Data plane BGP Hijack detection via latency measurement

- **Danilo Cicalese**
- Dario Rossi

Google Faculty Research Awards

mPlane

Workshop on Active Internet Measurements

UC San Diego

BGP hijack and anycast detection

Credits: renesys

- IP anycast
  - Syntactically equivalent to a BGP hijack in the BGP lingo [1,2,3]
  - Only difference: router authorized to advertise the prefix or not
  - We could use iGreedy [AIMS'15, Infocom'15, CoNext'15]
  - BGP hijack detection via latency measurement analysis

Reactive scan on BGP announces

• Analyse BGP feed (eg BGPstream) and issue analysis on suspicios

Problem

• BGP Hijacks are of short duration

• Control plane information may arrive late at some monitor

Proactive Internet-wide scan

•Scan all /24 prefixes every minute (or an IPv4 subset with opt-in/opt-out)

Problem

•Over 100x faster than current speed (but detection easier than geolocation)

•More challenging, hence more fun!

Anycast: detection and geolocation

Background: [Infocom'15]

IP address
(e.g. 141.101.125.63)

Active measurements

RIPE NCC

PLANETLAB

Detection/Geolocation
Anycast instances

Detect
Speed of light violations

Enumerate
Solve MIS

Geolocate
Classification

Iterate
Feedback

Enumeration/Geolocation

Detection:
speed of light violation

✔ Open Source SW+ dataset with ground truth published

Software available at http://goo.gl/7ESrCR

iGreedy SW overview

```
./igreedy -m 204.79.197.215 -p datasets/planetlab-vps -r datasets/ripe-vp.cityAccuracy-selected -b
```

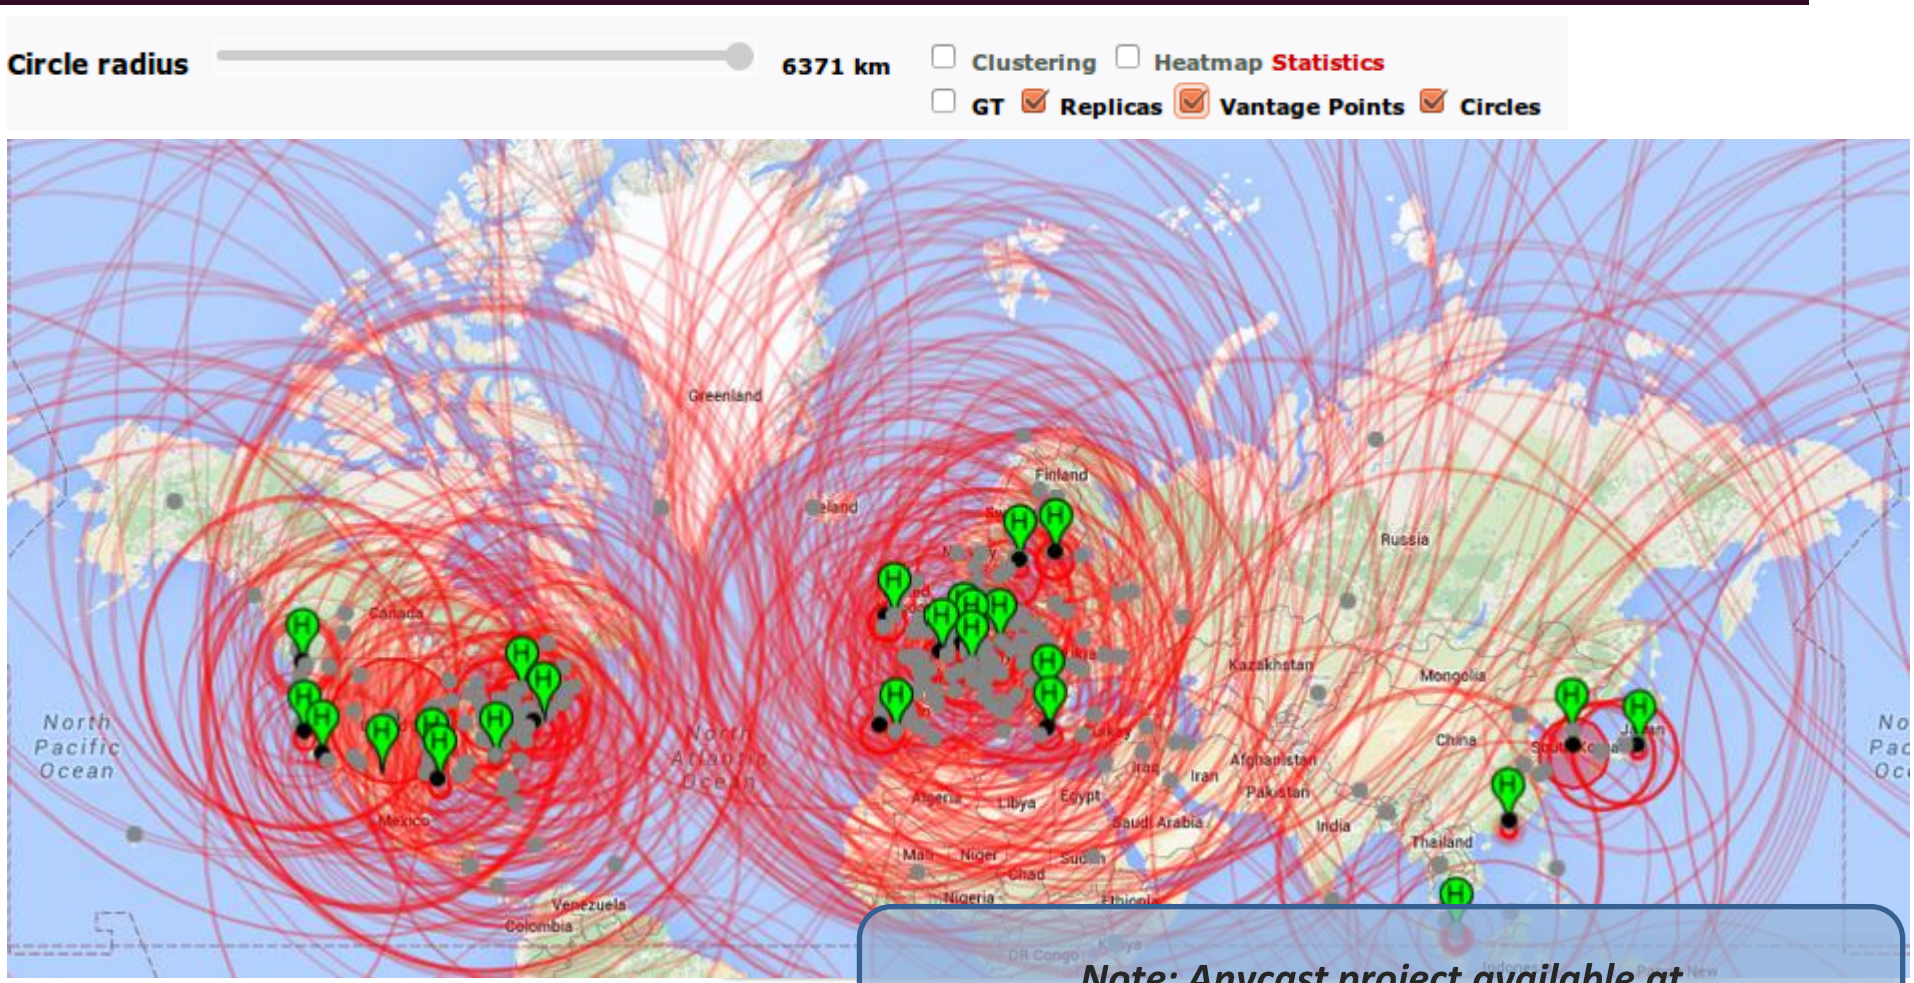

Note: Anycast project available at http://www.telecom-paristech.fr/~drossi/anycast

Anycast: scale up to a census

Background: [AIMS'15,CoNext'15]

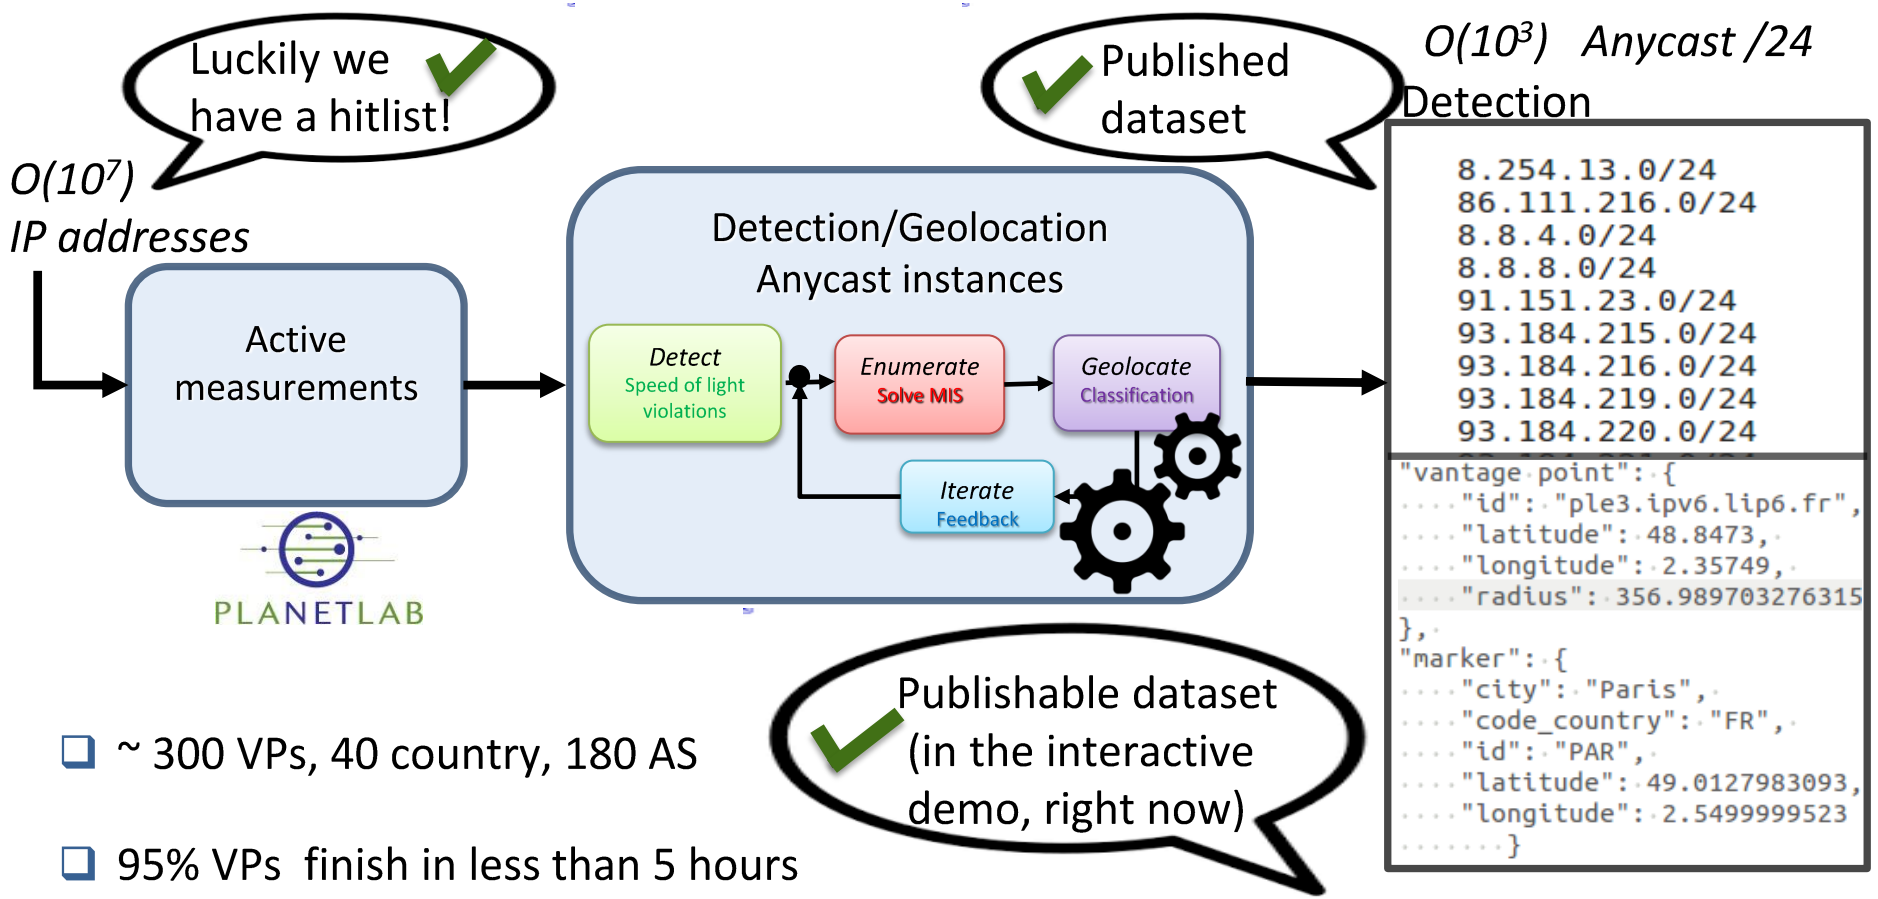

- ~ 300 VPs, 40 country, 180 AS
- 95% VPs finish in less than 5 hours
- ~ 6GB per census (~ 20MB per host)
- Analysis in 3 hours

Anycast: top 50 anycast deployements

Background: [AIMS'15,CoNext'15]

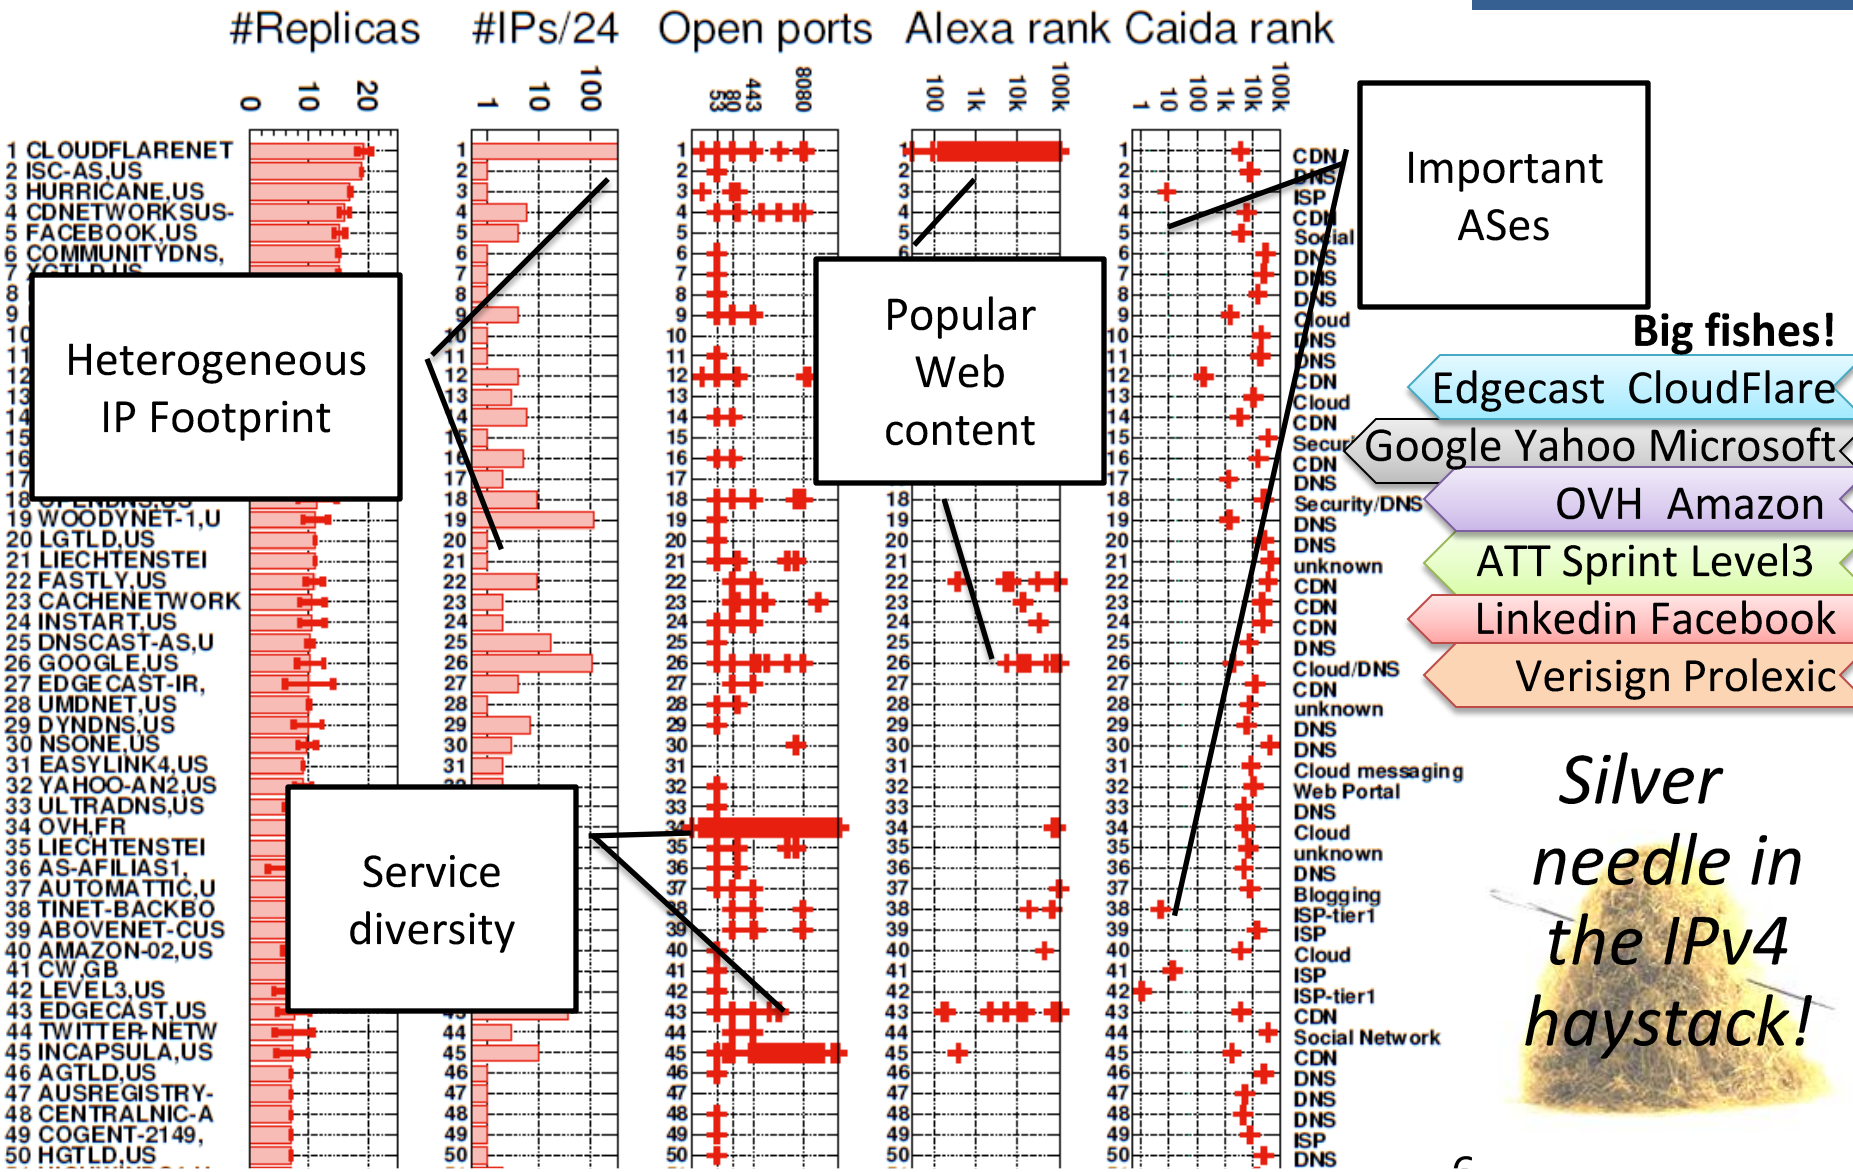

Demo overview

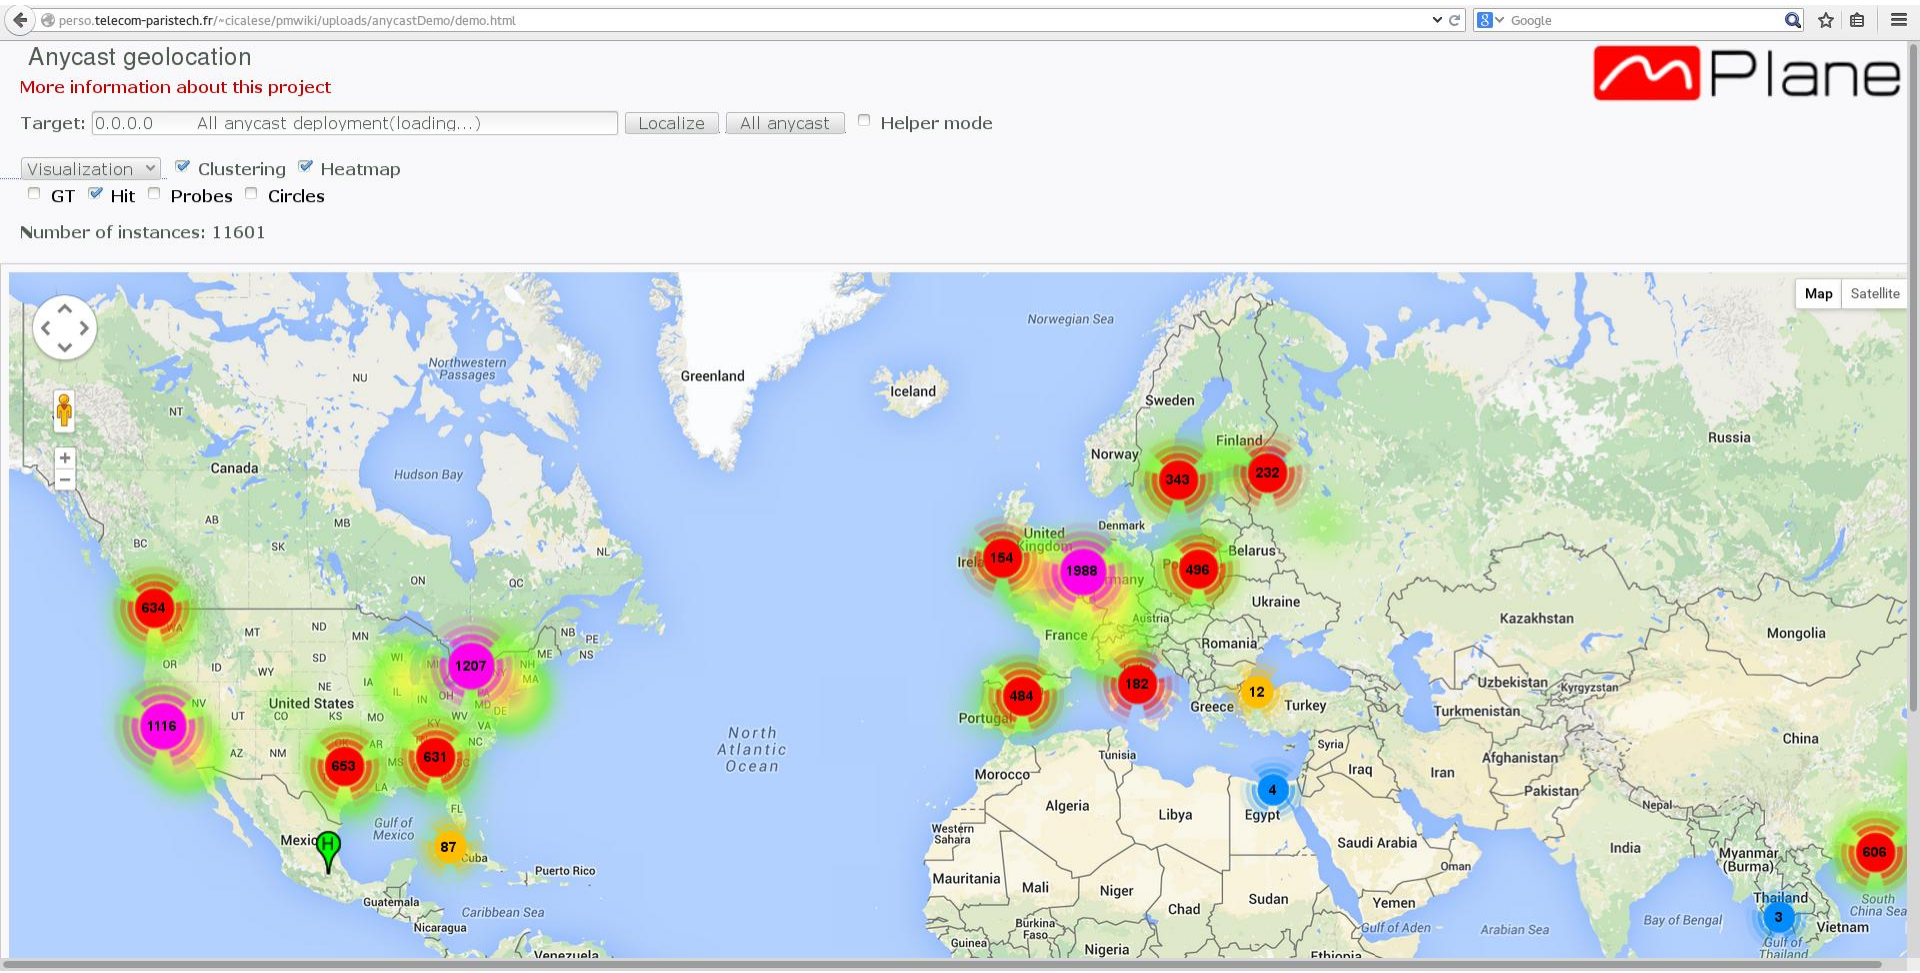

Anycast: monthly Census

Ongoing

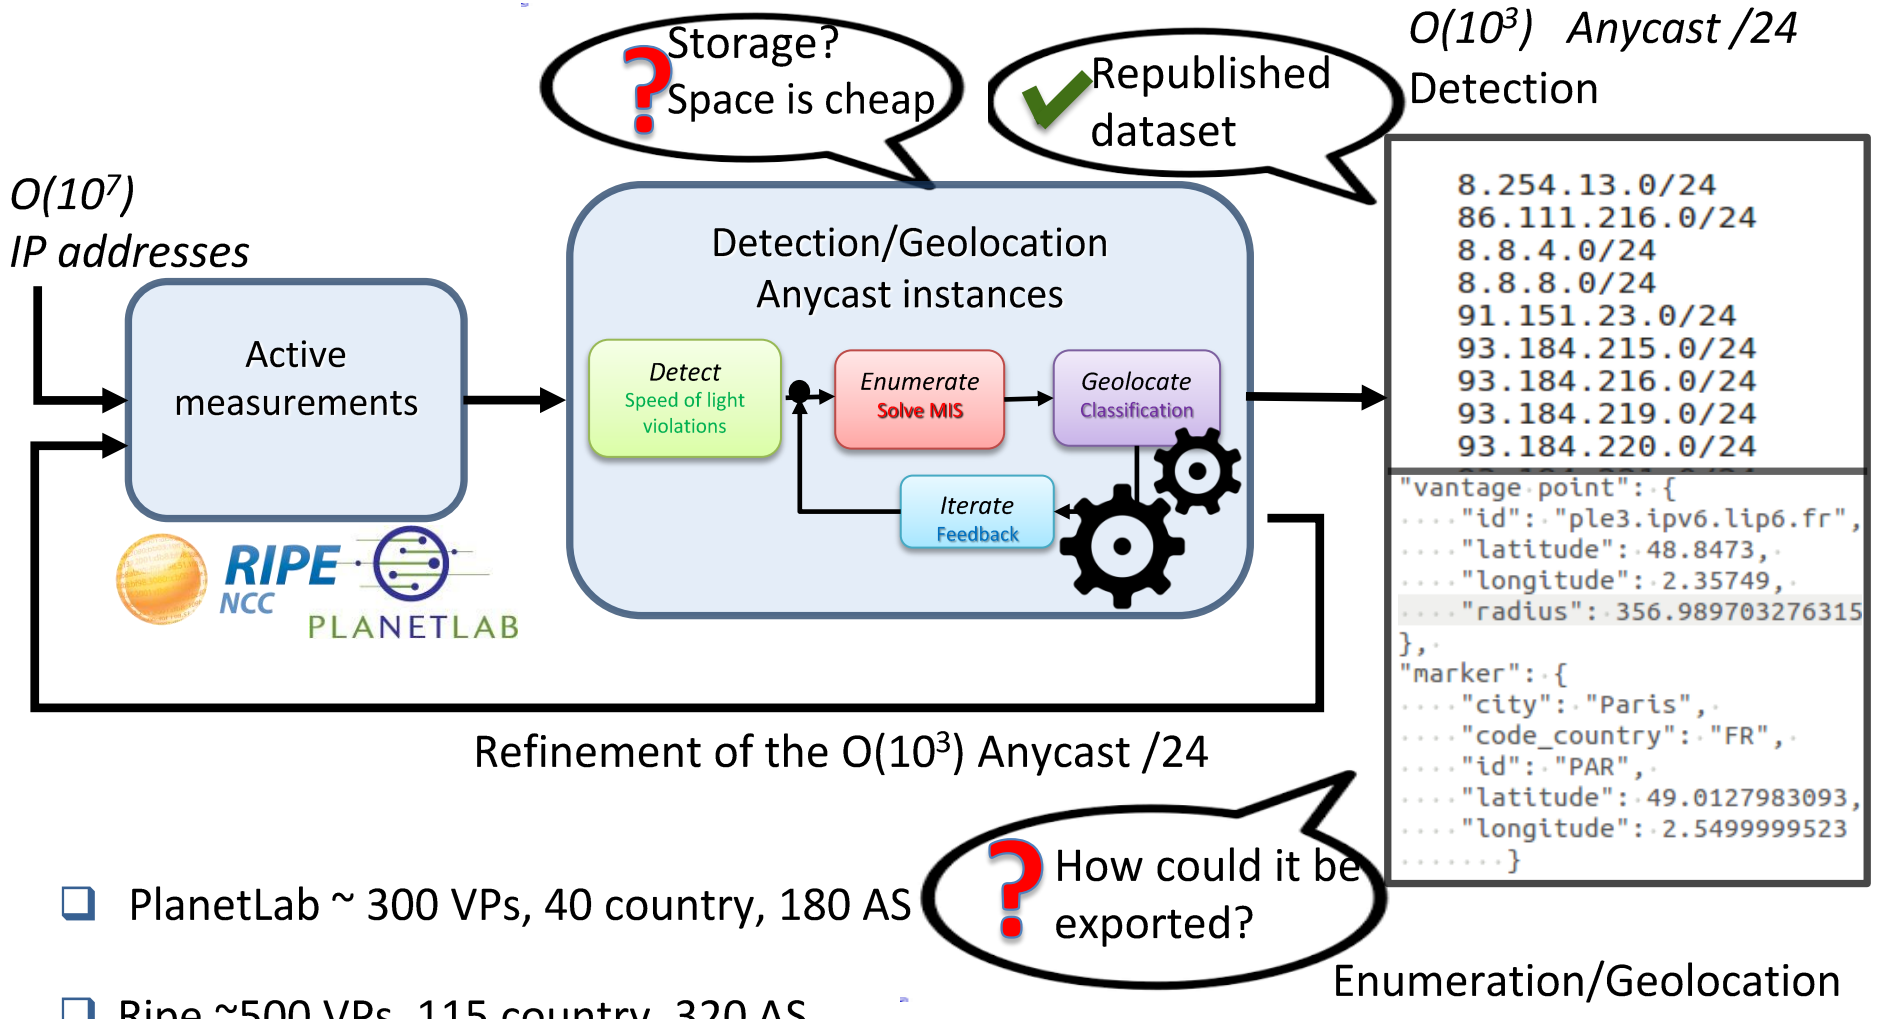

- PlanetLab ~ 300 VPs, 40 country, 180 AS
- Ripe ~500 VPs, 115 country, 320 AS
- ~ 6GB per census (~ 20MB per host)

Anycast: continuos census & BGP hijack

Next Step

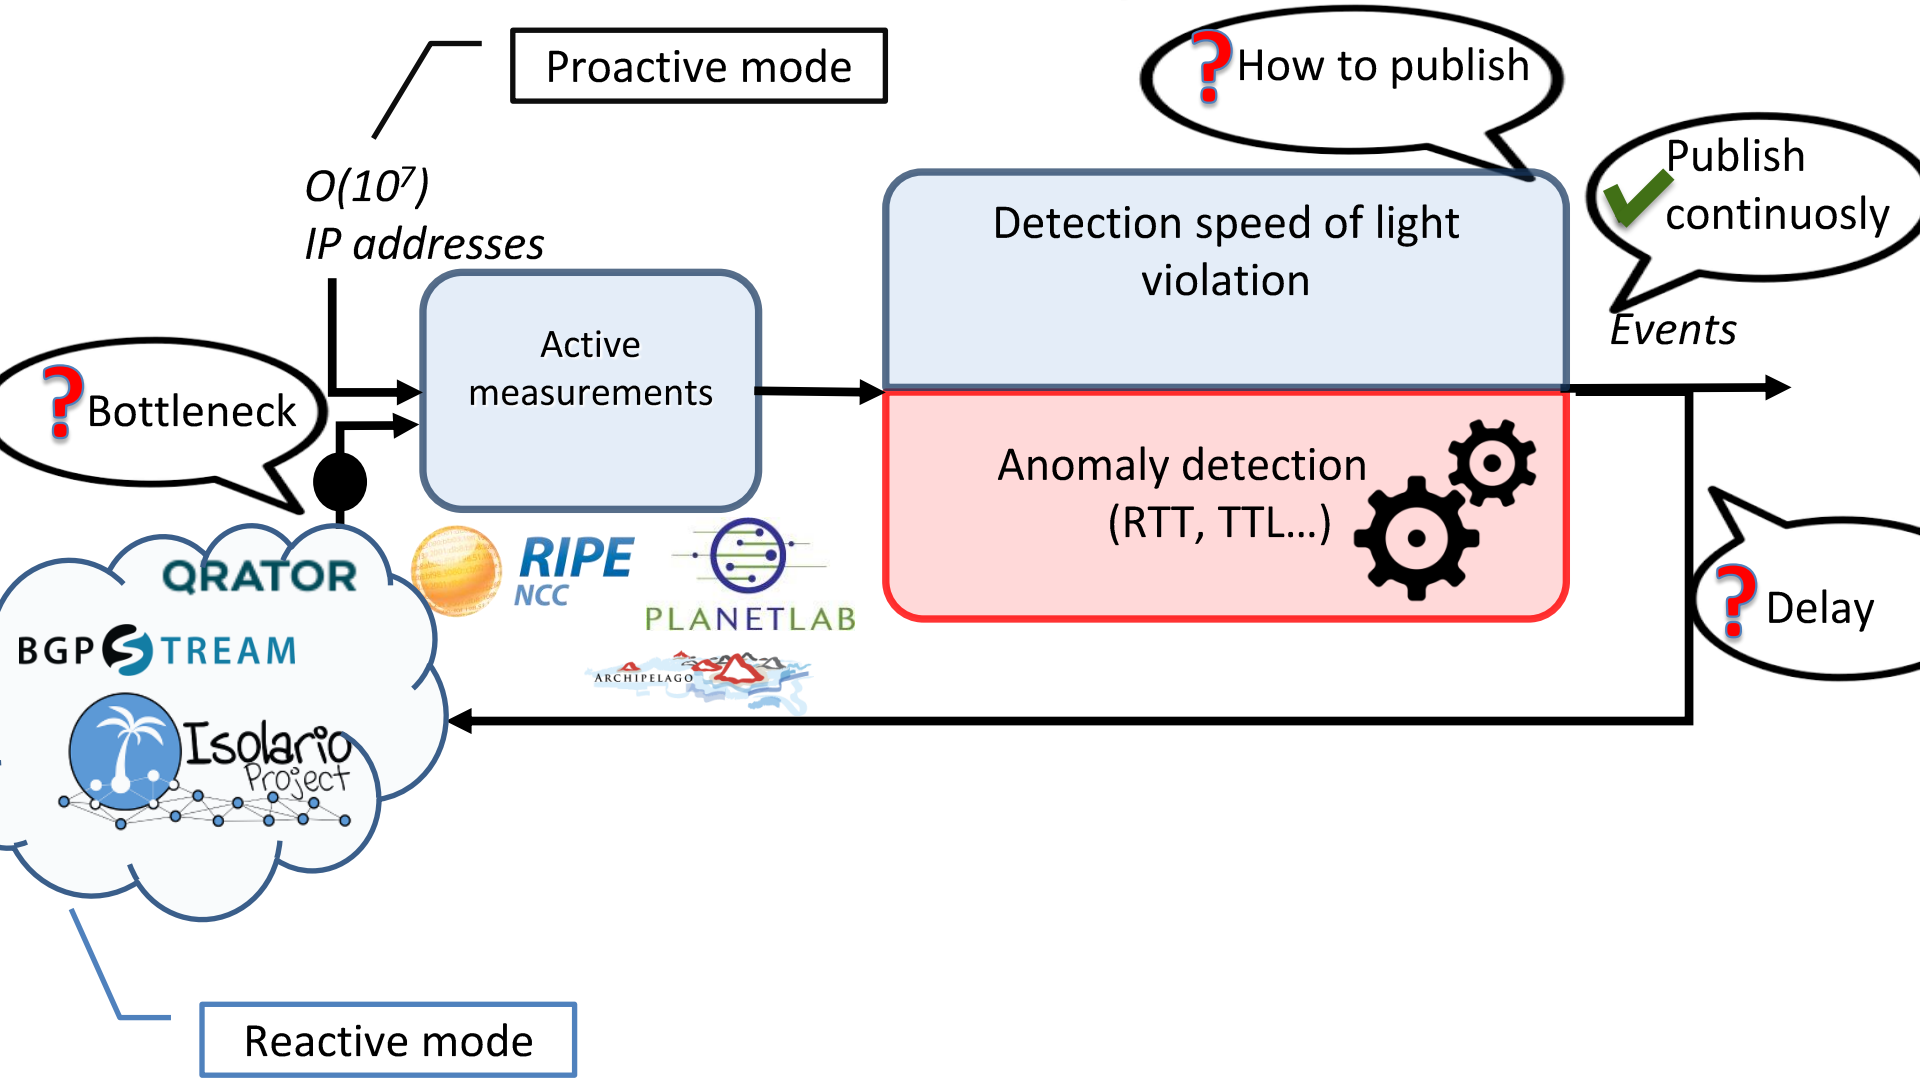

100=10x10

Clustering:

10X Space

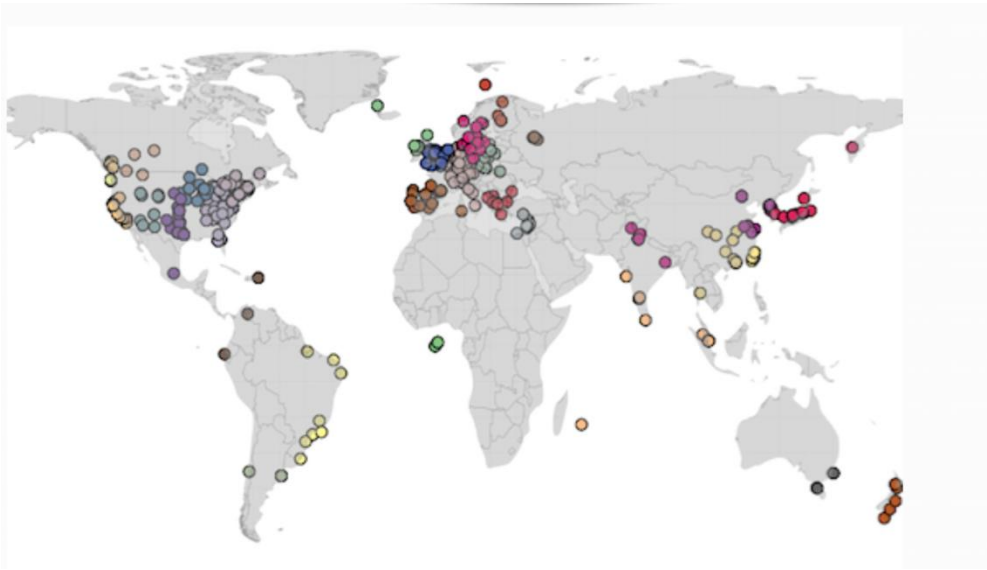

- Reason:
  - Less measurements
  - Less amount of data
  - Possible parallelization
- Centroid-based clustering:
  - K-means
  - 50 VPs per cluster
  - ~75% detection accuracy

10X Speed

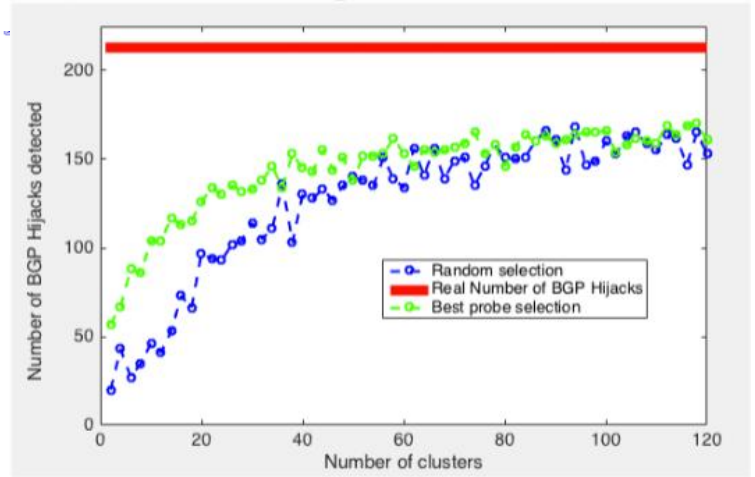

- Dynamically modulate the rate

Data plane BGP Hijack detection via latency measurement

http://www.telecom-paristech.fr/~drossi/anycast

- **Danilo Cicalese**
- Dario Rossi

Google Faculty Research Awards

mPlane

Workshop on Active Internet Measurements

UC San Diego