Mapping the Internet by graphing BGP

The Art of Pretty Pictures

Dean Pemberton – NSRC

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Who am I

- Dean Pemberton
  - **NSRC**
    - Trainer/Network Engineer
  - **Victoria University of Wellington**
    - Adjunct Researcher
  - **InternetNZ**
    - Technical Policy Advisor
  - **APNIC**

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Why

- Why would we want to map the Internet?
- What is it about maps that makes them so awesome?

What do maps do?

- Show us things
- Hide things from us

Show us things

- Maps can show us views which we would never be able to see with our own eyes.
- Only a small number of people have ever seen the Earth from space, but a much larger number have viewed a world map.
- You can't even see the whole surface of the earth from space but… you can on a map.

Show us things

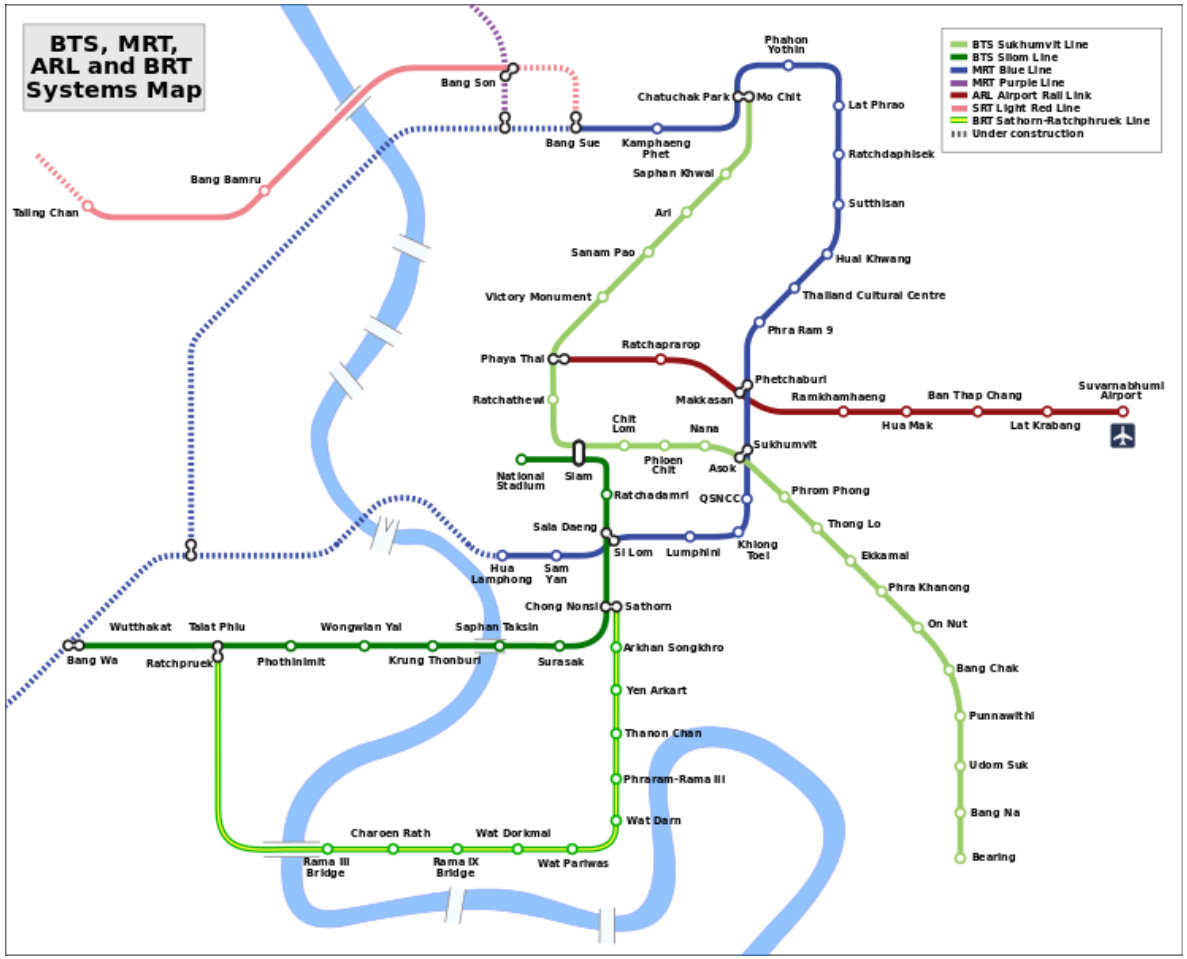

Hide things from us

Hide things from us

- Sometimes there is just too much complexity and information.
- We need to strip some out so that we can make sense of other areas.
- No use in mapping every laptop and smartphone for example

Validating what we currently know

- We all have some assumptions about things.
- Maps can help us to validate these.
  - "Oh that station is just next to this other one"
  - "Oh yes – that AS has lots of peers"

Finding out new things

- Once we've learnt how to read a map and validate the things we already know, we can use it to find similar but previously unknown structures.
  - "Oh I had no idea that those stations were within walking distance"
  - "Oh I had no idea that o3b was one of the major upstreams for the Pacific."

Pretty Pictures

- At the end of the day I think they look pretty cool =)

Asia/Pacific (APNIC)

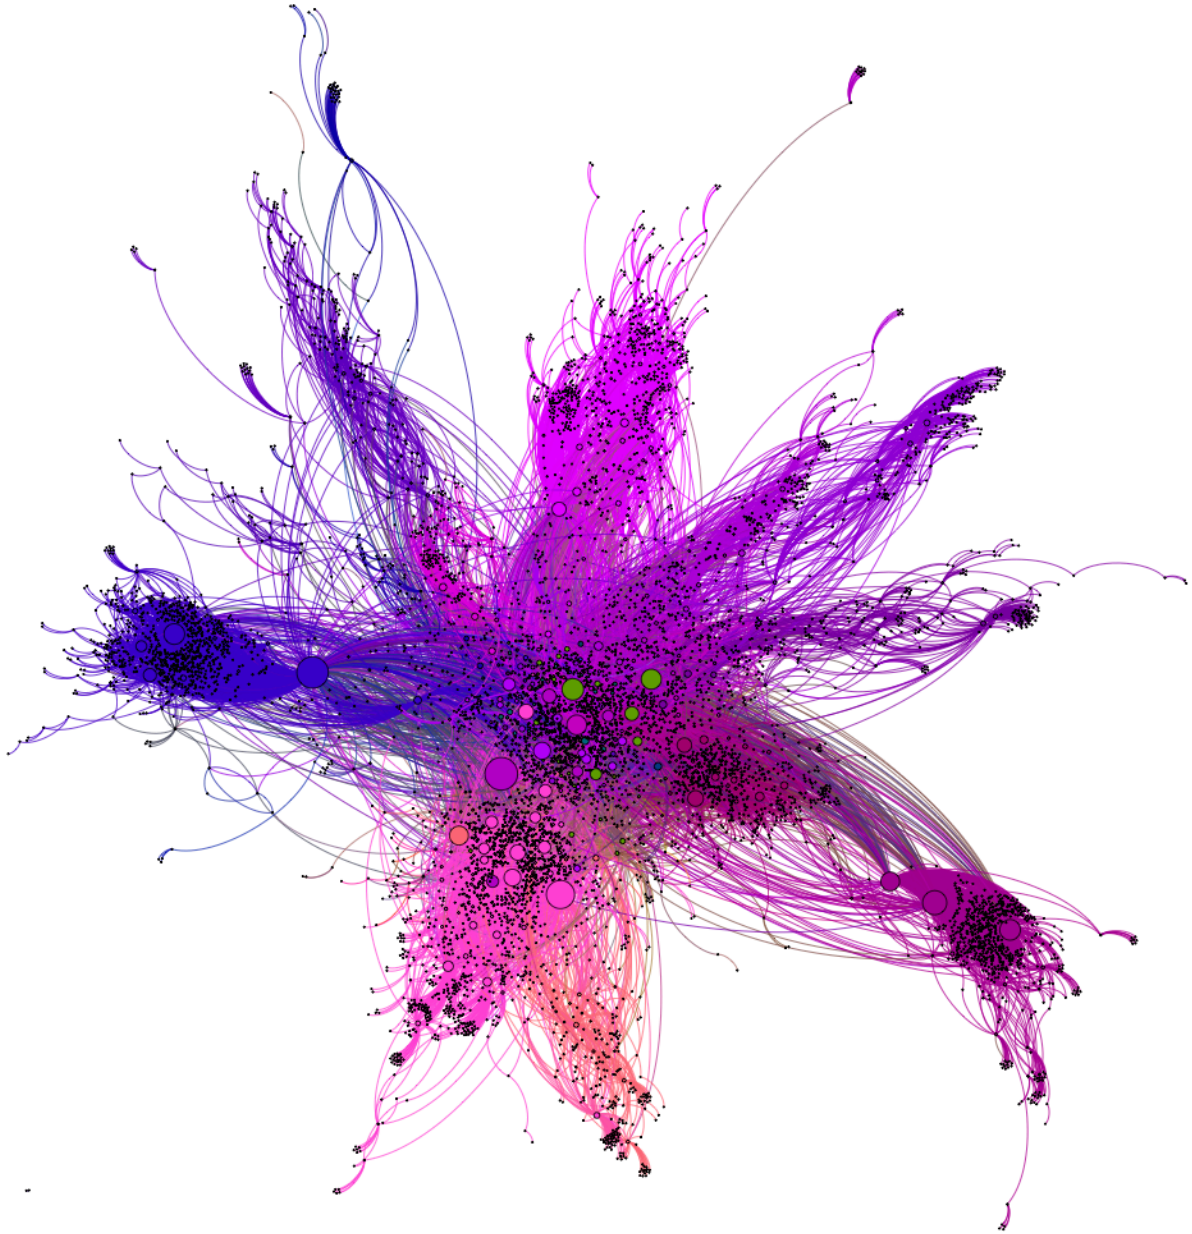

Europe (RIPE NCC)

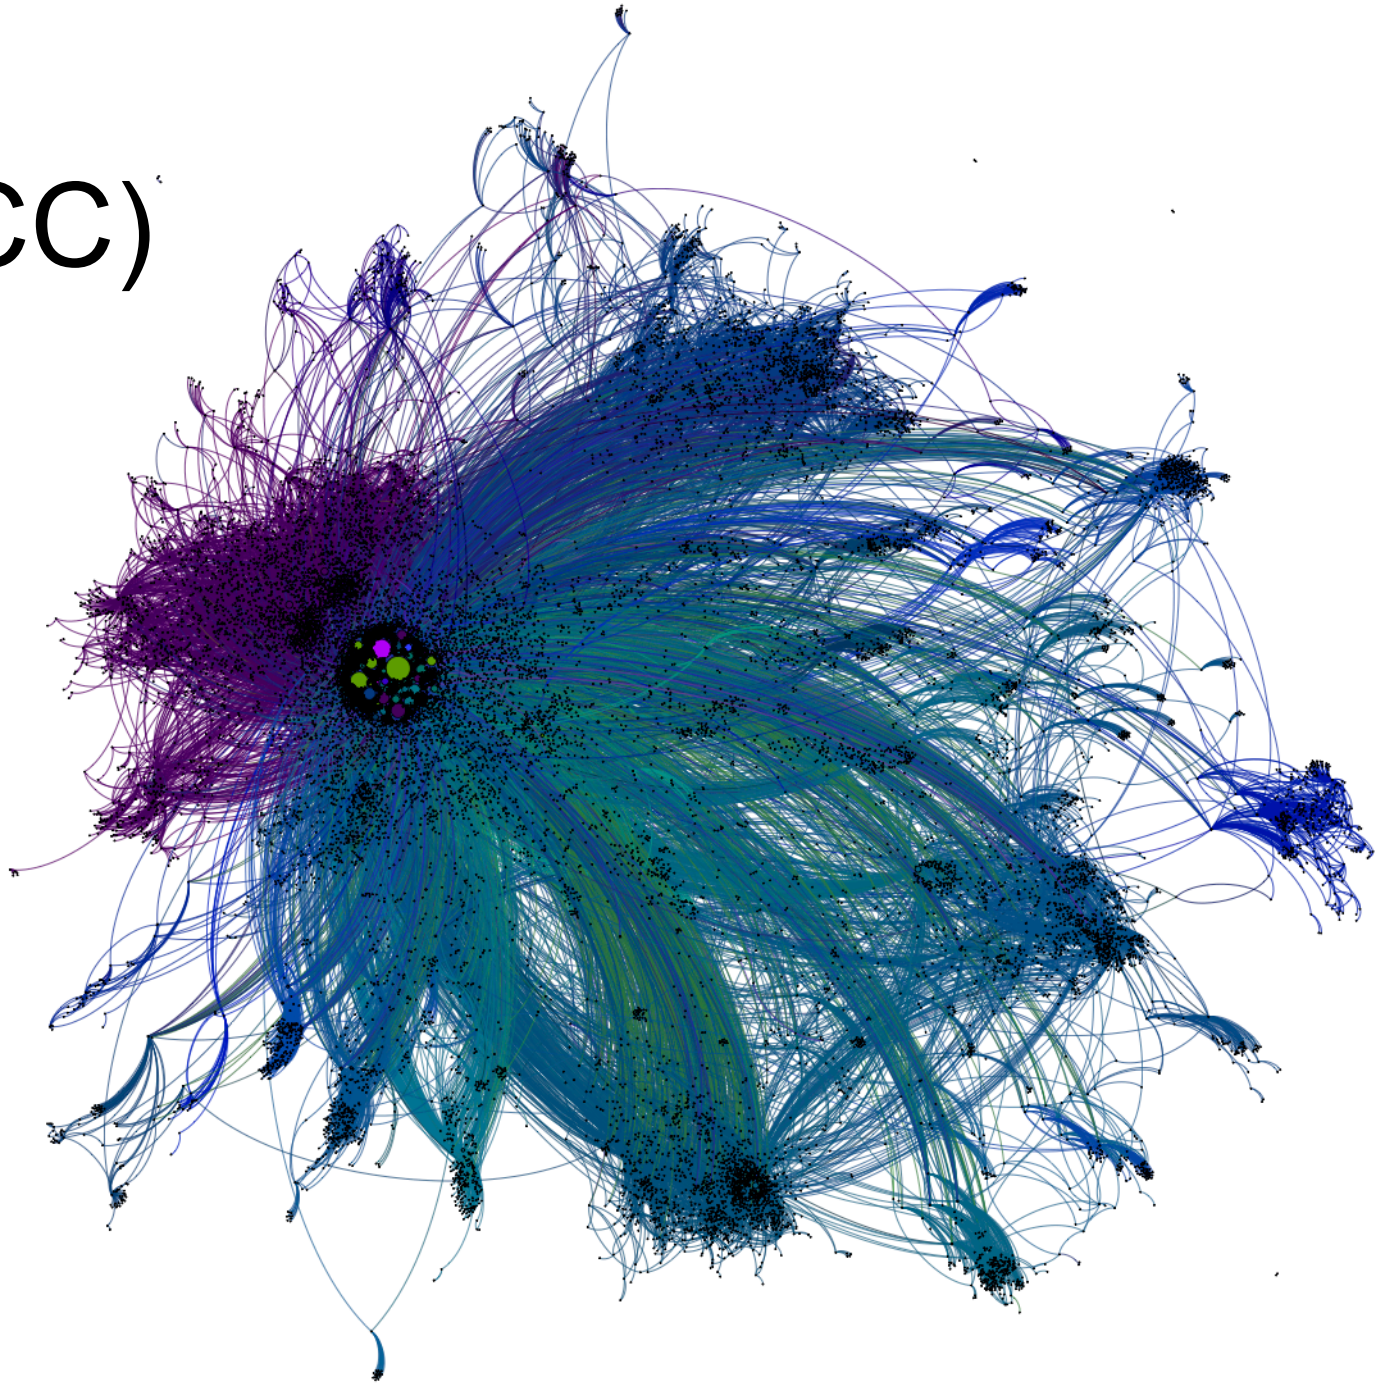

UNIVERS

North America (ARIN)

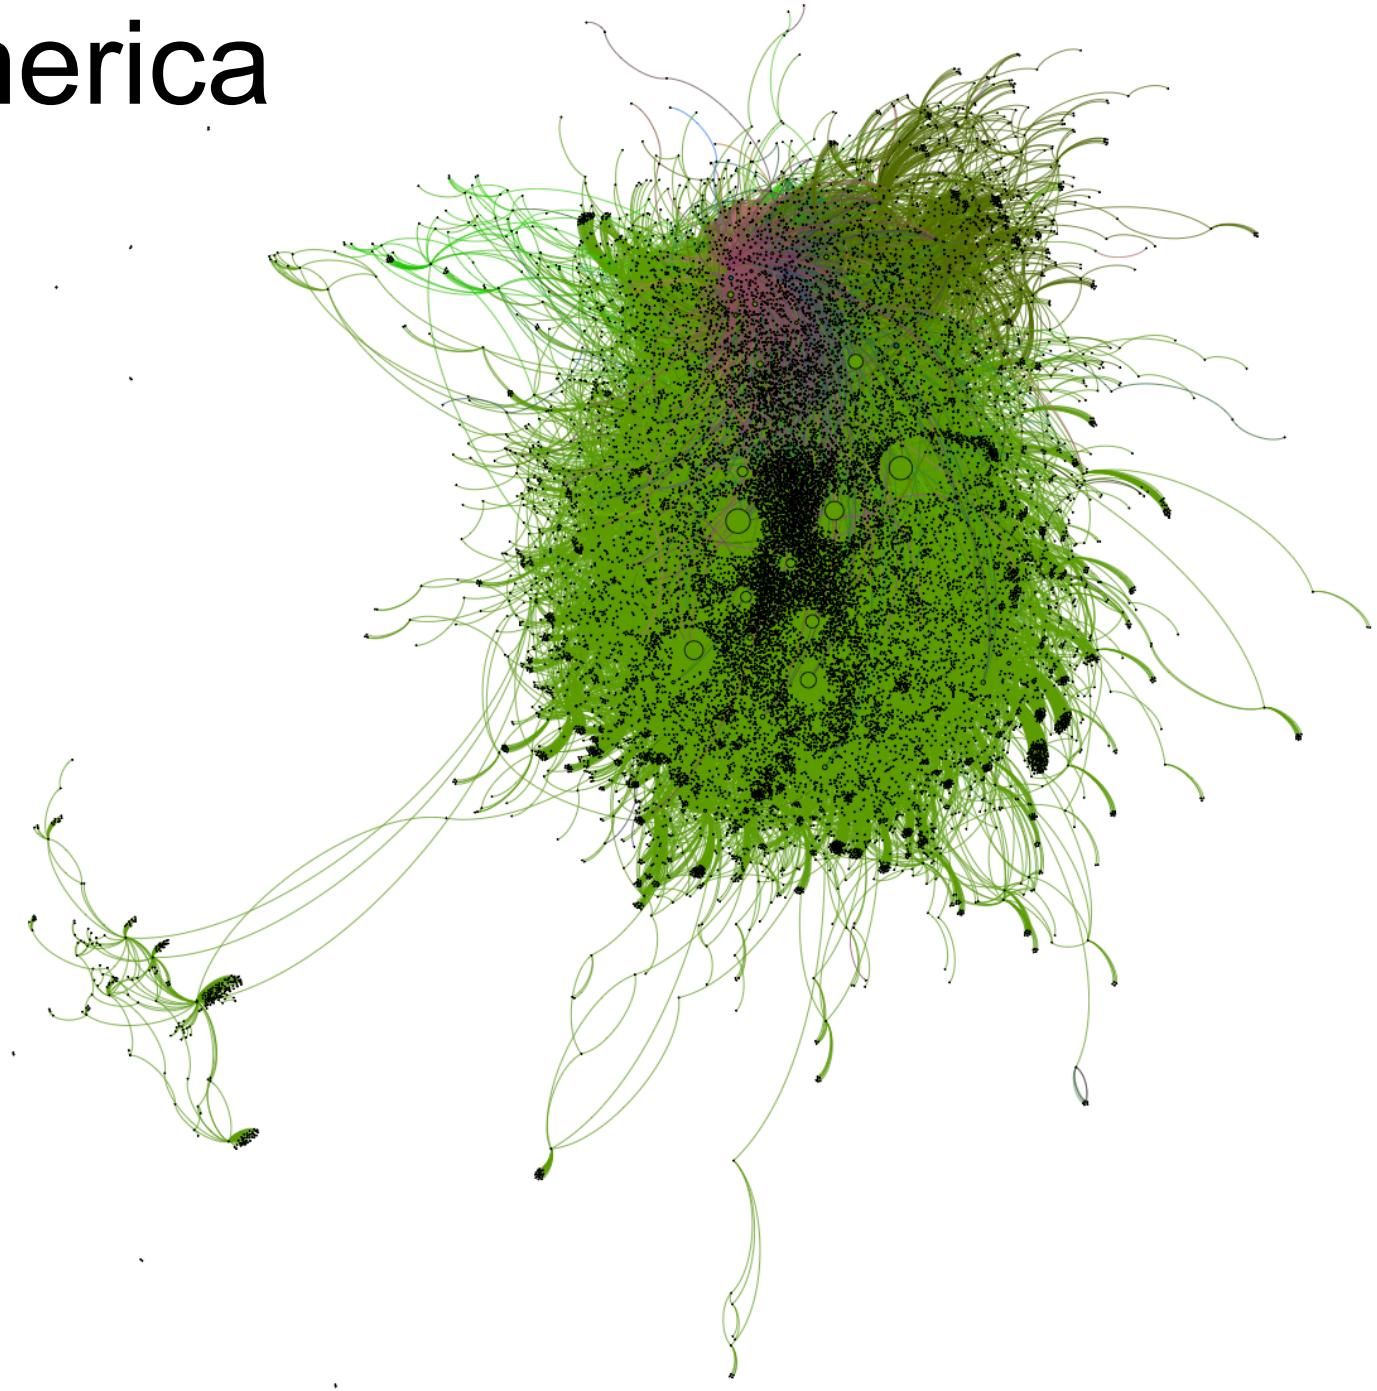

UNIVERSITY

Africa (AfriNIC)

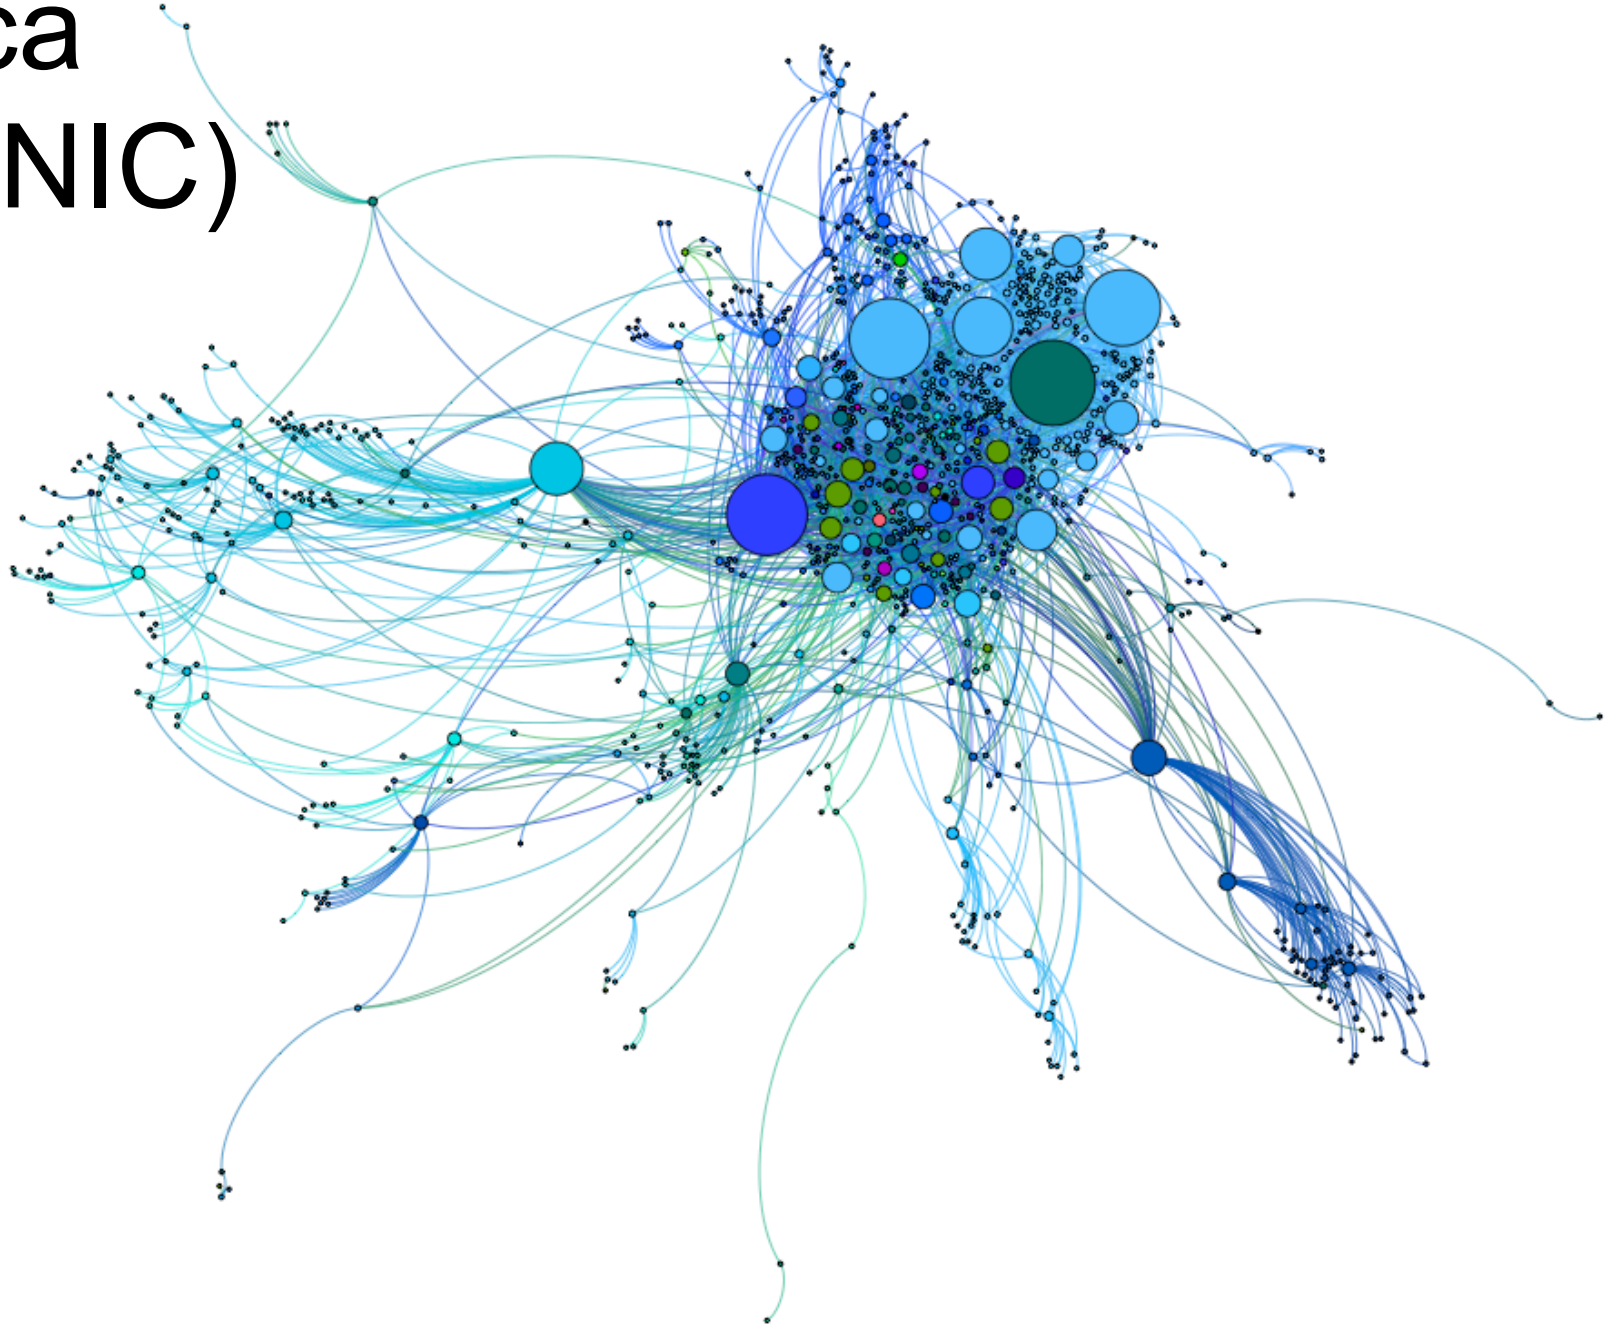

Latin America (LACNIC)

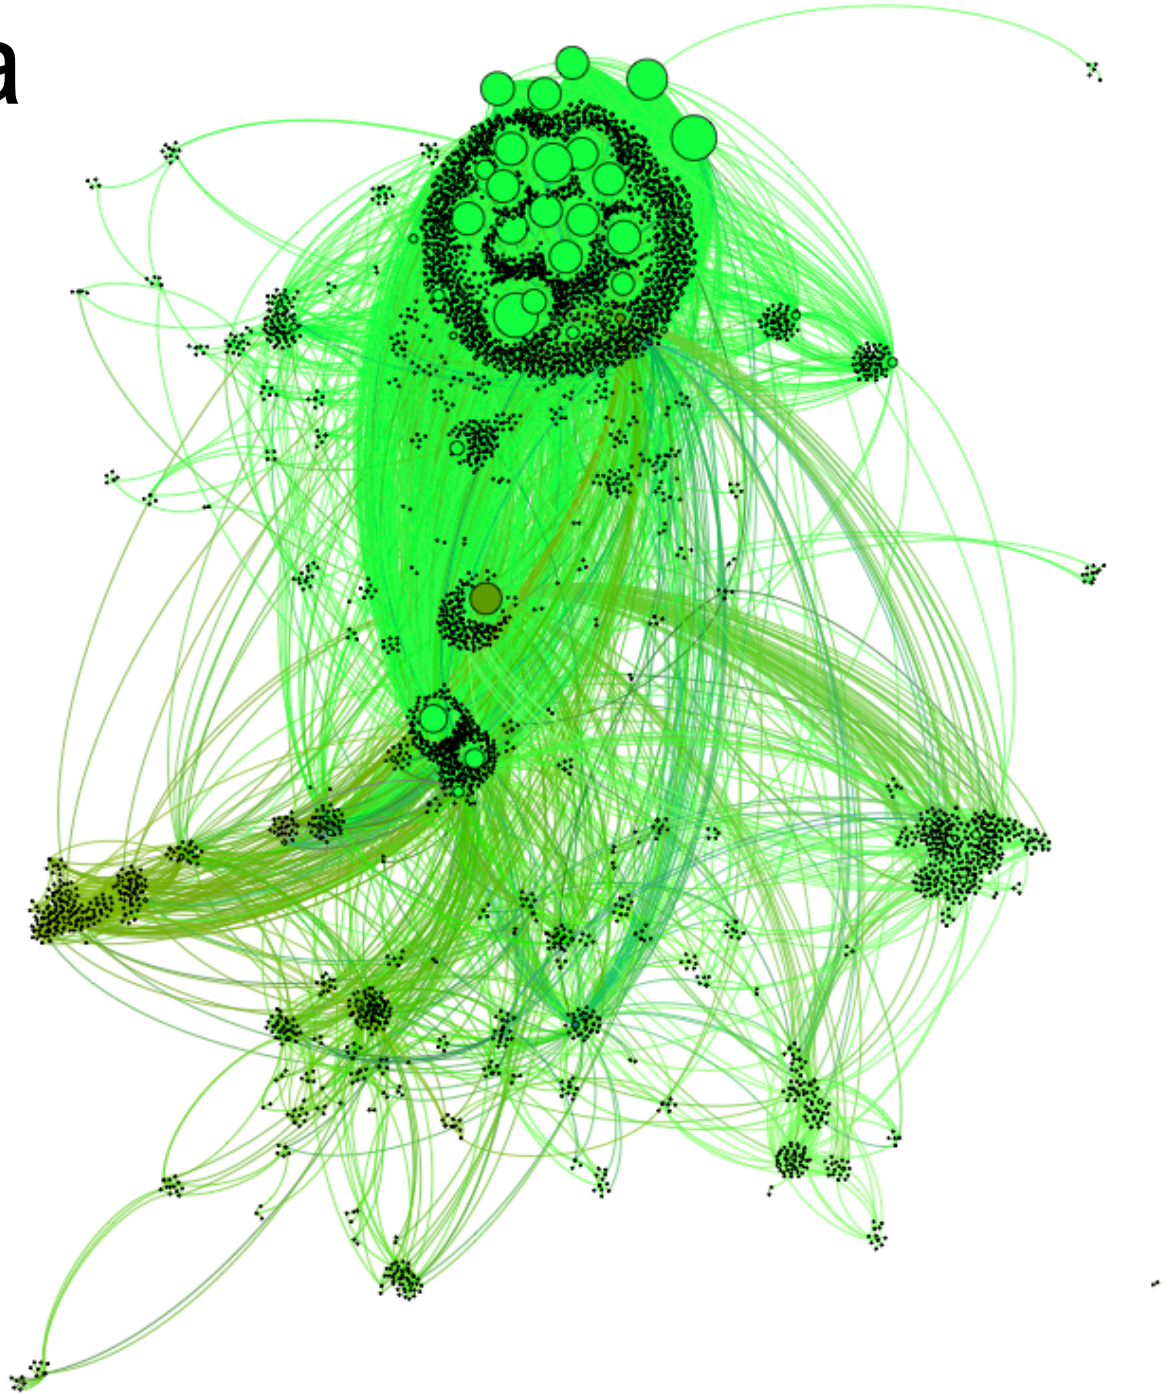

Middle East

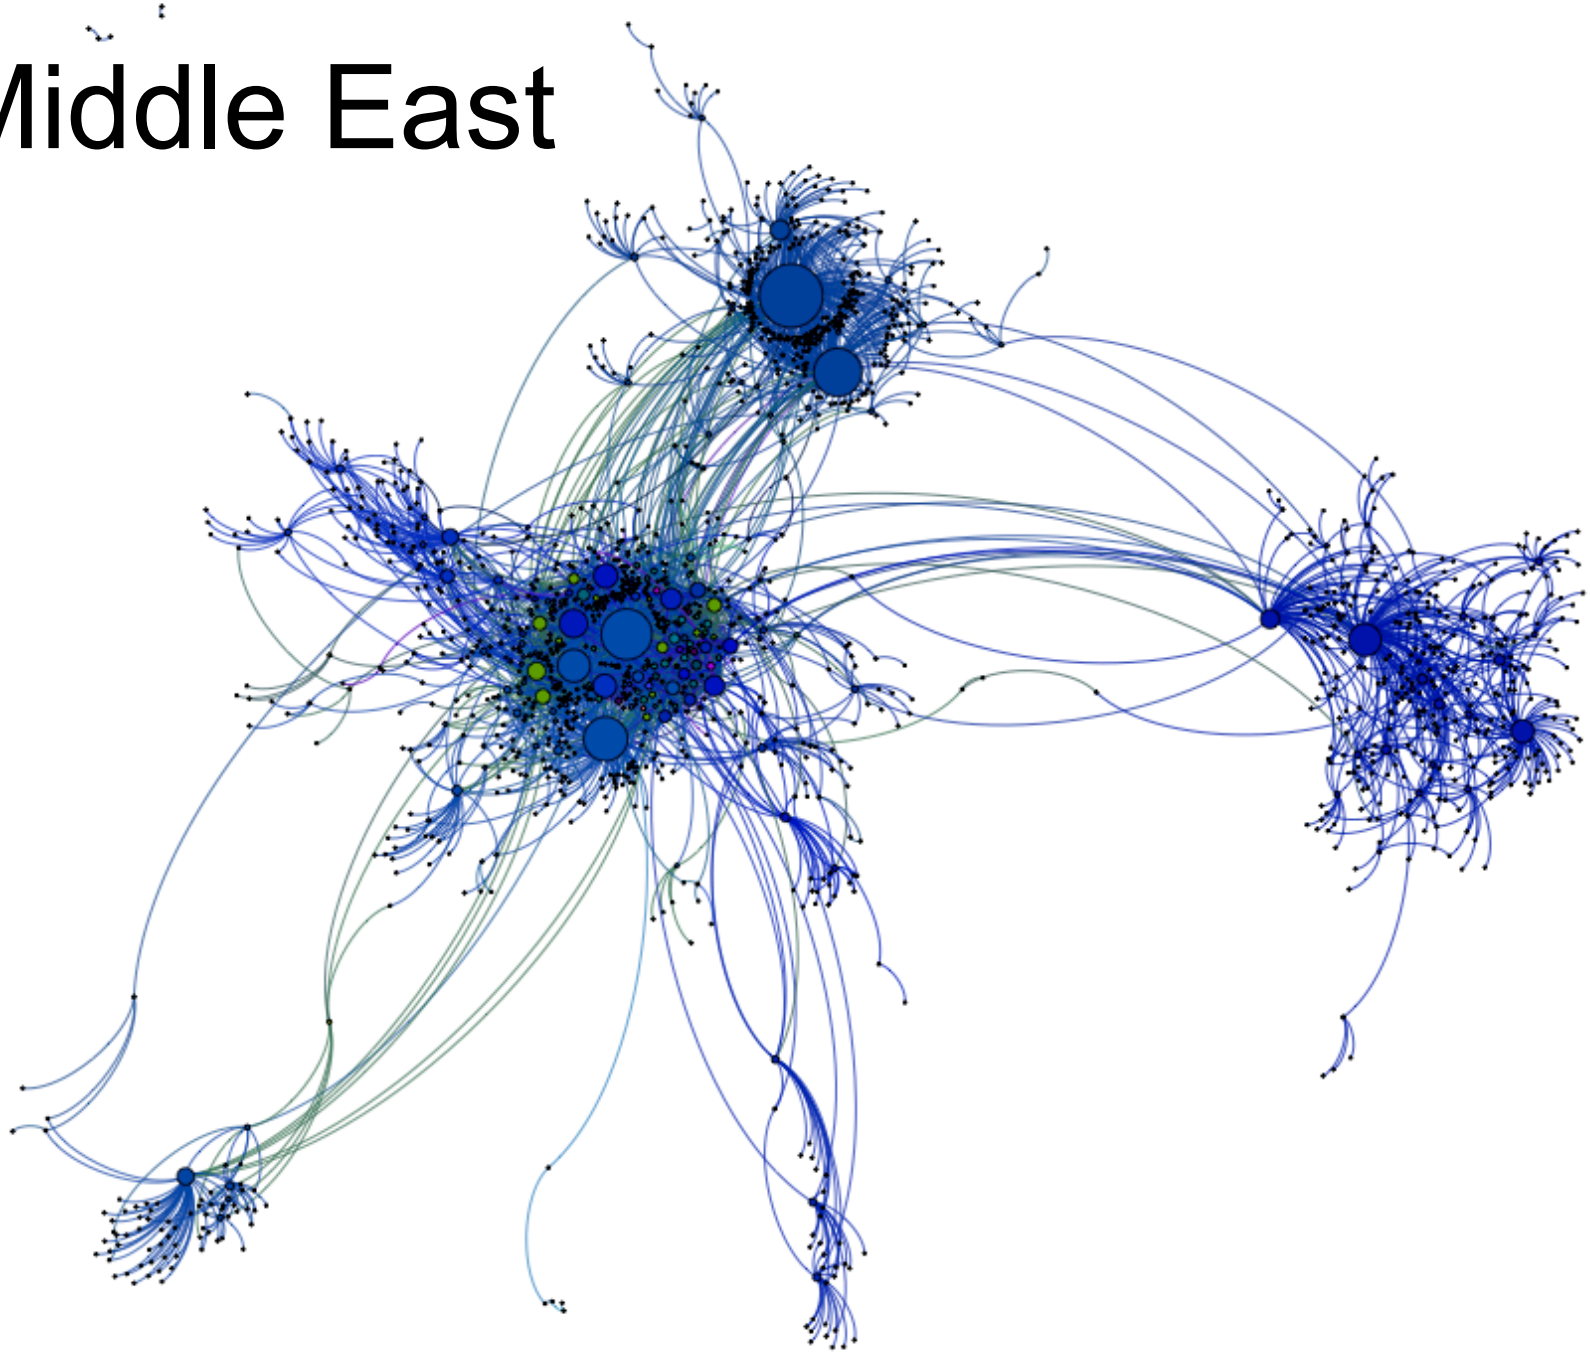

Oceania

(minus AU,NZ)

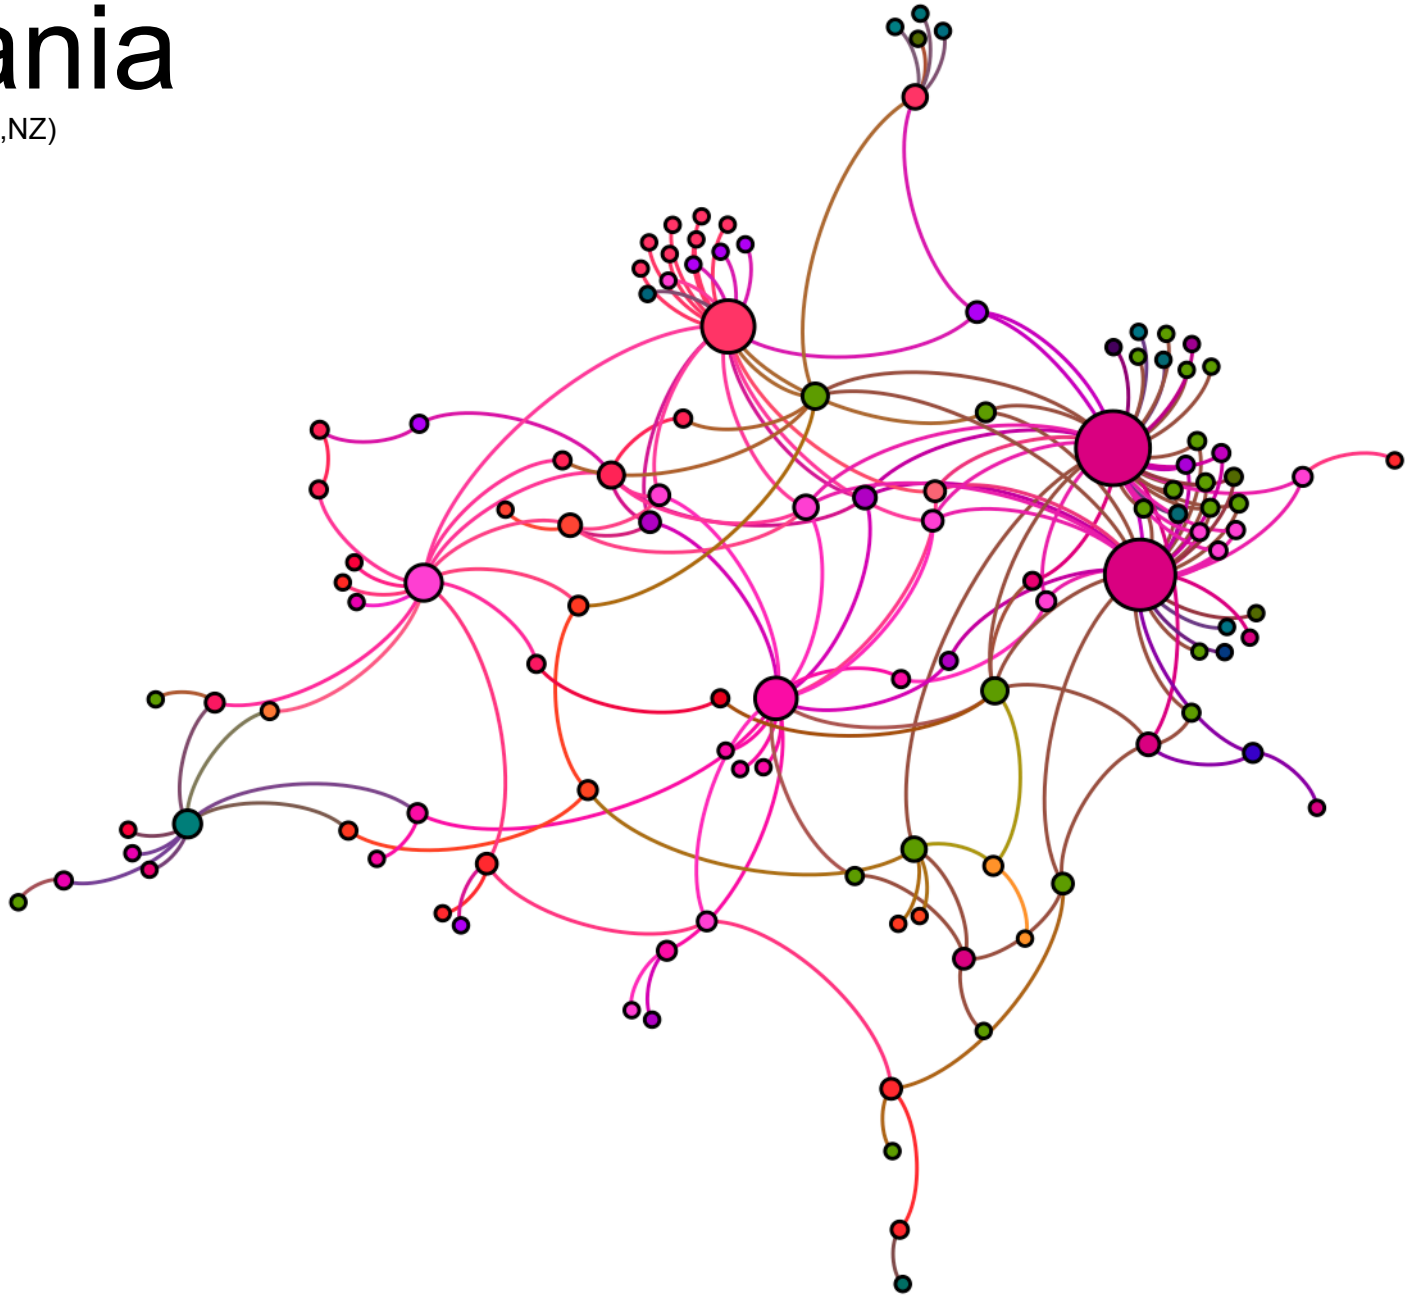

UNIVERSI

The Internet

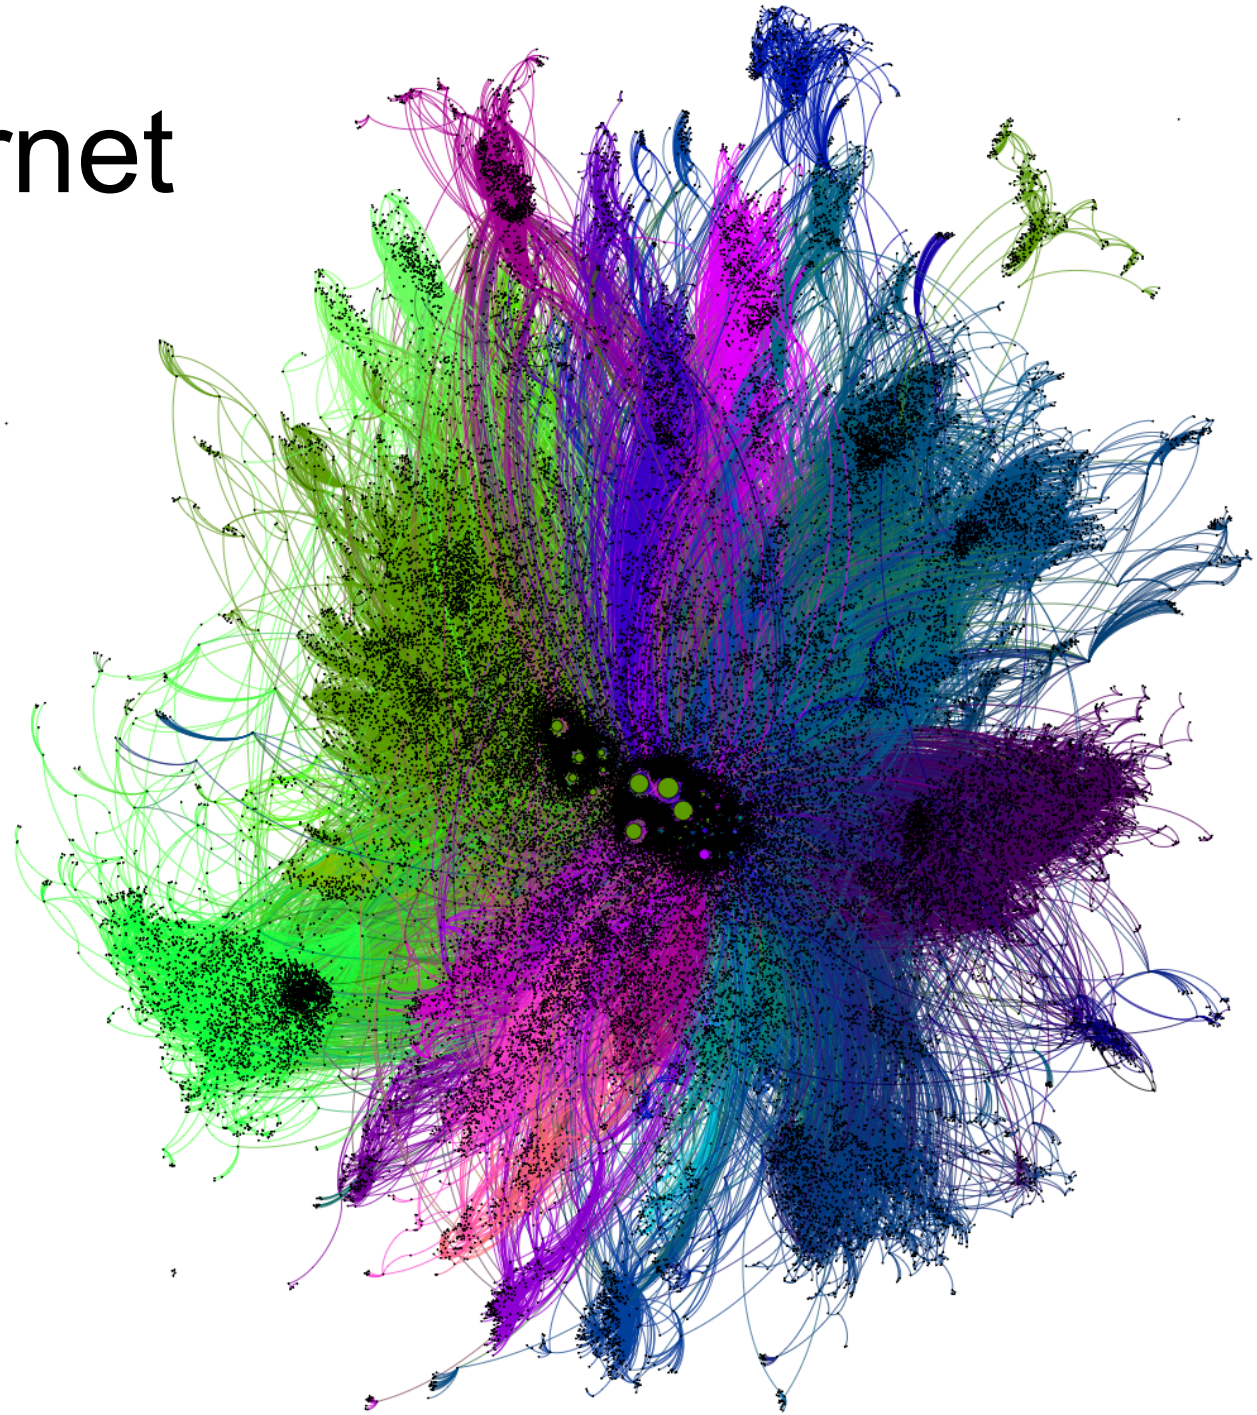

UNIVERSITY OF ORE

What

- What are we going to map when we try and visualise the Internet?
- Every cable?
- Some cables?
- Every Computer?
- Just companies?
- ??????

Autonomous Systems

- We already have a concept of entities on the Internet.
- RFC1771
  - "The classic definition of an Autonomous System is a set of routers under a single technical administration"
- ASNs use BGP sessions to exchange routes.
- If we could map these sessions we would have a map of the Internet.

Different to traffic flow

- This isn't exactly the same as looking at something like a traceroute.
- But it's much easier to get BGP AS Path data than a traceroute between every ASN on the Internet.
- That would be great though.

Optimal paths only

- BGP hides things.
- If there are multiple paths between two ASNs, BGP will only how you the best.
- We use multiple views to get around this problem.
- We still miss things.

How

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Where to get paths from

- So it's going to all come down to getting AS Paths.
- The more paths we can get the more neighbourings we can infer.
- The more neighbours the better the graph looks.

Where are some places we can get AS Paths from?

RouteViews

- http://www.routeviews.org/
- Over 18 collection points including:
  - "U of Oregon, Eugene Oregon, USA"
  - "Equinix, Ashburn, VA"
  - "ISC (PAIX), Palo Alto CA, USA"
  - "KIXP, Nairobi, Kenya"
  - "LINX, London, GB"
  - "NWAX, Portland, Oregon"
  - "DIXIE (NSPIXP), Tokyo, Japan"
  - "SYDNEY (SYD1 Equinix), Sydney, Australia"
  - "SAOPAULO (PTT Metro, NIC.br), Sao Paulo, Brazil"
  - "TELXATL (TELX Atlanta), Atlanta, Georgia"
  - "Colorado State University Fort Collins, CO"

RIS

- The RIPE NCC collects and stores Internet routing data from several locations around the globe, using the Routing Information Service (RIS), established in 2001.

RIS

- rrc00.ripe.net at RIPE NCC, Amsterdam
- rrc01.ripe.net at LINX, London
- rrc02.ripe.net at SFINX, Paris.
- rrc03.ripe.net at AMS-IX, Amsterdam
- rrc04.ripe.net at CIXP, Geneva
- rrc05.ripe.net at VIX, Vienna
- rrc06.ripe.net at Otemachi, Japan
- rrc07.ripe.net in Stockholm, Sweden
- rrc08.ripe.net at San Jose (CA), USA.
- rrc09.ripe.net at Zurich, Switzerland
- rrc10.ripe.net at Milan, Italy.
- rrc11.ripe.net at New York (NY), USA.
- rrc12.ripe.net at Frankfurt, Germany.
- rrc13.ripe.net at Moscow, Russia
- rrc14.ripe.net at Palo Alto, USA.
- rrc15.ripe.net at Sao Paulo, Brazil.
- rrc16.ripe.net at Miami, USA

‘show ip bgp’

- Various other sources of AS Paths
- Wellington Internet Exchange (WIX)
- http://nzix.net/cgi-bin/lg.cgi

```
rs1.wix.nzix.net-bgp> show ip bgp regex ^
BGP table version is 0, local router ID is 202.7.0.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 13.64.0.0/11     202.7.0.220             100             0 8075 i
*> 13.96.0.0/13     202.7.0.220             100             0 8075 i
*> 13.104.0.0/14    202.7.0.220             100             0 8075 i
*> 13.107.1.0/24    202.7.0.220             100             0 8075 8075 8068 i
*> 13.107.10.0/24   202.7.0.220             100             0 8075 8068 i
*> 13.107.44.0/24   202.7.0.220             100             0 8075 8068 i
*> 13.107.58.0/24   202.7.0.220             100             0 8075 8068 i
```

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

RIBs vs BGP UPDATES

- RIB = Routing Information Base
  - Pretty much what most people thing about as the ‘Route Table’
- BGP UPDATES
  - These are the actual UPDATE and WITHDRAW messages as they are communicated by the BGP protocol

Why UPDATES

- As routes are added and withdrawn, we may get to see some of those sub-optimal routes which are normally hidden by BGP.

MRT format

- https://tools.ietf.org/html/draft-ietf-grow-mrt-09
- This format was developed in concert with the Multi-threaded Routing Toolkit (MRT) from whence the format takes it name. The format can be used to export routing protocol messages, state changes, and routing information base contents.
- A lot of the RouteViews and RIS files are in MRT dump format. We need to be able to parse them to extract the Paths.

Turning into CSV

- Parse the MRT files and output one line per path.
- https://github.com/YoshiyukiYamauchi/mrtparse
- 1.0.0.0,24,10026 15169
- 1.0.0.0,24,10102 15169
- 1.0.0.0,24,11039 15169
- 1.0.0.0,24,11666 15169
- 1.0.0.0,24,11686 15169
- 1.0.0.0,24,1221 15169
- 1.0.0.0,24,1239 15169
- 1.0.0.0,24,1299 15169
- 1.0.0.0,24,13030 15169
- 1.0.0.0,24,13237 15169
- 1.0.0.0,24,133165 15169

Sort – Uniq

- If we take that output from a lot of different sources then we are going to get duplicates.

- We need a way to sort these CSV files and remove duplicates

GNU Parallels

```
$ cat *.csv | \
 parallel --pipe --files sort | \
 parallel -Xj1 sort -m {} ';' rm {} | \
 uniq > route-views.large

$ wc -l route-views.large
 68386677 route-views.large

$ du -h route-views.large
2.6G route-views.large
```

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

From Path to Peering

If we have the following path:

1.0.128.0,18,19151 3356 3257 38040 9737

We now know that the following peerings exist

- 19151 3356
- 3356 3257
- 3257 38040
- 38040 9737

Make a list of ASNs

- From this we get a list of all the ASNs in all the paths.
- We can use this to make a list of ‘nodes’ for the graphs.

```
Id,Country,color,Label,Lat,Lon
10000001,US,#5e9e00,AS1-US-LVLT-1 - Level 3 Communications, Inc.,,
10000010,US,#5e9e00,AS10-US-CSNET-EXT-AS - CSNET Coordination and
10000100,US,#5e9e00,AS100-US-FMC-CTC - FMC Central Engineering
10001000,CA,#556d00,AS1000-CA-GONET-ASN-17 - GONET,,
10010000,JP,#a00070,AS10000-JP-NCM Nagasaki Cable Media Inc.,,
```

Make a list of Neighbours

- We can also take all the peerings and build a list of ‘edges’

```
Source,Target,Type,Id,Label,Weight,Routes
10031353,10049426,Undirected,1,,10,116.0
10031353,10048890,Undirected,2,,10,48.0
10042979,10002588,Undirected,3,,10,9444.0
10042979,10024651,Undirected,4,,10,2047.0
10042979,10024607,Undirected,5,,10,949.0
10042979,10041424,Undirected,6,,10,240.0
10042979,10043615,Undirected,7,,10,89.0
10042979,10008194,Undirected,8,,10,556.0
```

What about those colours though?

- I didn't want to tie the ASNs to geographic locations on the map.
- BGP has no idea of where you are, so I wanted to lay put the graph using the BGP topology, not our physical boarders.
- I knew I'd have to represent countries somehow though. I chose to use colour.

Country Colours

- Different Country = Different Colour.
- Close countries needed similar colours.
- How to do this?

- HSL colour model
- Hue 0-360 deg
- Longitude -180 – 180 deg
- Hmmmmmm I can work with that.

Country Colours

- Latitude = -90 – 90 deg.
- Saturation = 0-100
- Brightness (or Level) = 0-100

- I can make Northern Hemisphere more or less bright.
- And Southern Hemisphere more or less saturated

Colours for Countries

Gephi

- https://gephi.github.io/

Gephi
makes graphs handy

Download | Blog | Wiki | Forum | Support | Bug tracker

Home | Features | Learn | Develop | Plugins | Services | Consortium

The Open Graph Viz Platform

Gephi is an interactive visualization and exploration platform for all kinds of networks and complex systems, dynamic and hierarchical graphs.

Runs on Windows, Linux and Mac OS X. Gephi is open-source and free.

Learn More on Gephi Platform »

Download FREE
Gephi 0.8.2-beta

Release Notes | System Requirements

- Features
- Quick start
- Screenshots
- Videos

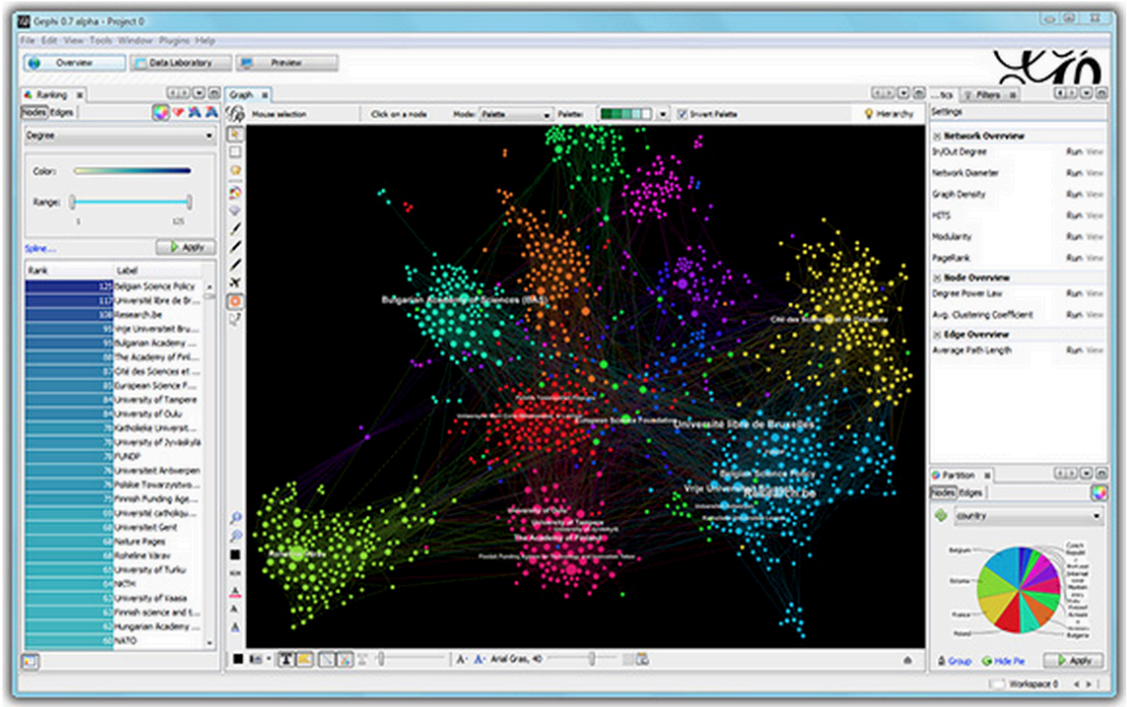

Plugins

- OpenOrd Layout
  - https://marketplace.gephi.org/plugin/openord-layout/
- Give Colour To Nodes
  - https://marketplace.gephi.org/plugin/give-color-to-nodes/

Layouts – Combination of Art and Science

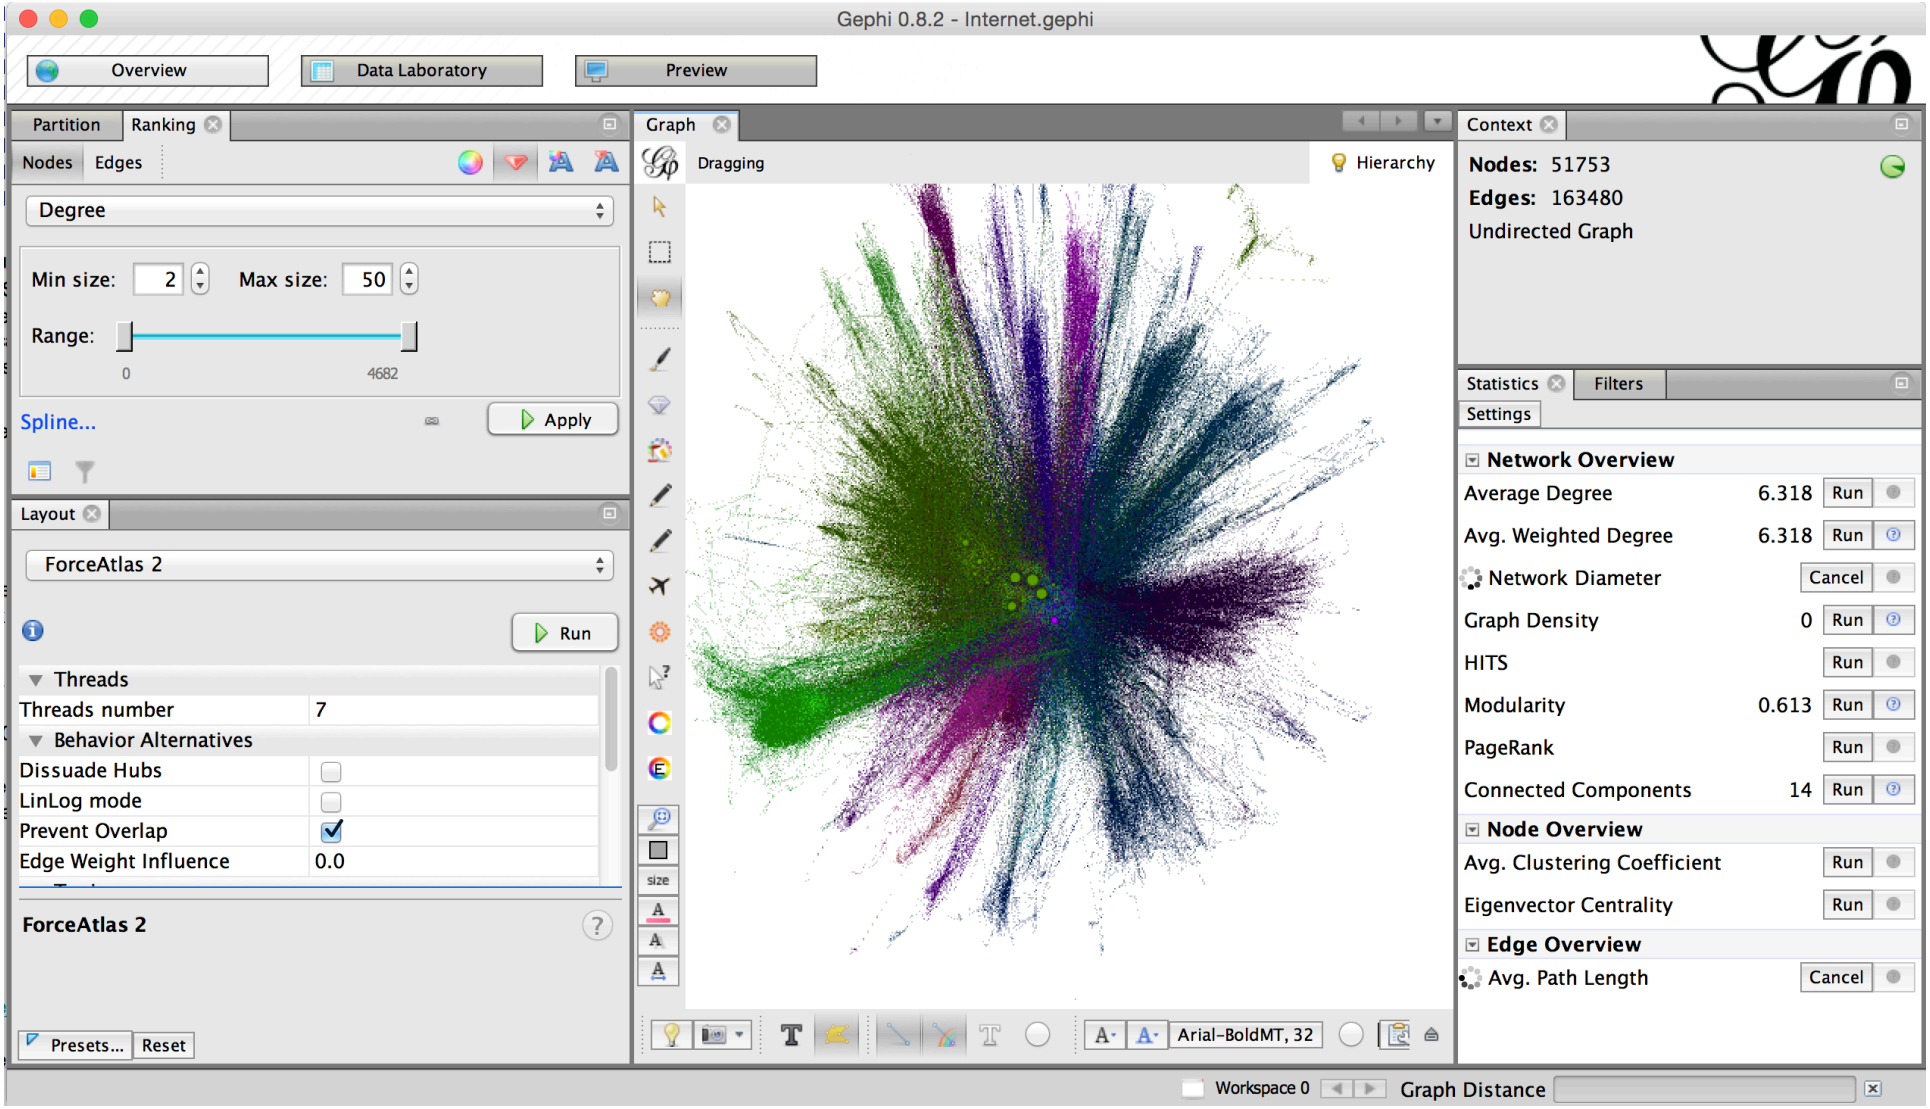

Output

- PDF
  - High Resolution
  - Vector
  - Able to add names but it gets REALLY BUSY
- PNG
  - Nice for mailing to people
  - Smaller size
- JavaScript Interactive website.

Interactive

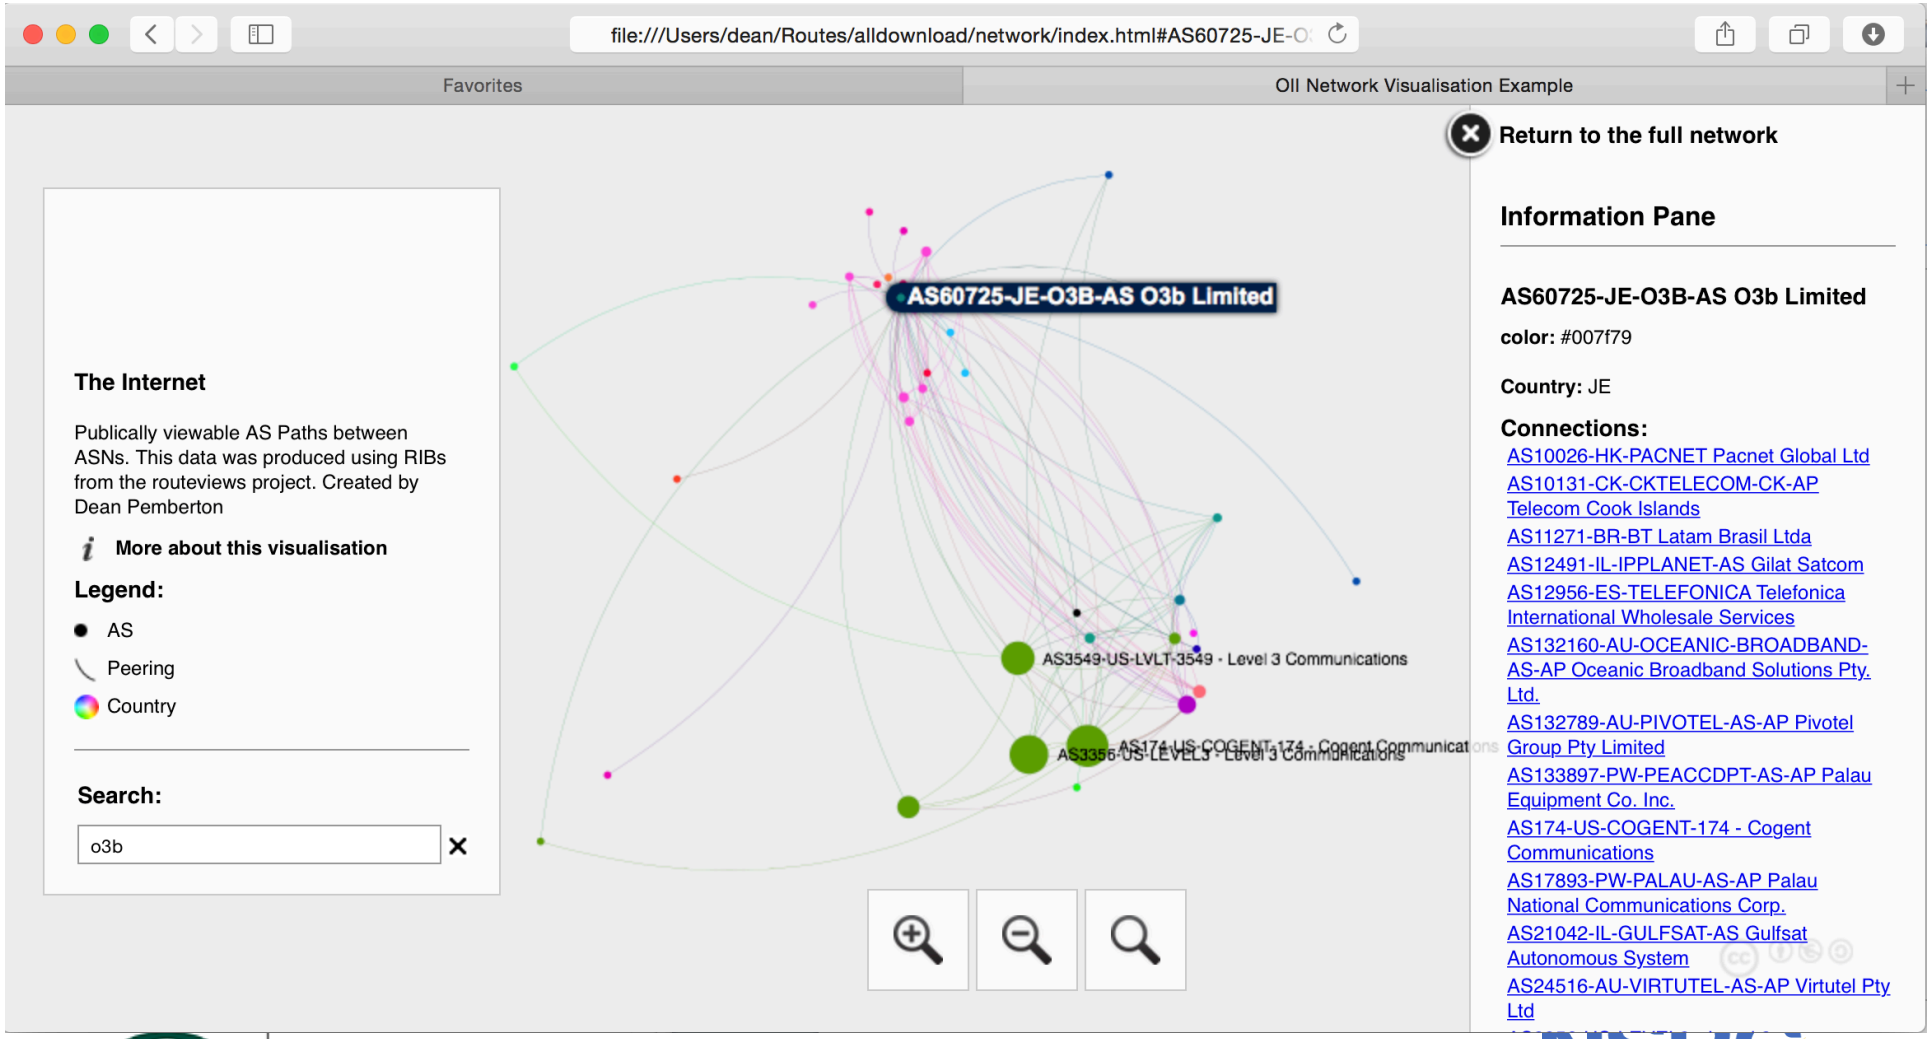

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Demo

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Interactive versions

- http://thyme.apnic.net/BGP/