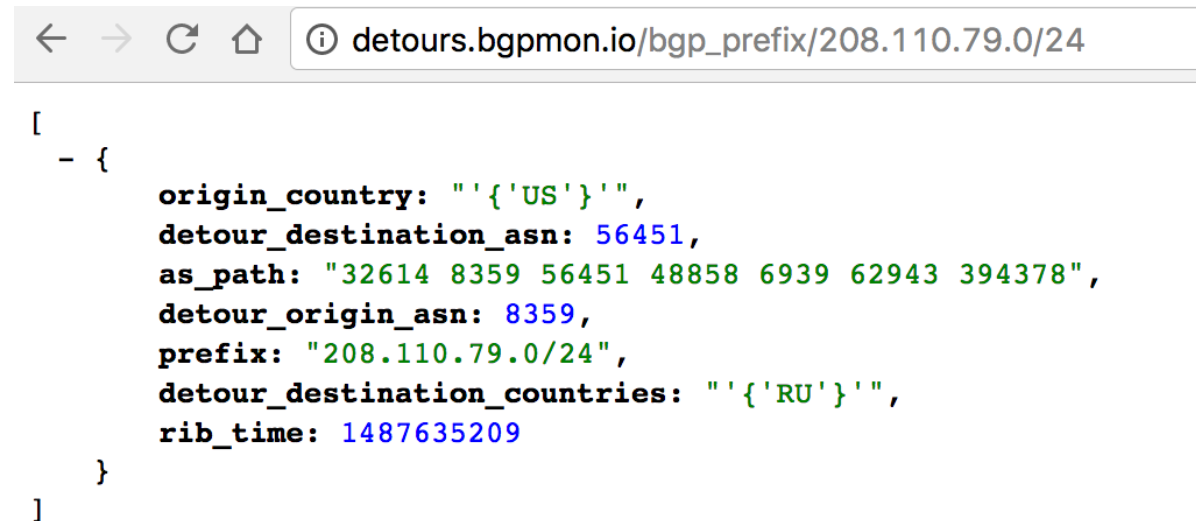


BGPmon Detours API

- Our geolocation data is publically available via simple API @
 - Works now, try it: <http://detours.bgppmon.io>

- Type of queries API will supports:
 - Detours for a prefix
 - Detours from a country
 - Detours by day

- Example:
 - **curl http://detours.bgppmon.io/bgp_prefix/208.110.7**
 - **curl http://detours.bgppmon.io/all_detoured_prefixes/20170227/**



The screenshot shows a web browser window with the address bar containing the URL `detours.bgppmon.io/bgp_prefix/208.110.79.0/24`. The page content displays a JSON array with one object. The object contains the following fields: `origin_country` (US), `detour_destination_asn` (56451), `as_path` (32614 8359 56451 48858 6939 62943 394378), `detour_origin_asn` (8359), `prefix` (208.110.79.0/24), `detour_destination_countries` (RU), and `rib_time` (1487635209).

```
[
  - {
    origin_country: "'{US}'",
    detour_destination_asn: 56451,
    as_path: "32614 8359 56451 48858 6939 62943 394378",
    detour_origin_asn: 8359,
    prefix: "208.110.79.0/24",
    detour_destination_countries: "'{RU}'",
    rib_time: 1487635209
  }
]
```


Summary

- Our methodology provides quick insight into International detours
 - At a global scale
 - Using only control-plane data
- Data plane measurements can be used to complement detected detours
- Further detection capabilities can be improved with a larger deployment of BGP peers and RIPE probes

Public Availability

- Data:
 - BGPmon archive: <http://bgpmon.io/archive/help>
 - AS Geolocation: <http://geoinfo.bgpmon.io>
 - Detours: <http://detours.bgpmon.io>
 - Send feature requests!
- Code:
 - BGPDataEngine: <https://github.com/akshah/bgpDataEngine>
 - Netra: <https://github.com/akshah/netra>
- Verify and send updates on AS geolocation:
 - <http://geoinfo.bgpmon.io/feedback>

Thank you!

akshah@cs.colostate.edu

RouteViews / RIPE RIS

- For control plane information
- A binary dump of routing tables from 400+ routers in 30+ countries

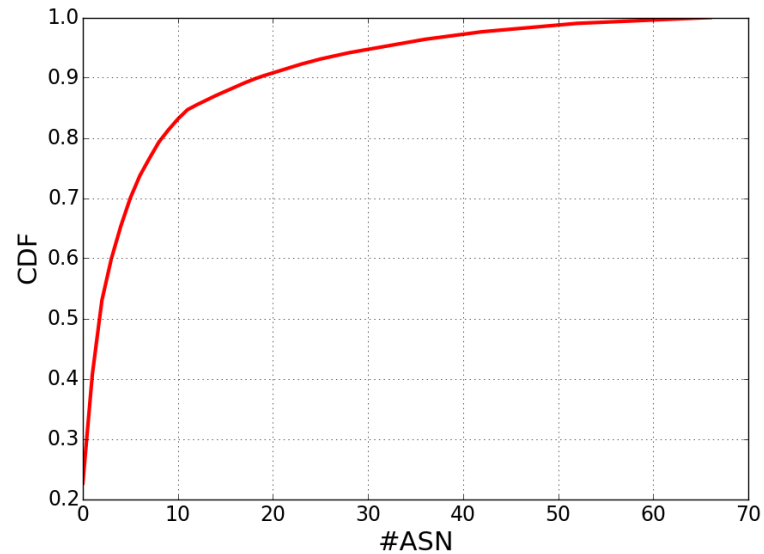


RIPE Atlas

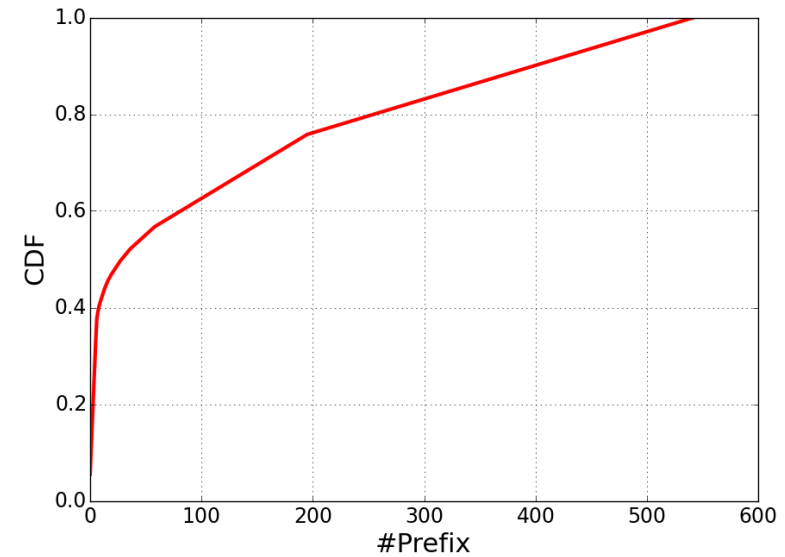
- For data plane information
- 10K+ probes in 178 countries
- Periodic measurements
 - Root Servers
 - Anchors (Special Probes)
- Users can launch Pings/Traceroutes as well



Top Transient Detour Origins and Prefixes



- 3-4 ASes originate 50% of the transient detours



- 30 prefixes account for 50% of the transient detours

Top Transient Detour Origins and Prefixes

Transient Detour Origin AS	Total %	Frequent Detour Destination AS	% to frequent destination
9002 (RETN-AS RETN Limited,RU)	22.64%	2914 (NTT America)	99.07%
6939 (Hurricane Electric,IT)	10.94%	8551 (Bezeq International)	100%
1299 (TELIANET,IT)	10.87%	8708 (RCS-RDS)	100%