RIPE 88, Kraków
20-24 May 2024

University of York

Architecture and Routing in a Geopolitical World

Joshua Levett, Dr Vassilios Vassilakis, Dr Poonam Yadav

SYSTRON Lab

SEtS

UKRI Engineering and Physical Sciences Research Council

What does the Internet look like?

Approaches to visualisation

- AS core [1]
- Hosted content [2]
- Connectivity map [3]

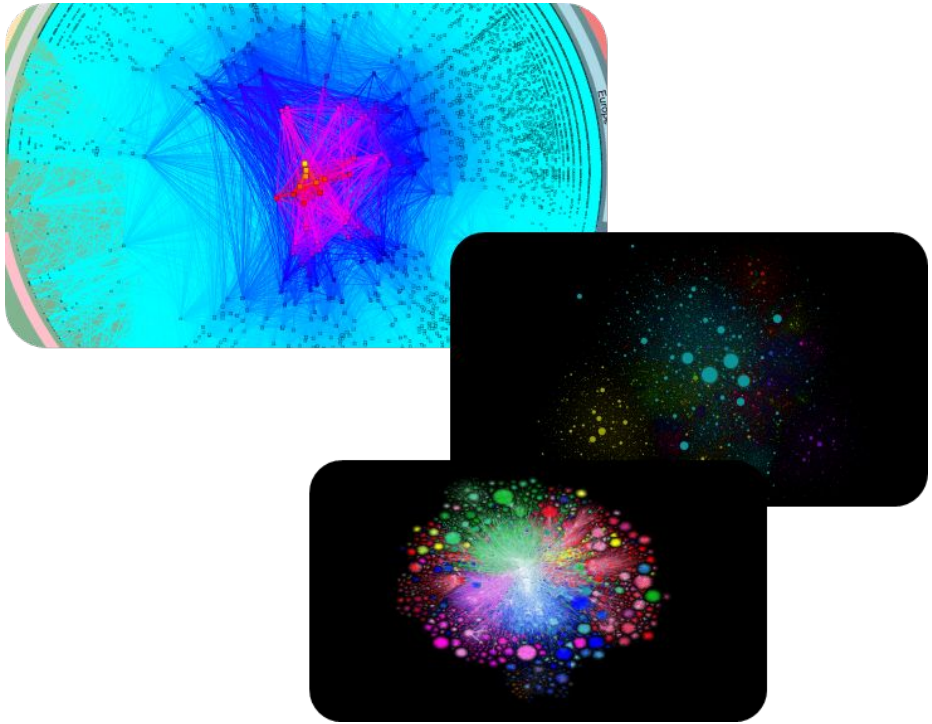

But why should we care?

[1] CAIDA, 2020, CAIDA's IPv4 and IPv6 AS Core.
[2] Ruslan Enikeev, Internet Map.
[3] Barrett Lyon, The Opte Project.

Objective: Utopian Internet

Predictable traffic movement

Low latency

Operational resiliency

Carrier transparency

Settlement-free peering

…a border-free, settlement-free, transparent, uncensored Internet?

Context: Local & Global Internet Governance

Internet Sovereignty

"the right of a state to govern its networks to serve national interests" [1]

NETmundail+10

"[Working] to create the networked global governance architecture that is human-centric, inclusive, environmentally friendly, and development-oriented, as the networked society demands." [2]

[1] James Lewis, Sovereignty and the Evolution of Internet Ideology.
[1] Preliminary NETmundail+10 Outcome Document, 2024.

Context: Regulatory Intervention

- Content
- Platform
- Network

What is the Impact?

Does the topology differ between countries?

How does the topology differ between countries?

Why does the topology differ between countries?

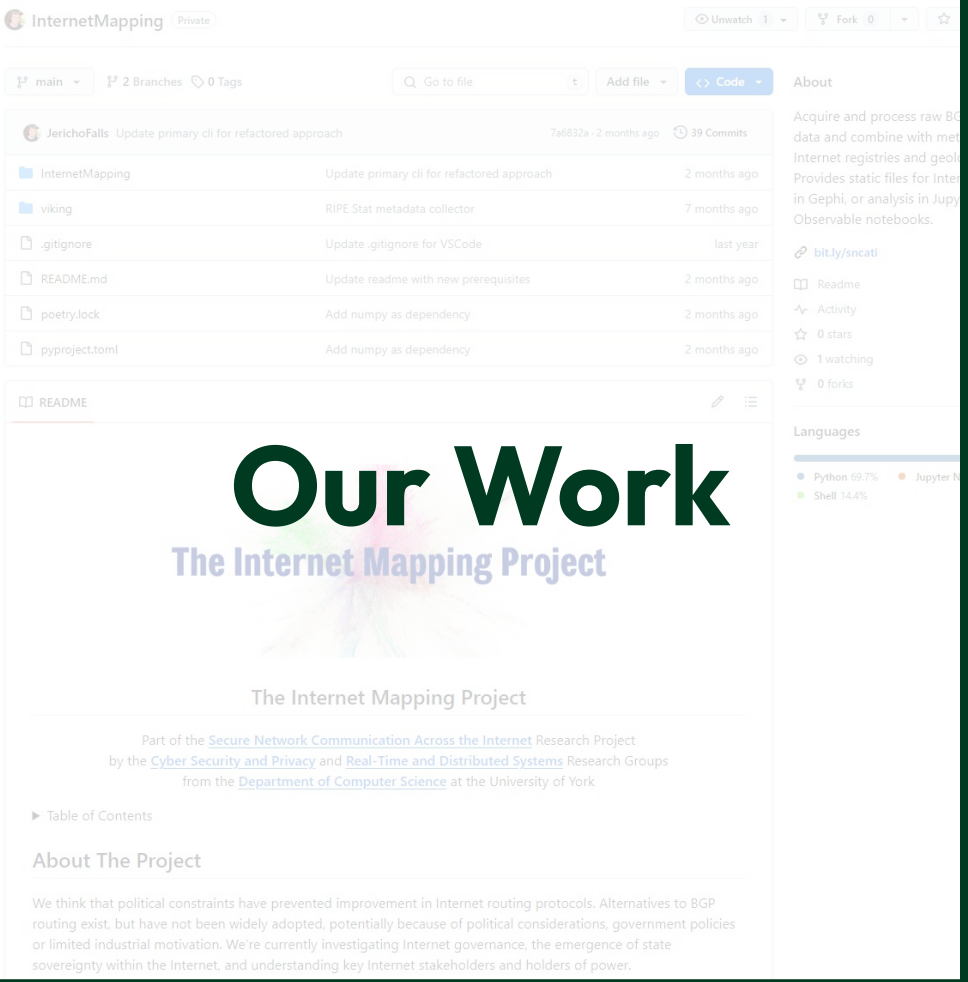

A new tool to create metadata-rich Internet topology graphs at higher completeness

systronlab.github.io/projects/internet-mapping

Releasing source code next week.

1.
Our Approach

Observing the Internet

+

+

System Architecture

- **Route Collectors**
  RIPE RIS, RouteViews, PCH
- **Metadata and Probes**
  RIPEstat, RIPE Atlas, GeoLite2
- **Data Processing**
- **Graph output**

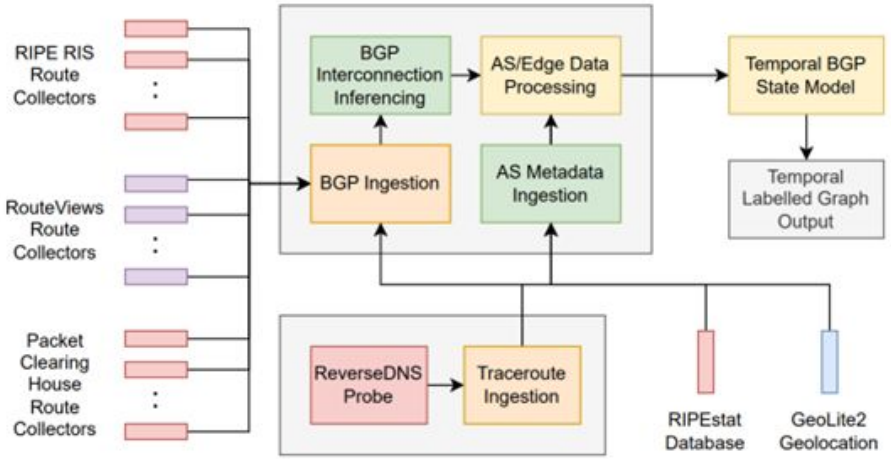

Capturing Location

- **Internet Registry**
  RIPEstat, pulling data from ARIN, LACNIC, RIPE NCC, AFRINIC, APNIC
- **Prefix Geolocation**
  MaxMind GeoLite2 City

MaxMind also uses some WHOIS data

Adding Metadata

- **Registered Owner**
  Processed into a usable format and sibling ASes (same owner) detected (using orgName, orgId, opaqueId, PeeringDB)
- **Registered Location (country)**
- **State Ownership Data**
  Using state_owned_as dataset to identify majority state-owned ASes

Output

- **Node**
  AS with metadata (owner, registered location, geolocation, etc)
- **Edge**
  Adjacency between ASes

2.
The Internet in 2024

Our Snapshot: 1 May 2024

Our Topology

- 84,266 ASes
  (72% of NRO assignments)

Comparisons

- 116,977 NRO assignments
- 84,042 ASes seen in RIS

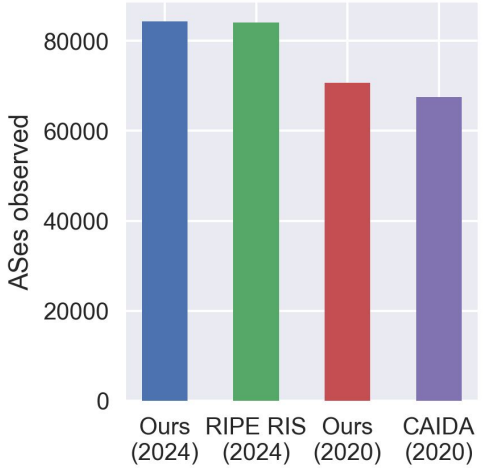

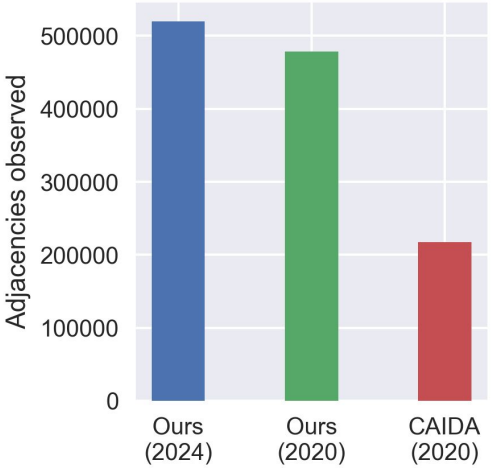

Maximum Degree

CCDF plot of the maximum degree of an AS combined with its probability (log-log scale)

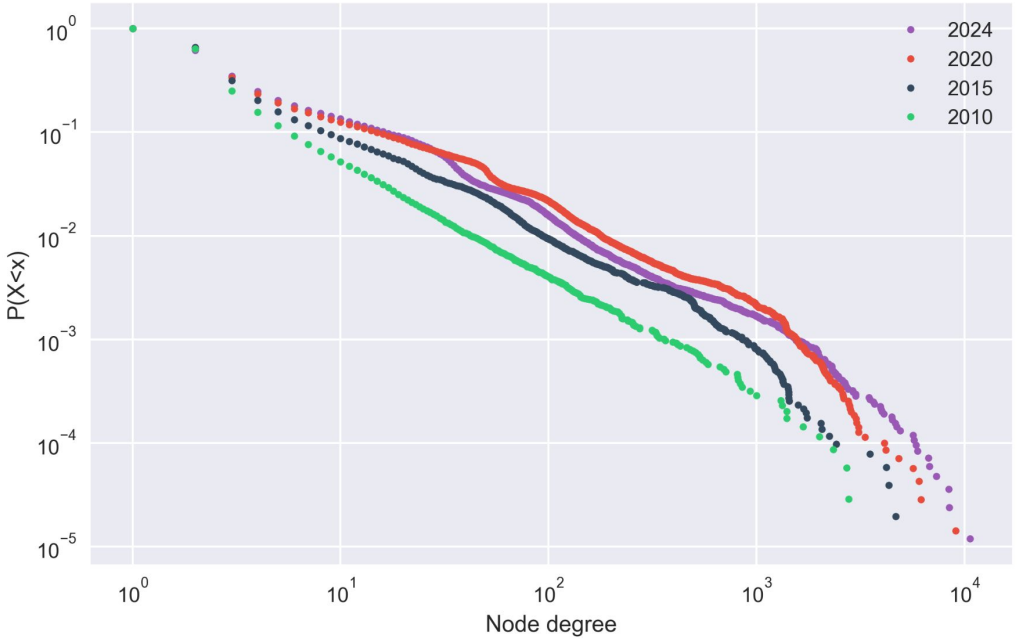

Maximum Degree

Decrease in Medium-Degree
Between 2023 and 2024, the volume of nodes in the mid-degree range reduces.

Increase in High-Degree
The number of the highest degree nodes grows.

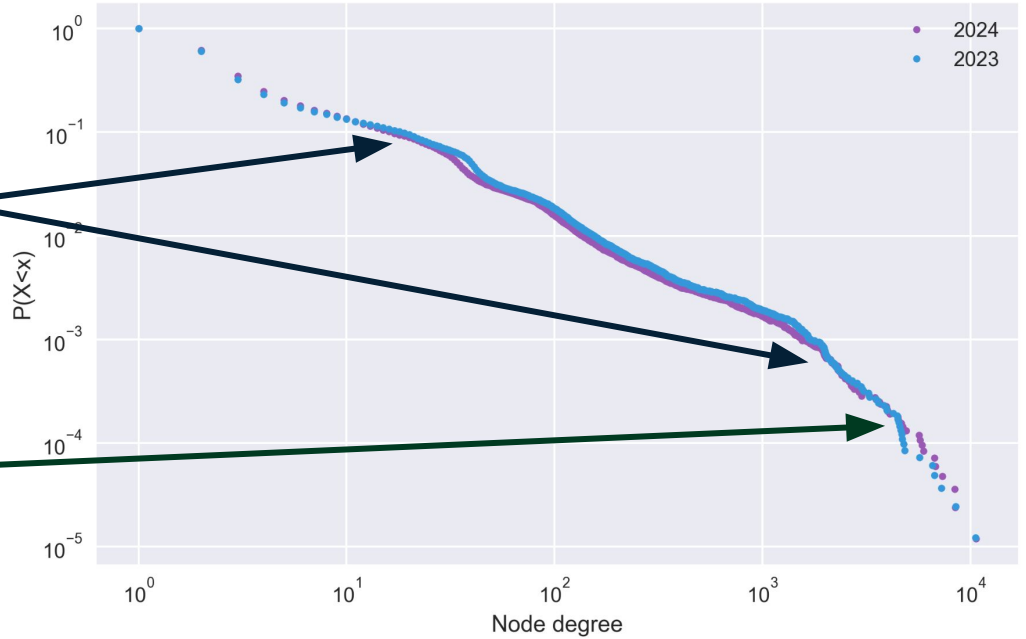

Maximum Degree & RPKI

CCDF plot of ASes:

- with an RPKI deployment (at least one ROA object)
- without an RPKI deployment

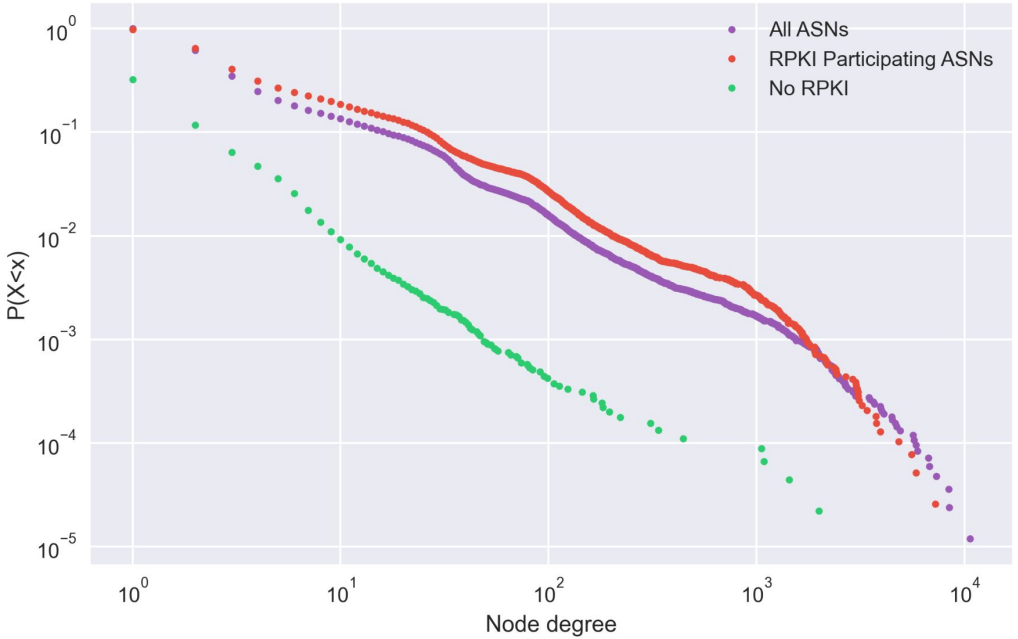

Average Neighbour Degree

CCDF plot of the average degree in a node's neighbourhood (e.g. the average degree of all nodes directly connected to a node)

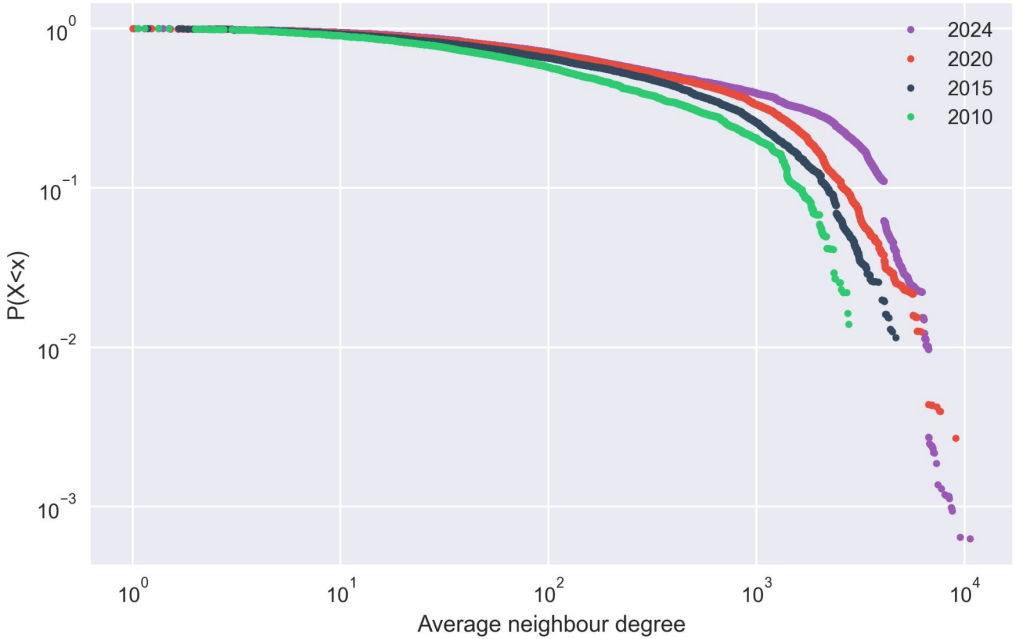

Most Influential*

Ranking the most eigencentral ASes within the topology.

- **How well-connected is a node, also considering its field of neighbours?**

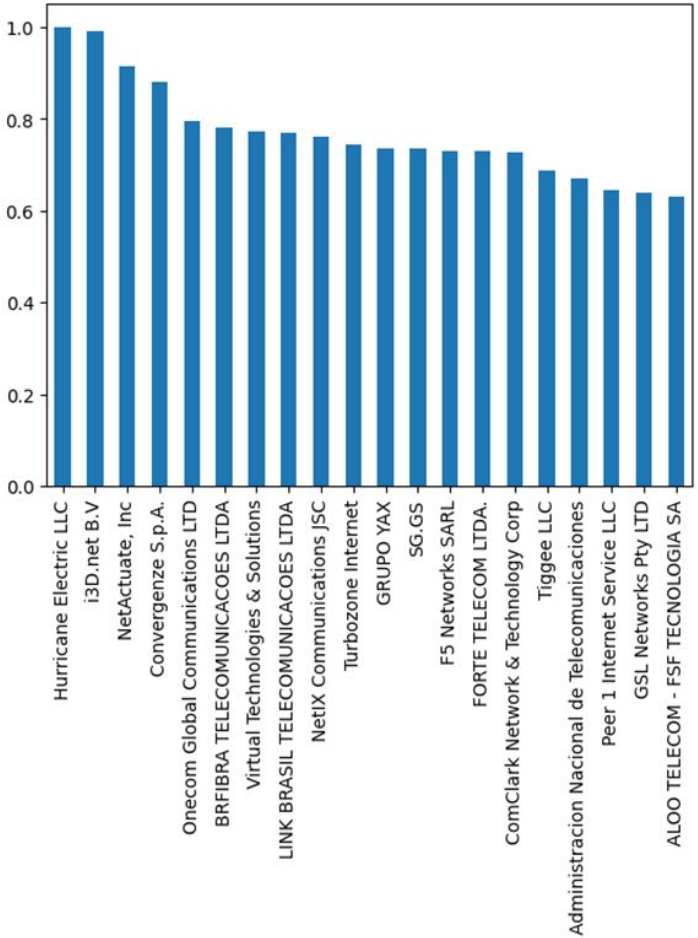

*Some vantage-point based bias exists, so results are only indicative rather than conclusive.

ASes and Peerings

The number of ASes (registered or observed) has increased by 411%

The number of (public) interconnections has increased 1,119%

Average path length has **increased from 3.7 to 12**.

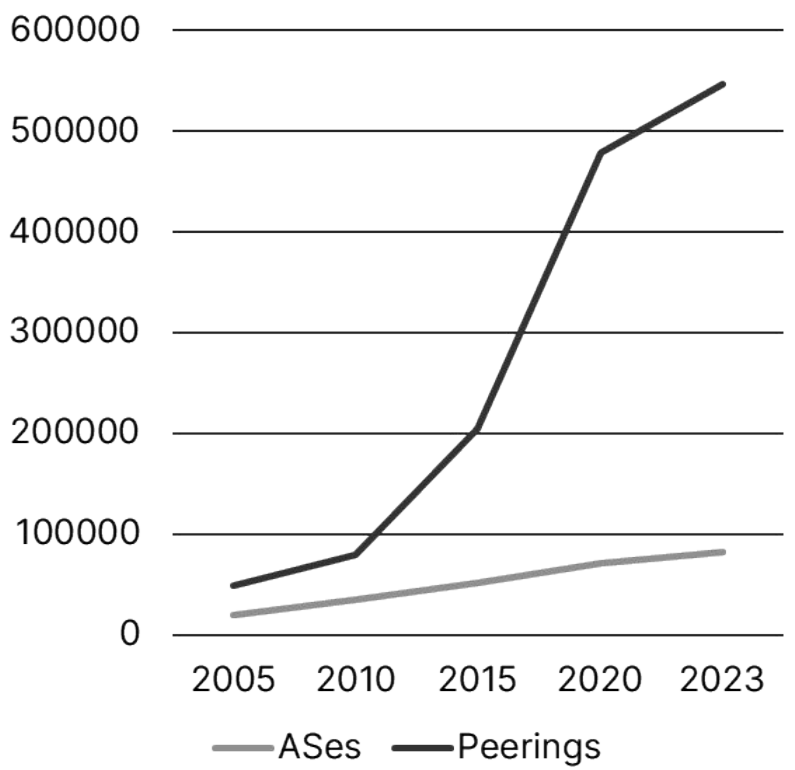

3.
The Does

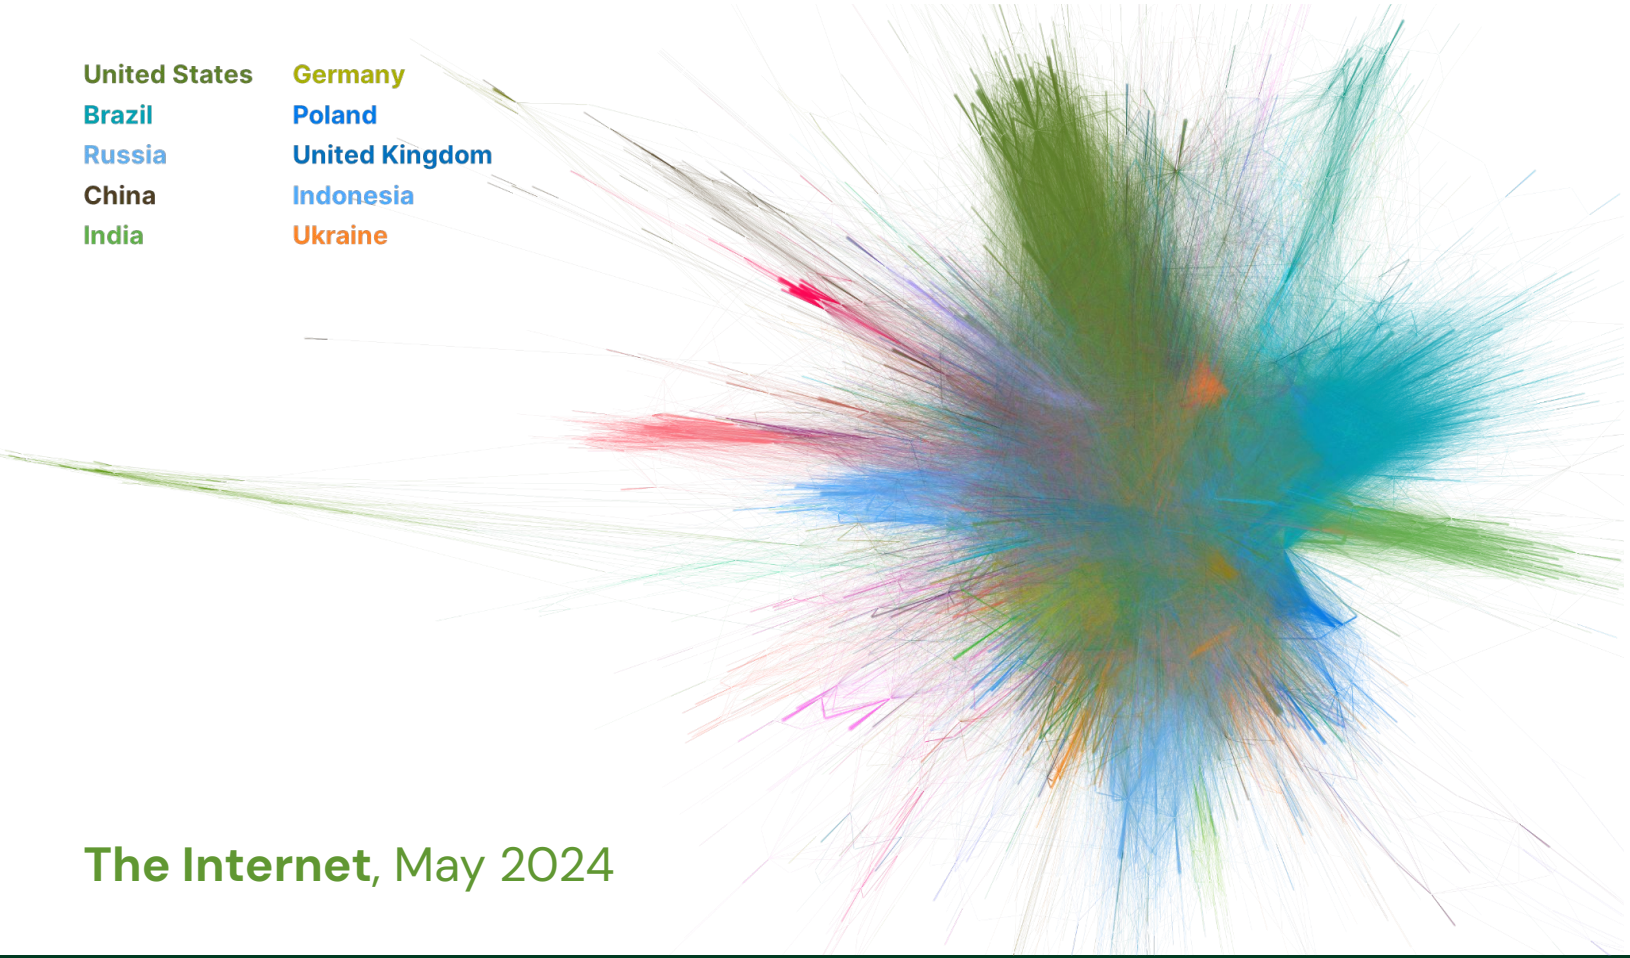

The Internet, May 2024

4.
The How

Foreign Neighbours

Unique Upstream Neighbours

Where X is a country of interest, N the Internet's ASes, and K the ASes within X:

$$K = \{k \in N \mid country(k) = X\}$$

Where I is the neighbours:

$$I = \{i \in N \mid \exists k \in K : i \in neighbours(k)\}$$

Where I_0 is the set of first-degree neighbours of country X, the sets $I_1 \rightarrow I_d$ where $d \in \mathbb{N}$:

$$I_1 = \{f \in N \mid \exists i \in I_0 : f \in neighbours(i)\} \setminus I_0$$

Or more generally:

$$I_d = \{f \in N \mid \exists i \in I_{d-1} : f \in neighbours(i)\} \setminus (I_{d-1} \cup \ldots \cup I_{d-q}) : q \in \mathbb{N} \wedge d > 0 \wedge d - q \geq 0$$

Unique Upstream Neighbours

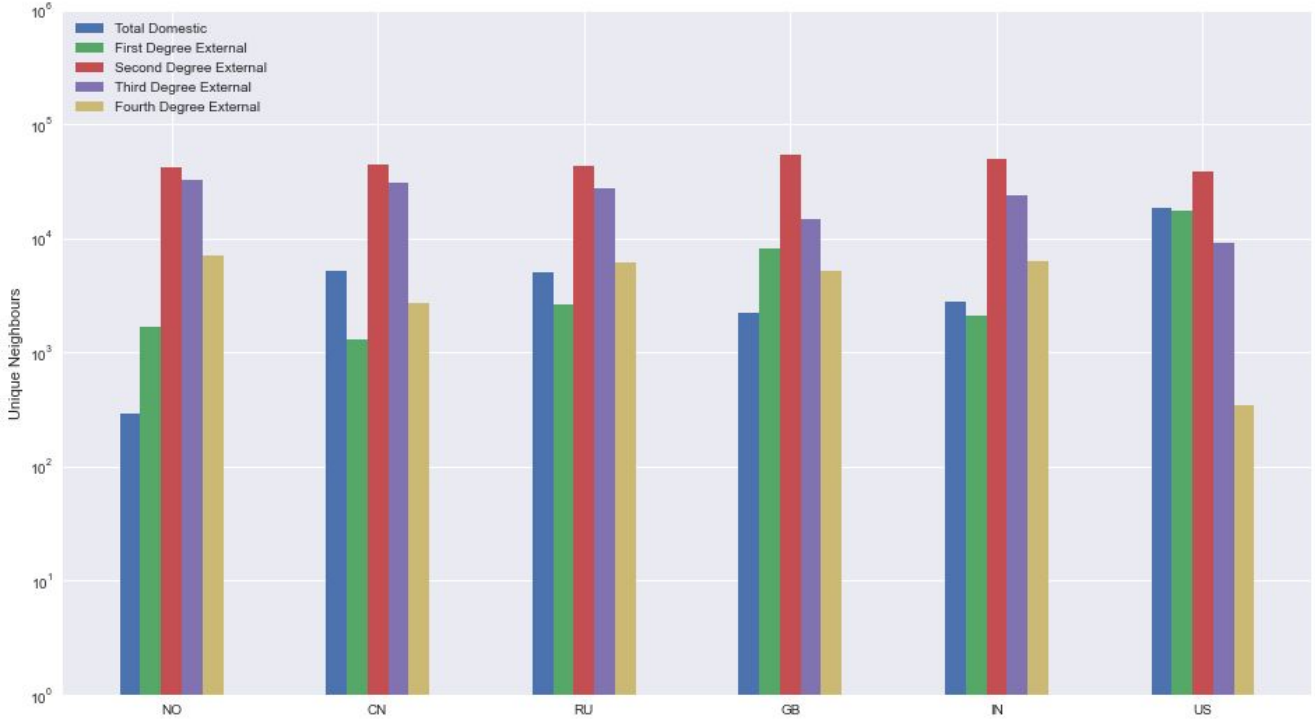

Visualising Connectivity

Using **Kamada & Kawai (1987)**

Greater distance
Lesser importance
(likely to be T2/3)

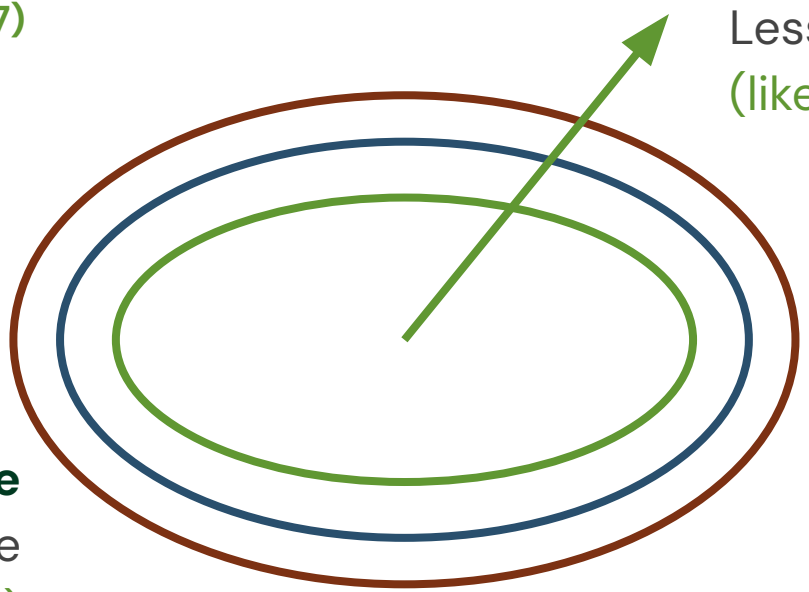

Closer to the centre
Greater importance
(more likely to be T1/2)

Iran

- 527 domestic ASes
- 57 foreign neighbours

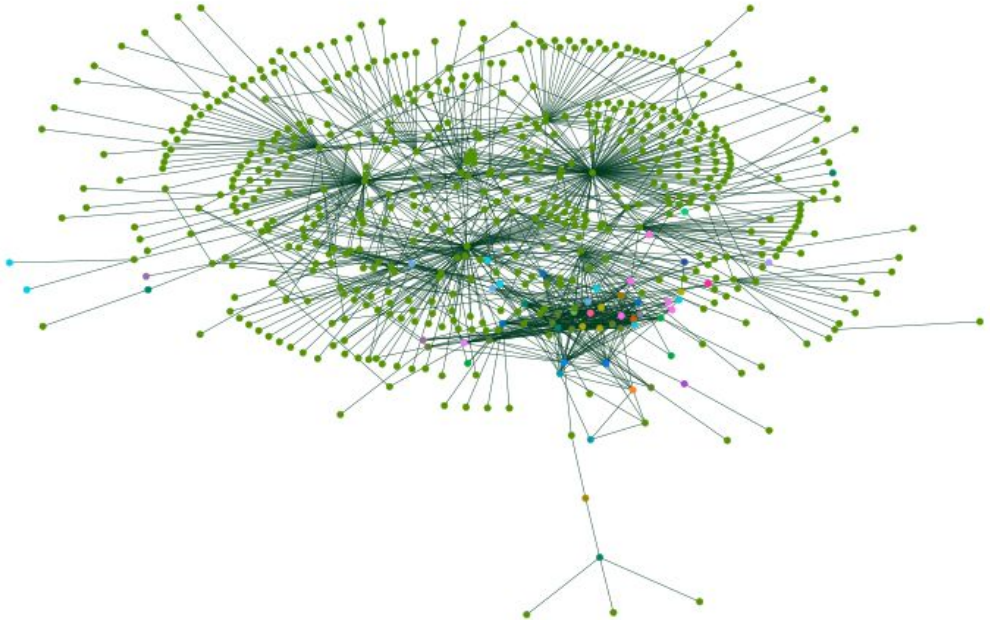

Norway

- 282 domestic ASes
- 489 foreign neighbours

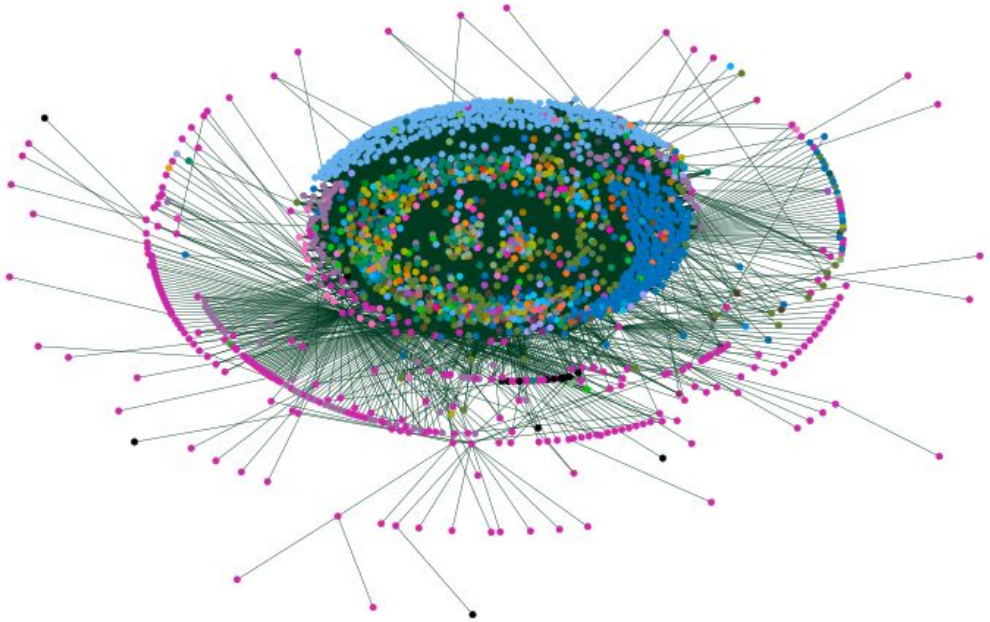

United Kingdom

- 2,115 domestic ASes
- 3,572 foreign neighbours

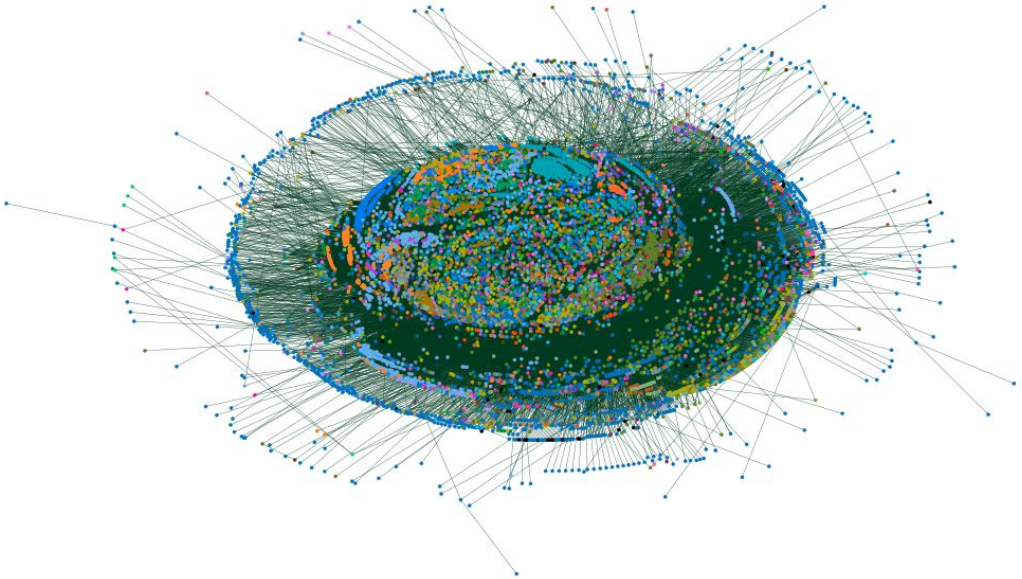

5.
Censorship: a Why?

Detecting Censorship

Explicit
Users are aware of censorship taking place.

- **Blockpages**
- **Incorrect DNS resolution**

Non-Explicit
Users are possibly not aware of censorship taking place.

- **HTTP failures**
- **DNS, TCP, HTTP anomalies**

Mitigation
Also considering the ratio of ‘normal’ traffic.

Relative Upstream Change

Capturing the change in the number of ASes at each step from the ‘border’.

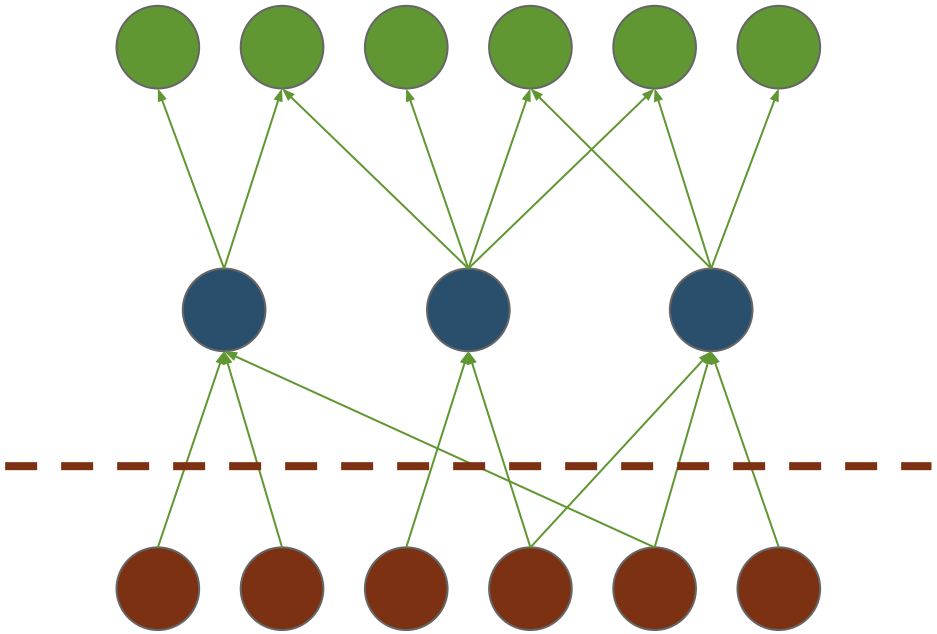

Relative Upstream Change

A relationship starts to become apparent as the degree increases...

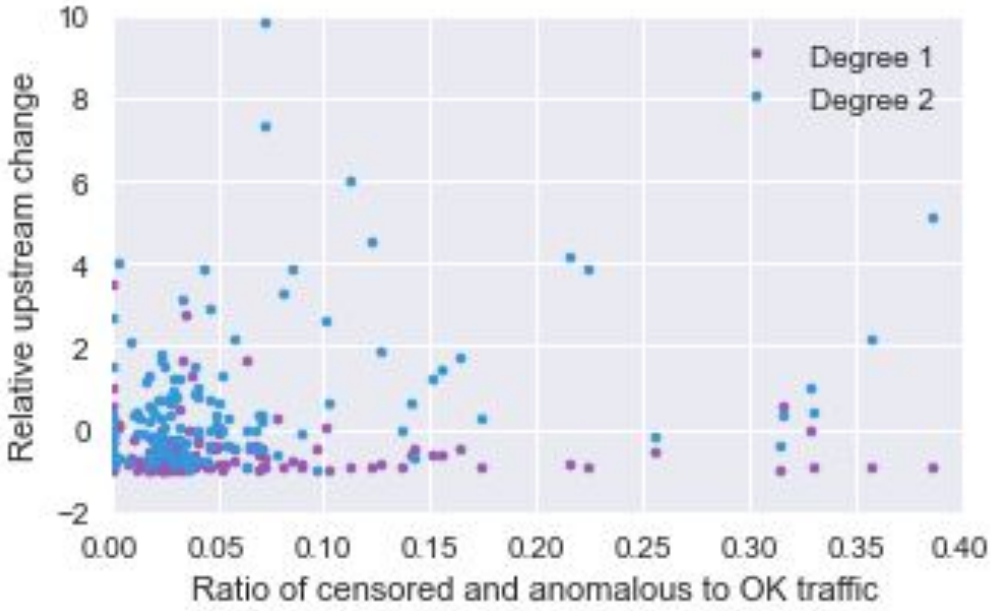

Cumulative Downstream Burden

Starting to see some connection (~0.5 correlation coefficient) between the downstream burden and censored traffic.

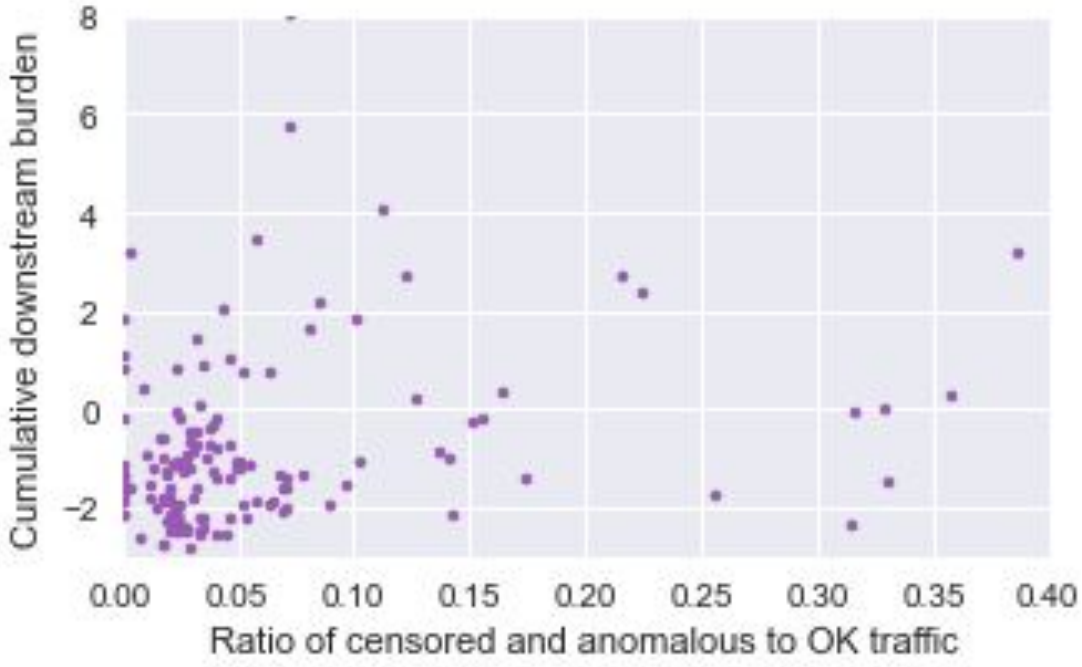

Funnelling

Based on the cumulative downstream burden, but weighted for network size.

Generally, higher censorship gives higher funnelling.

...with some exceptions (but remaining geopolitical)

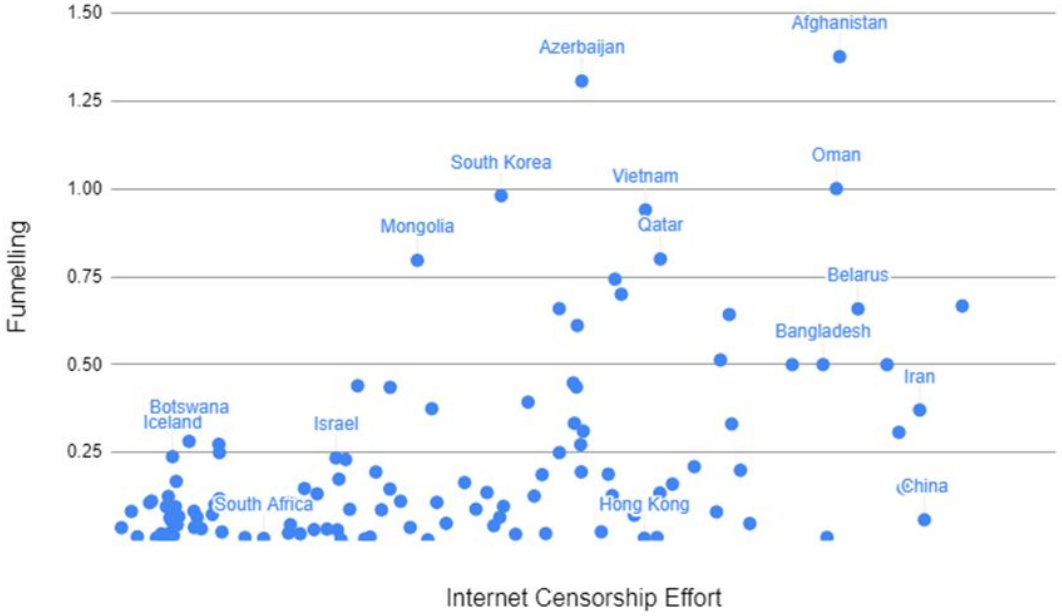

Summary

Internet topology fusion tool

The volume of highest-degree ASes is growing…
…and the volume of lowest-degree is also growing

Geopolitics has a statistically significant impact on topology structure

Questions & Contact

SYSTRON Lab
systronlab.github.io

Joshua Levett
joshua.levett@york.ac.uk

Dr Poonam Yadav
poonam.yadav@york.ac.uk

SEtS

UKRI Engineering and Physical Sciences Research Council