ProBGP:
Geographical Approximation of BGP Update Paths

Alex Ulmer

Fraunhofer IGD
Fraunhoferstraße 5
64283 Darmstadt

Tel.: +49 (0) 6151 155 – 418

alex.ulmer@igd.fraunhofer.de

www.igd.fraunhofer.de

Fraunhofer IGD

Overview

- Introduction
- BGP Update Visualizations
- Geographical Visualization of BGP Update Paths
  - Approximation of AS Data Center Locations
    - Ground Truth
    - GeoIP Clustering and Assumptions
    - Internal AS connections
  - Update Path Visualization
- ProBGP Live Demo

Introduction

- Alex Ulmer
- Fraunhofer Institute for Computer Graphics Research IGD Darmstadt, Germany
- Competence Center for Information Visualization and Visual Analytics

- Cybersecurity visualization as a focus topic for the past 4 years as part of the national research center ATHENE

ATHENE
National Research Center for Applied Cybersecurity

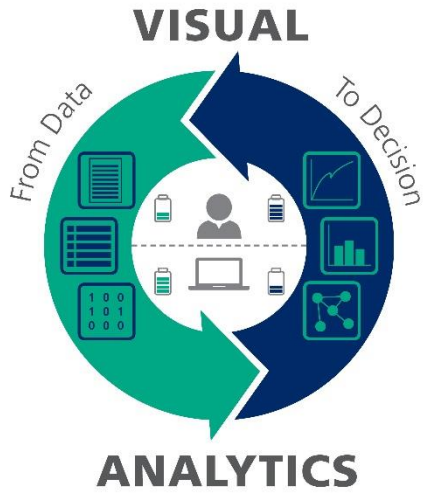

Introduction

Goal:

Present our prototype as a new way to visualize BGP update paths and discuss with domain experts on potential ways to improve the approach.

BGP Update Visualizations

- A BGP Update

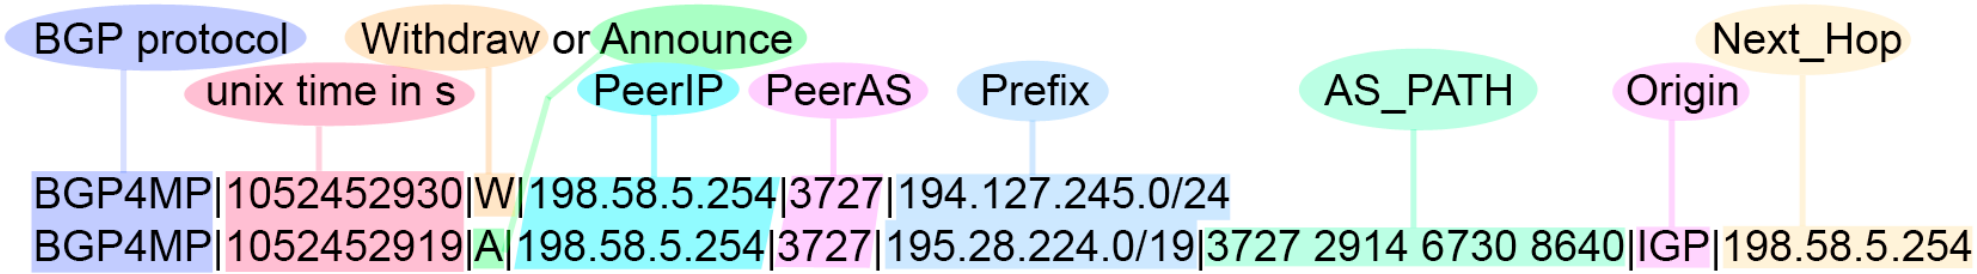

- Analyze the AS Path

- BGP Update Visualization Tools from the Research Community

BGP Update Visualizations

- BGPlay

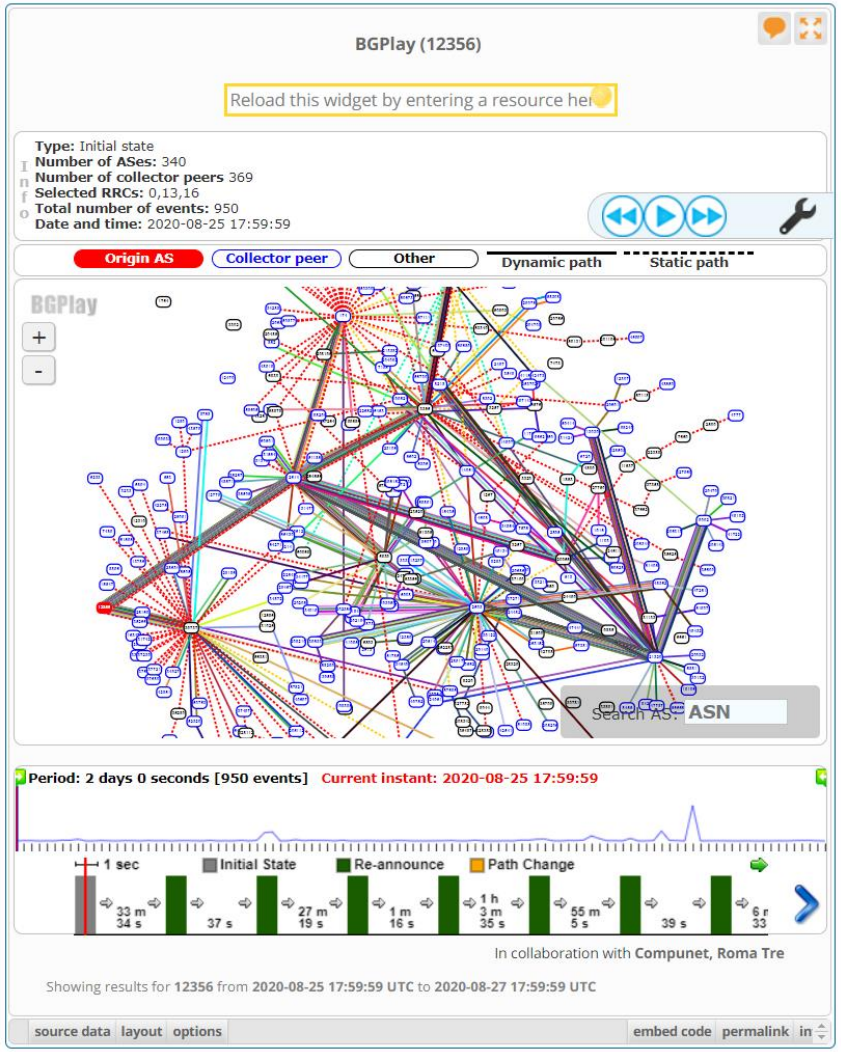

Screenshot from: https://stat.ripe.net/widget/bgplay#w.resource=12356

BGP Update Visualizations

- BPGViewer

Fig. 5. Country-view of the BGP Hijacking event on 24-Dec-2004 around 9:24 GMT. Turkey gains most of the Internet traffic while most of the rest of the world lose traffic. The monitoring point is the orange colored node (AS-3549) located in US.

Papadopoulos, Stavros, Konstantinos Moustakas, and Dimitrios Tzovaras. "BGPViewer: Using Graph representations to explore BGP routing changes." 2013 18th International Conference on Digital Signal Processing (DSP). IEEE, 2013.

BGP Update Visualizations

- Bigfoot

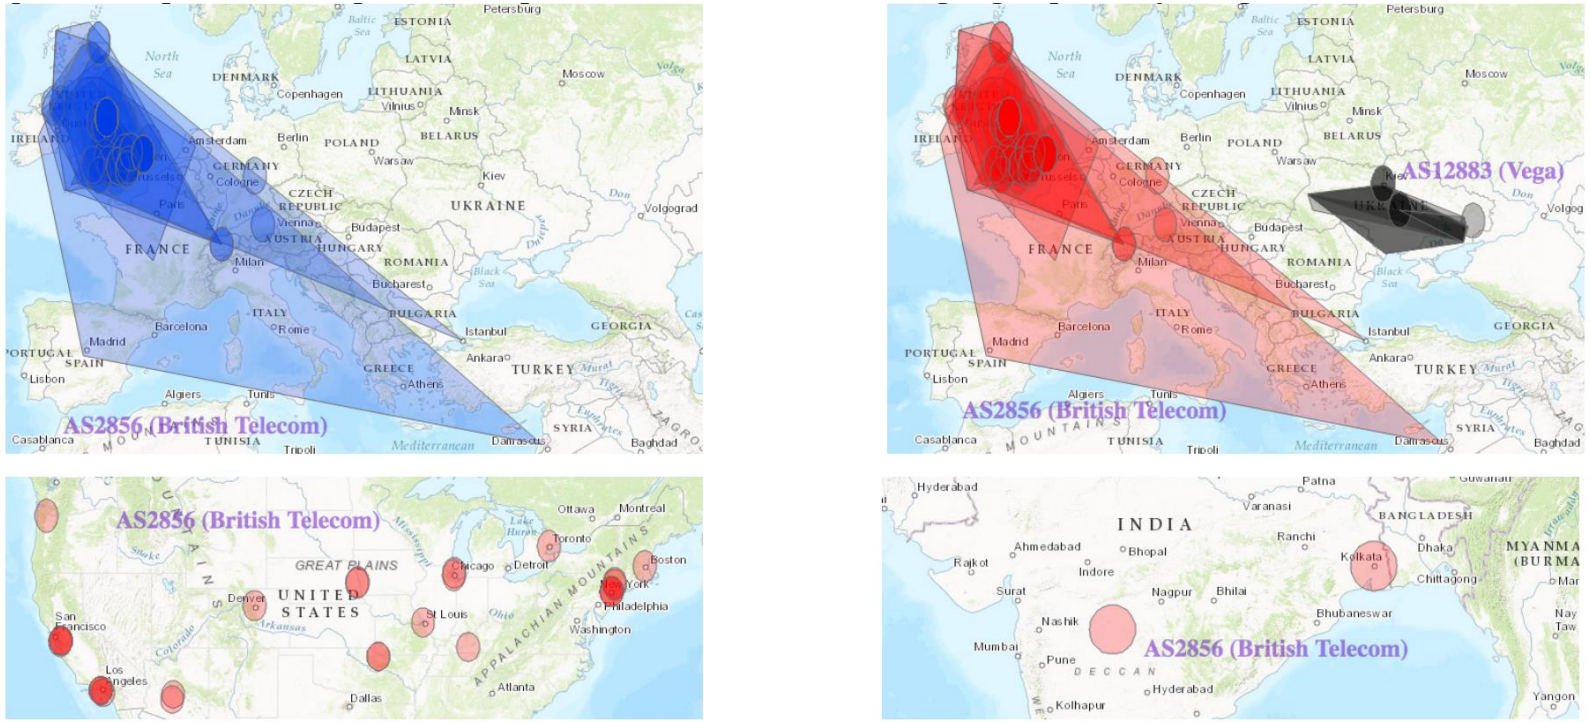

Figure 7: **Footprint of BGP update events produced before (top left) and during (top right) prefix hijacking (Ukraine) described in [25]. Customers from other continents including countries like India (bottom right) and USA (bottom left) were also attacked.**

Syamkumar, Meenakshi, Ramakrishnan Durairajan, and Paul Barford. "Bigfoot: A geo-based visualization methodology for detecting bgp threats." *2016 IEEE Symposium on Visualization for Cyber Security (VizSec)*. IEEE, 2016.

Geographical Visualization of BGP Update Paths

- Why geographical?
  - Path changes are easier to understand
  - Unusual detours are easier to spot

- How to do it more accurate than using bounding polygons or country level approximation?
  - GeoIP data for IP blocks with city precision
  - Highly dependent on GeoIP accuracy

- How do we visualize BGP update paths over multiple ASes?
  - Approximate intra and inter AS geographical route

Approximation of AS Data Center Locations

- Ground Truth
  - Only a few AS provide a network map of their data centers
  - No numerical ground truth dataset

https://www.cogentco.com/en/network/network-map

Approximation of AS Data Center Locations

From GeoIP Data to AS Data Center Locations

- Maxmind GeoIP2 City and ISP Databases

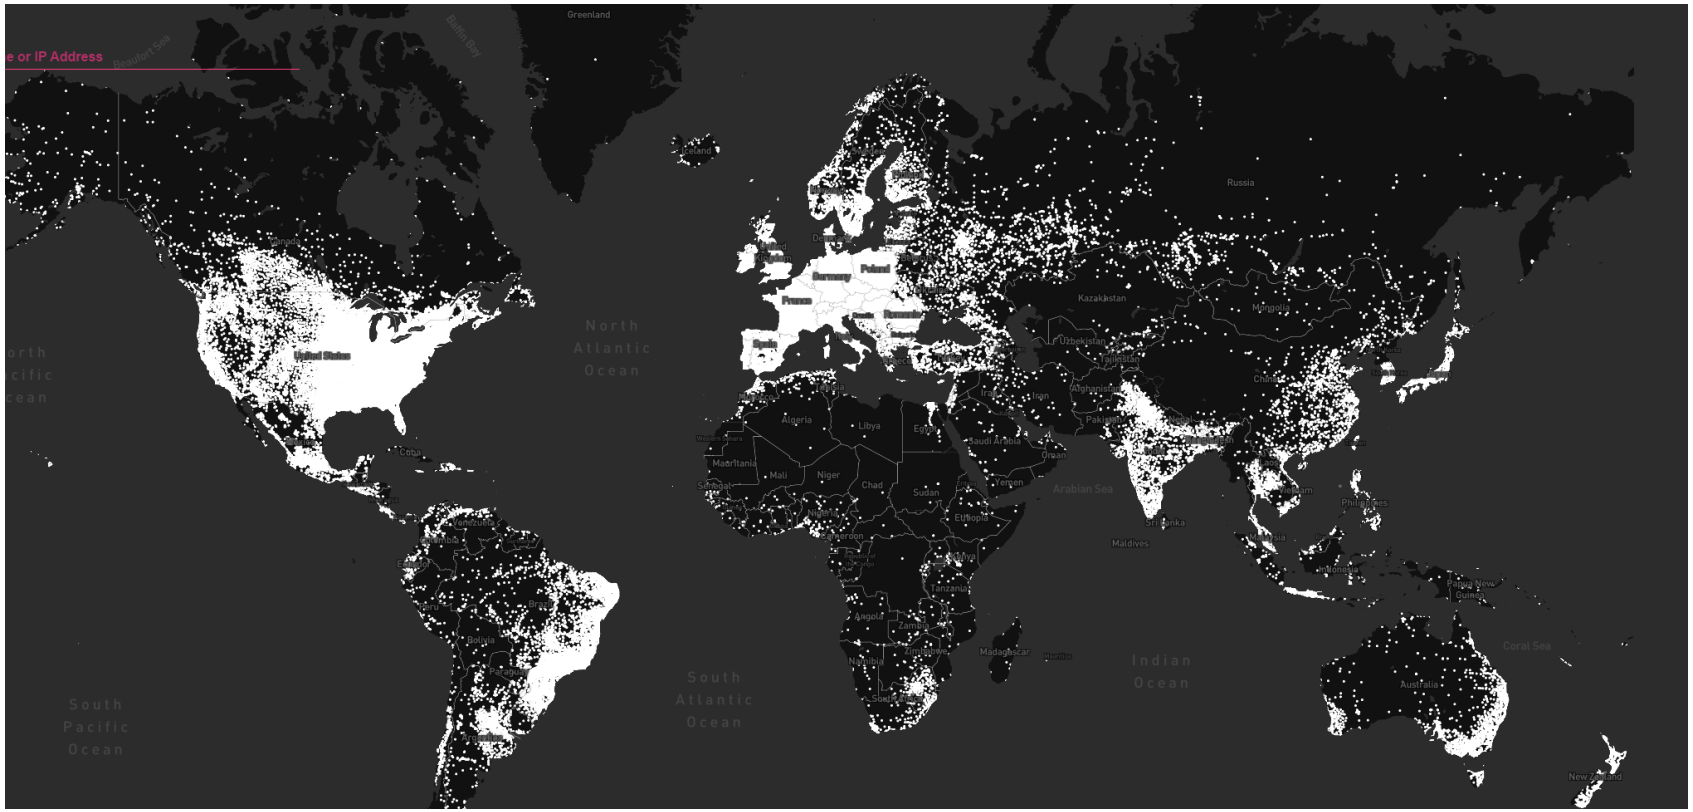

Fraunhofer IGD

Approximation of AS Data Center Locations

Clustering and Aggregation

- As we are doing an approximation, we need to make some assumptions.
- Some are good, some are weaker.
  Your feedback is very welcome, as our expertise in this domain is limited
- Assumptions:
  - (a) data centers are in proximity for most of their IP blocks
  - (b) two data centers of the same AS are not close to each other
  - (c) huge IP blocks indicate possible data center locations
  - (d) small IP blocks may qualify as data centers in sparse regions

Approximation of AS Data Center Locations

Clustering and Aggregation

1. Adaptive K-d-Tree over the whole map
2. On each K-d-cell an adaptive DBSCAN clustering
3. Post-processing to remove outlier and noise in the GeoIP data

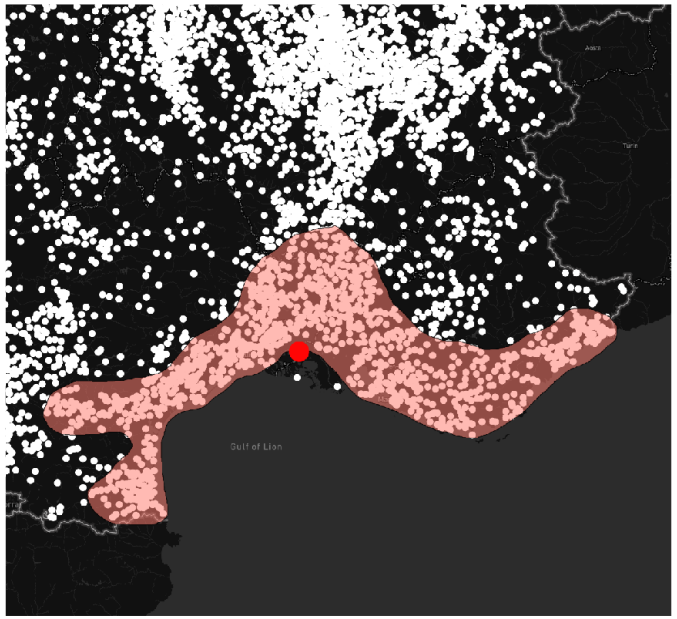

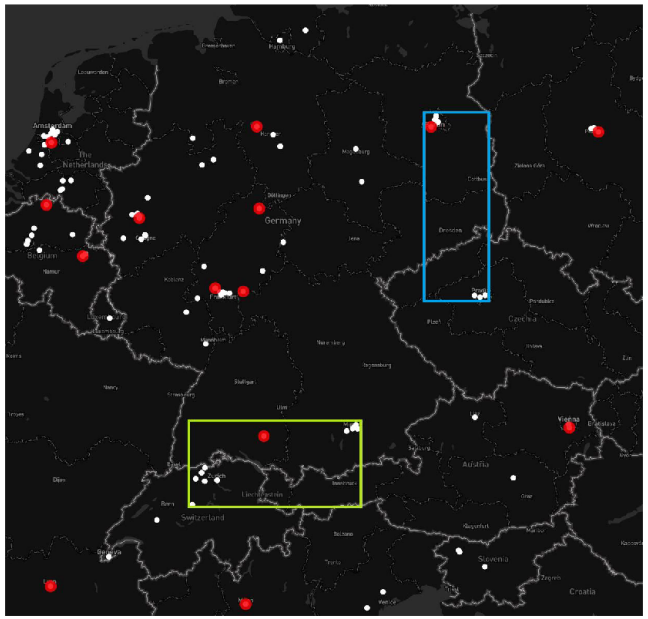

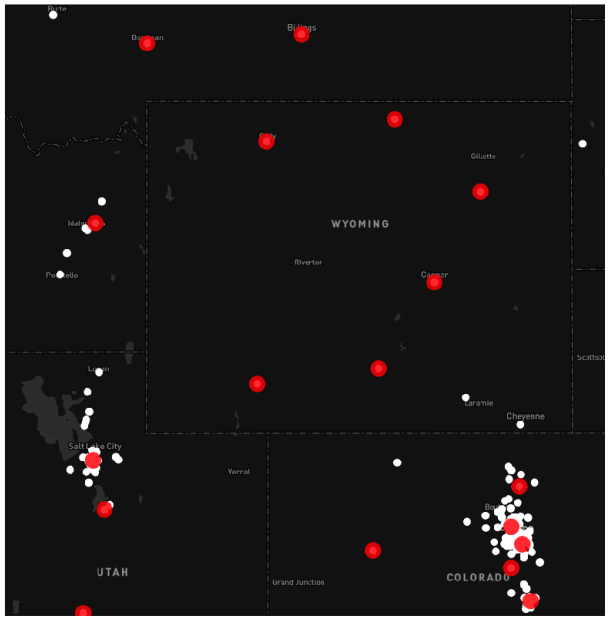

Approximation of AS Data Center Locations

Clustering and Aggregation

- Example: AS174 Cogent – IP block locations

Approximation of AS Data Center Locations

Clustering and Aggregation

- Example: AS174 Cogent – Data center approximation

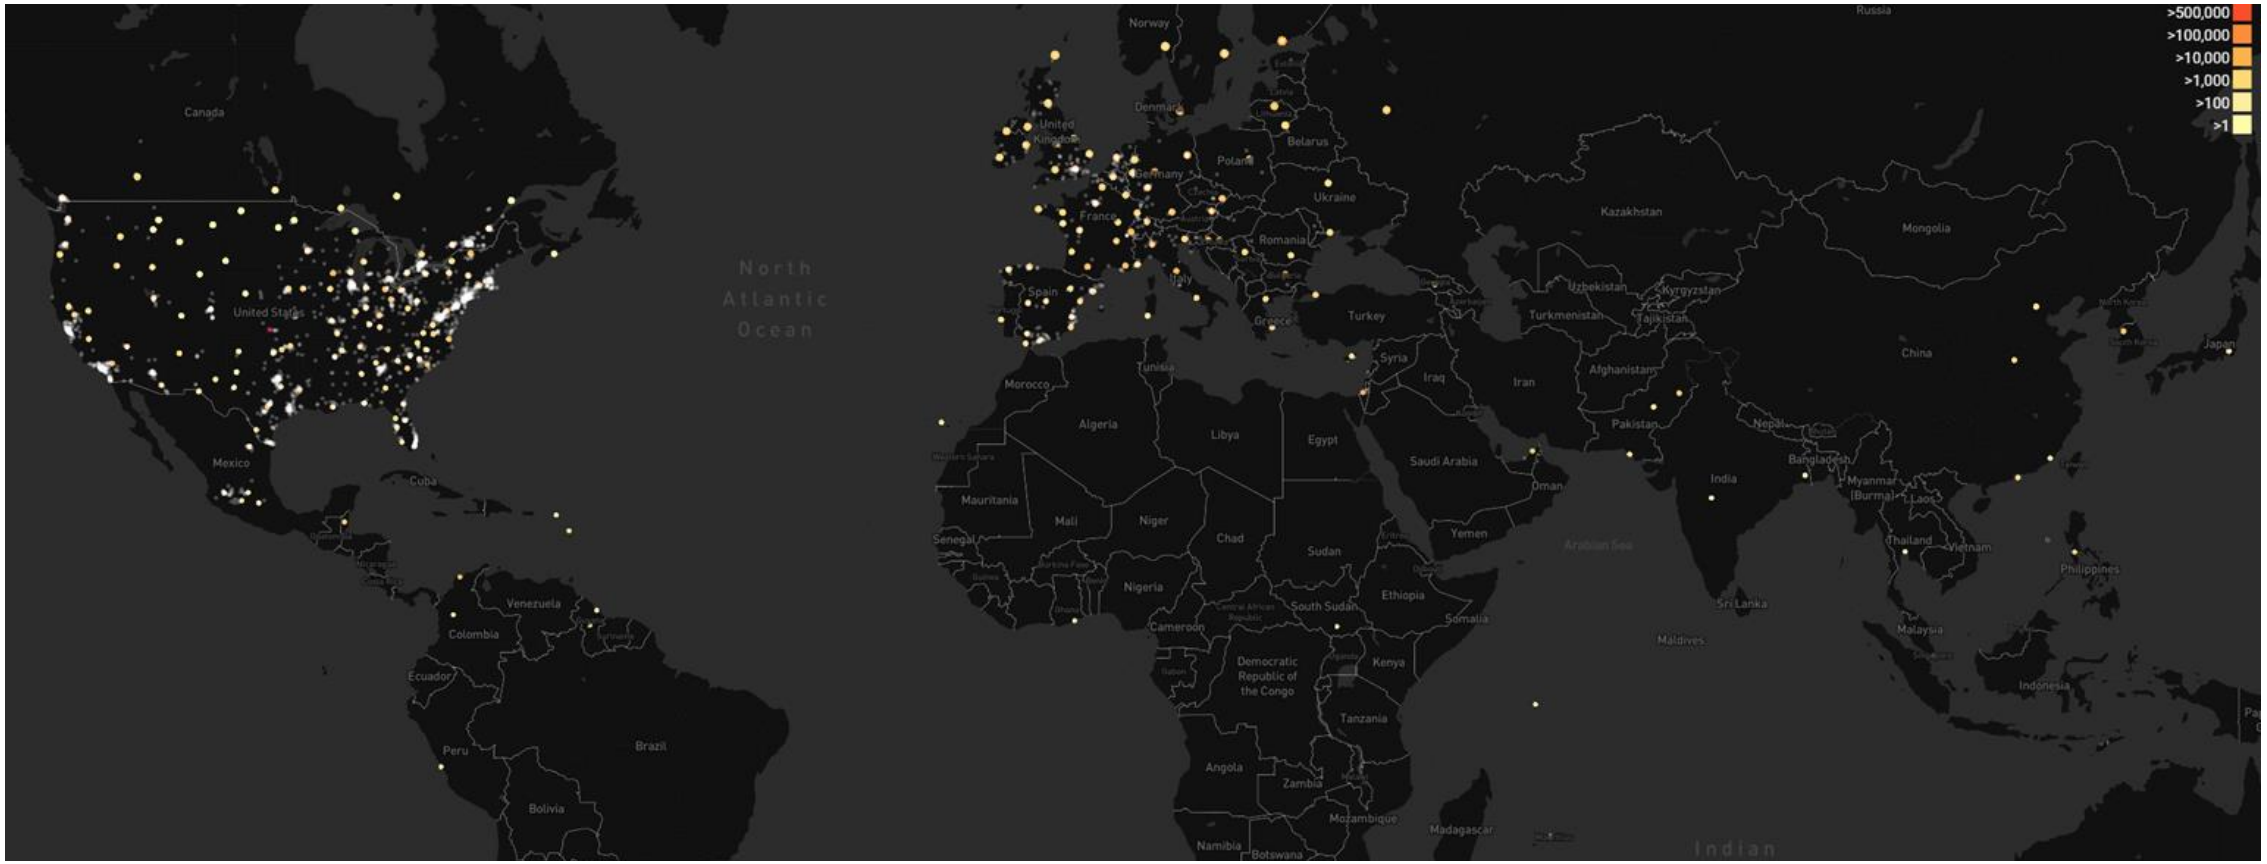

Geographical Visualization of BGP Update Paths

Internal Connections

- Assumptions:
  - (e) short connections are prioritized
  - (f) connections between big data centers are prioritized
  - (g) data centers should reach others with a low amount of hops

1. Kruskal's minimal spanning tree
2. Adding edges which reduce the distance between to data centers

- Your feedback is very welcome again

Approximation of AS Data Center Locations

Internal Connections

- Example: AS174 Cogent – Data center approximation

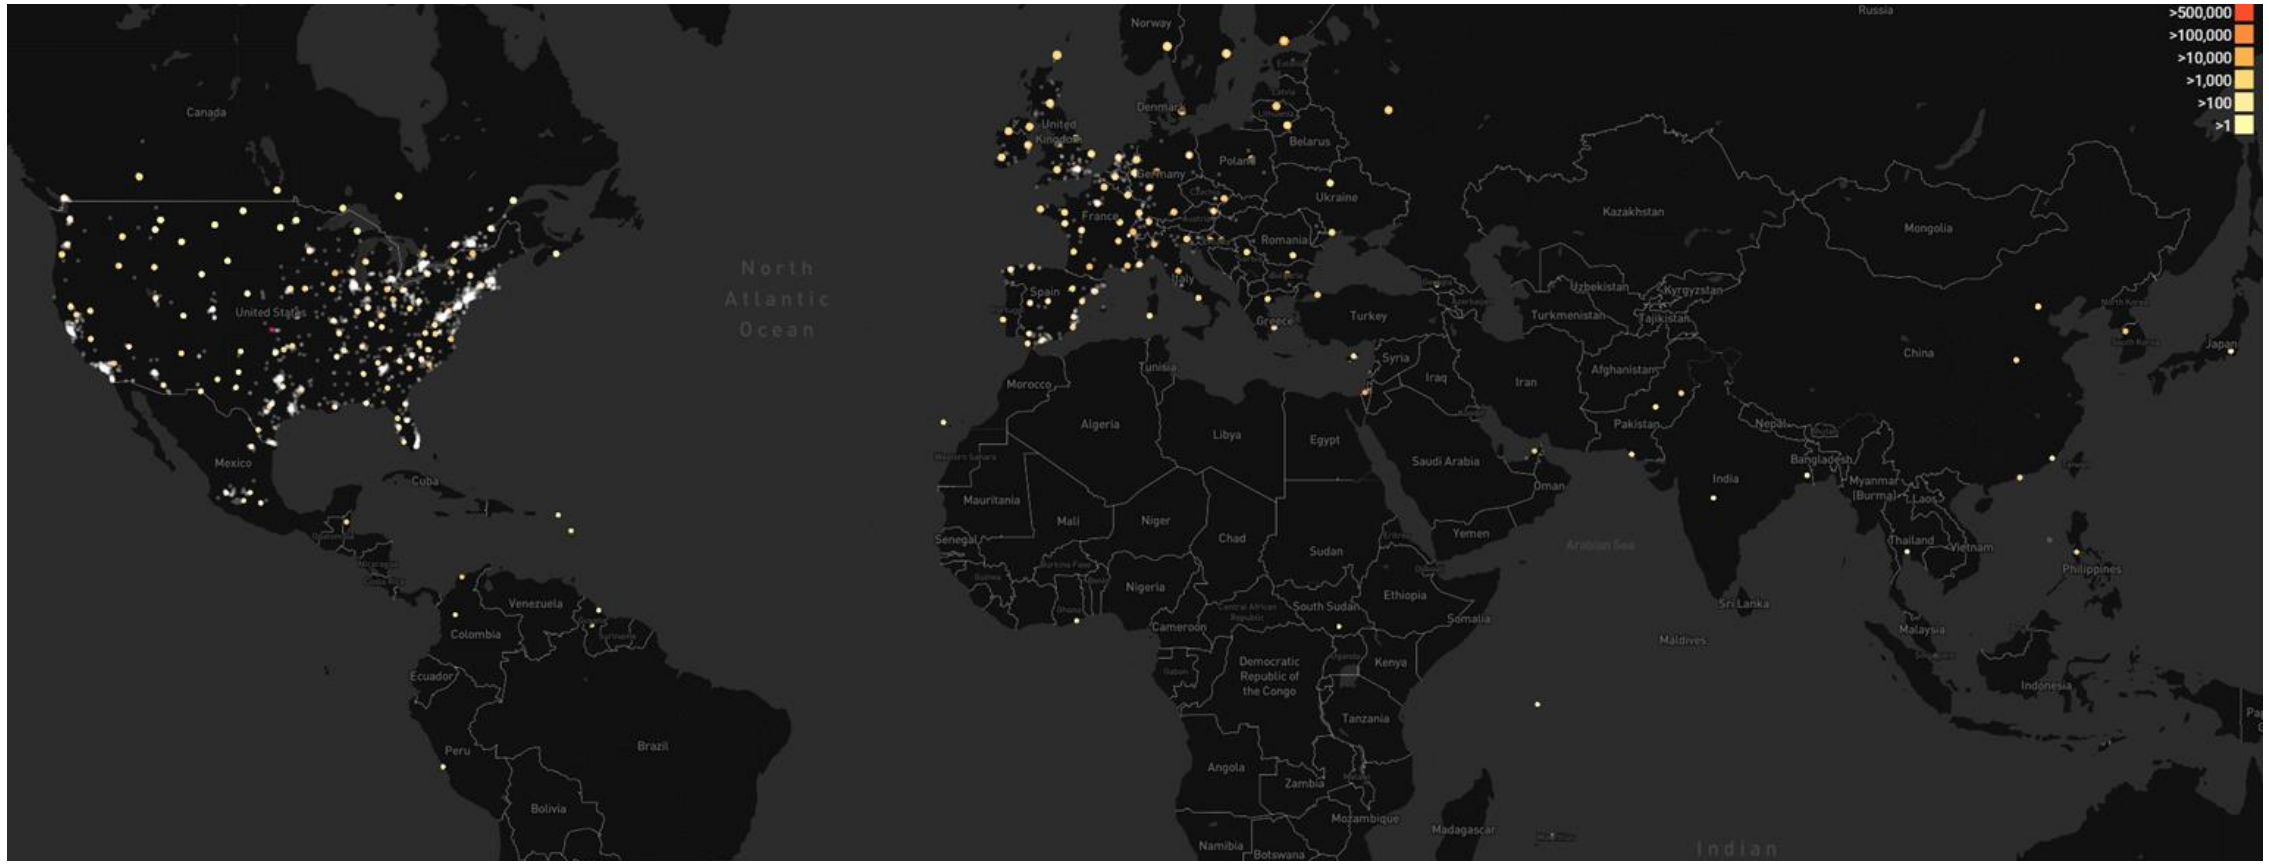

Approximation of AS Data Center Locations

Internal Connections

- Example: AS174 Cogent – Internal connection approximation

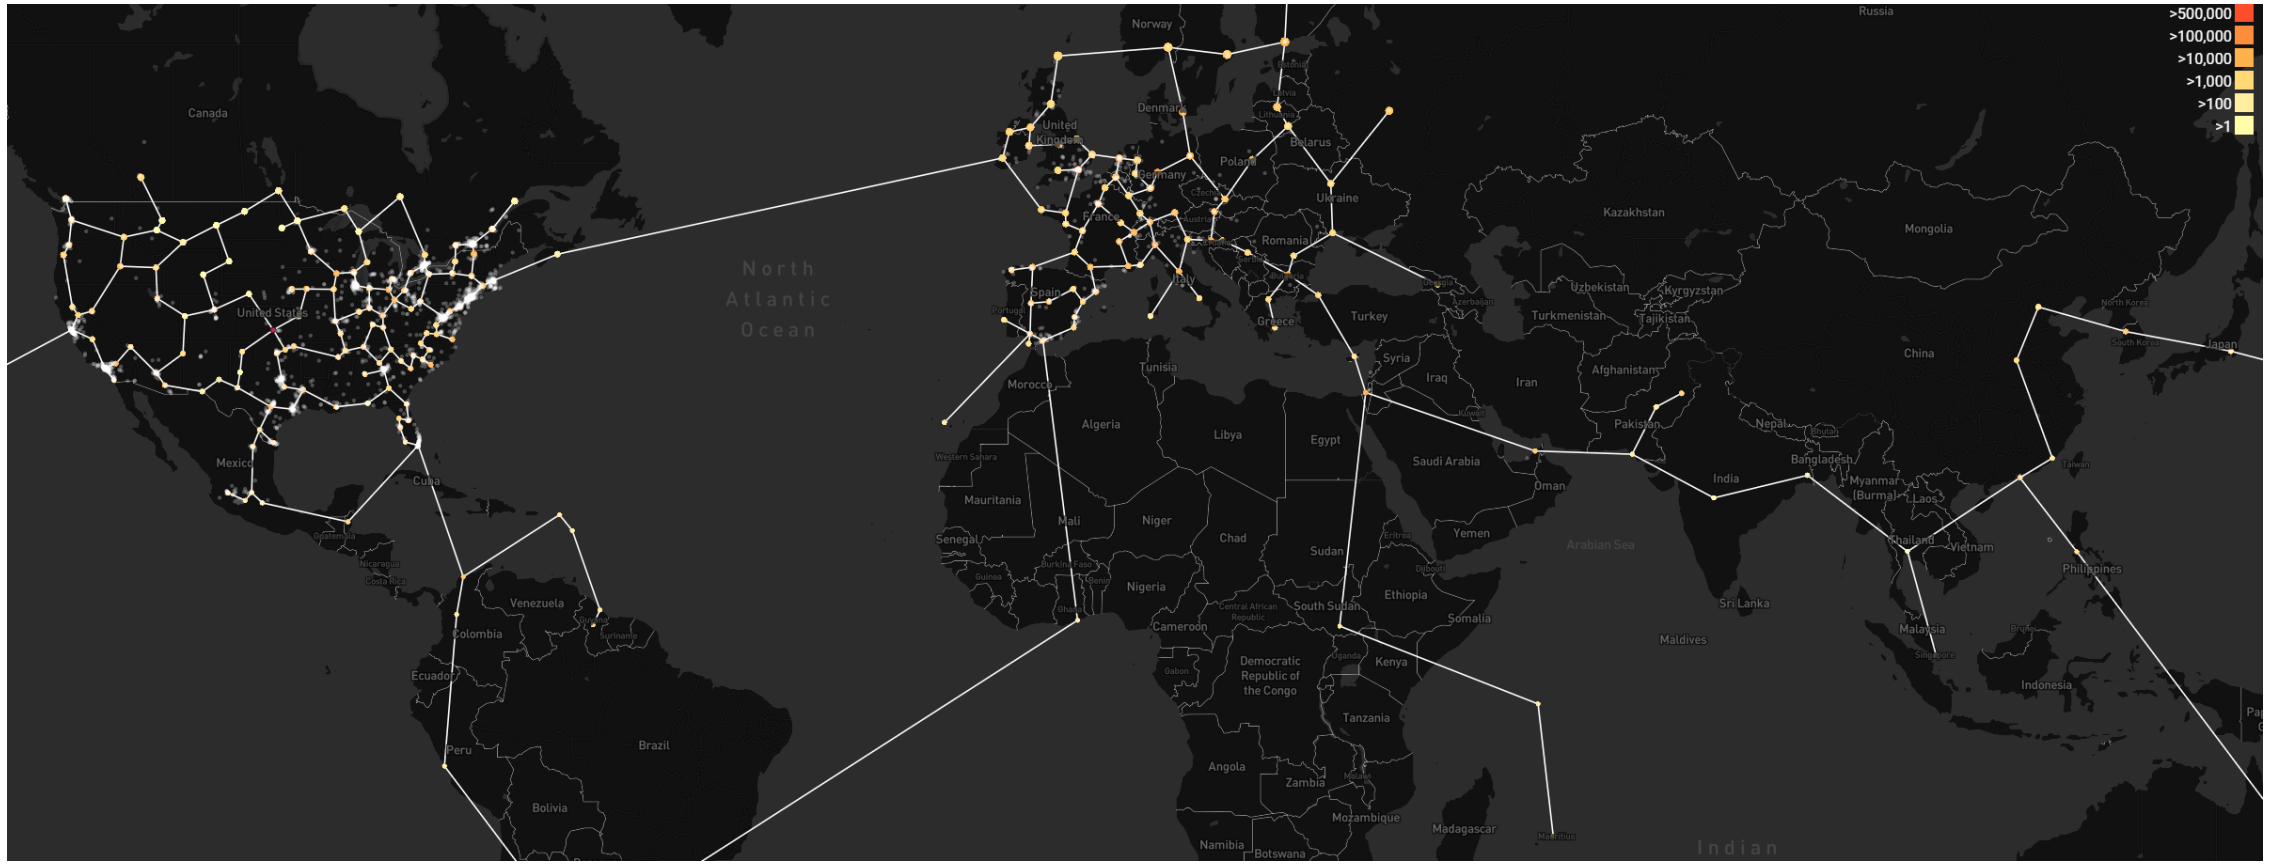

Approximation of AS Data Center Locations

Data Centers and Internal Connections

- Comparison to the Infographic from Cogent

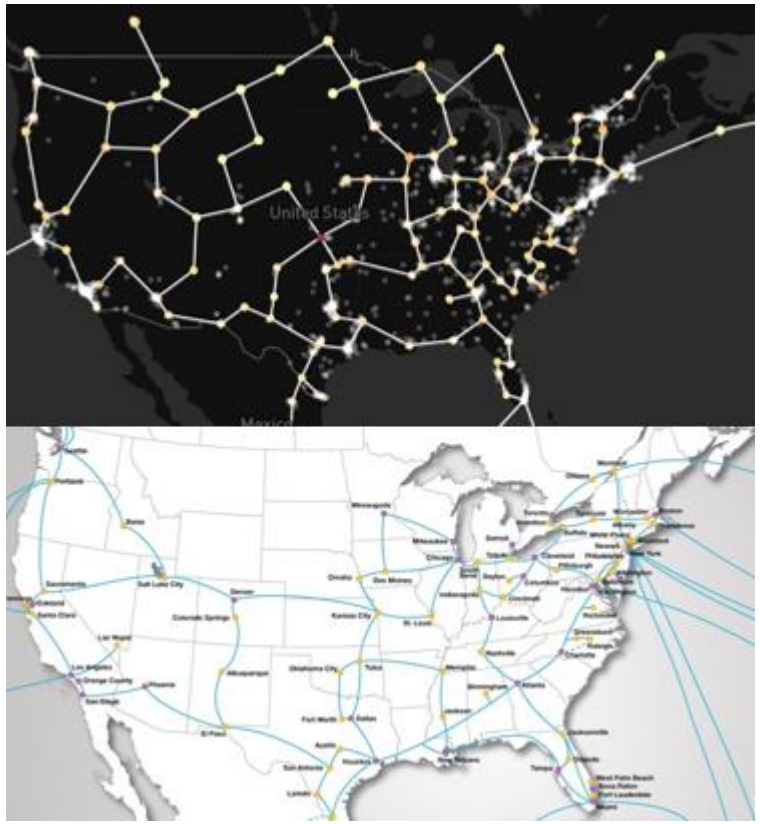

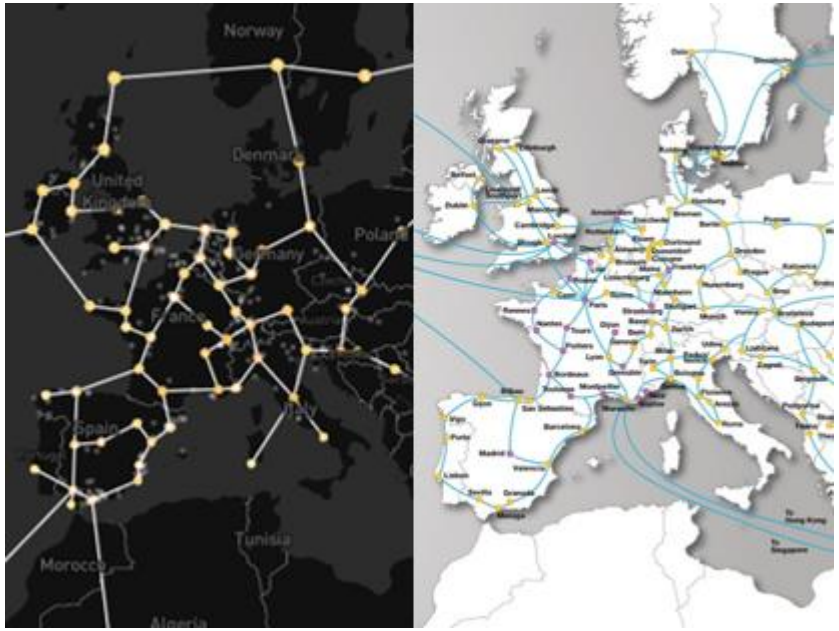

Geographical Visualization of BGP Update Paths

Update Path Visualization

Start at the next hop IP address location and create a path up to the origin AS

- Assumptions:
  - (h) prioritize long paths on own AS network
  - (i) neighboring ASes are locally connected
  - (j) prioritize routing in the direction of the target prefix
  - (k) prioritize shortest path

- Your feedback is very welcome again

Geographical Visualization of BGP Update Paths

Update Path Visualization

- Getting the BGP Update raw data:
  1. Get the archive download links by using CAIDA BGPStream Broker for the queried time frame
  2. Download the files and process them with ISOLARIO BGPScanner
  3. Filter the requested IP prefix and call the approximation algorithm for each AS
  4. Cache archives and approximation results for faster future calls

ProBGP Live Demo

https://probgp.igd.fraunhofer.de/