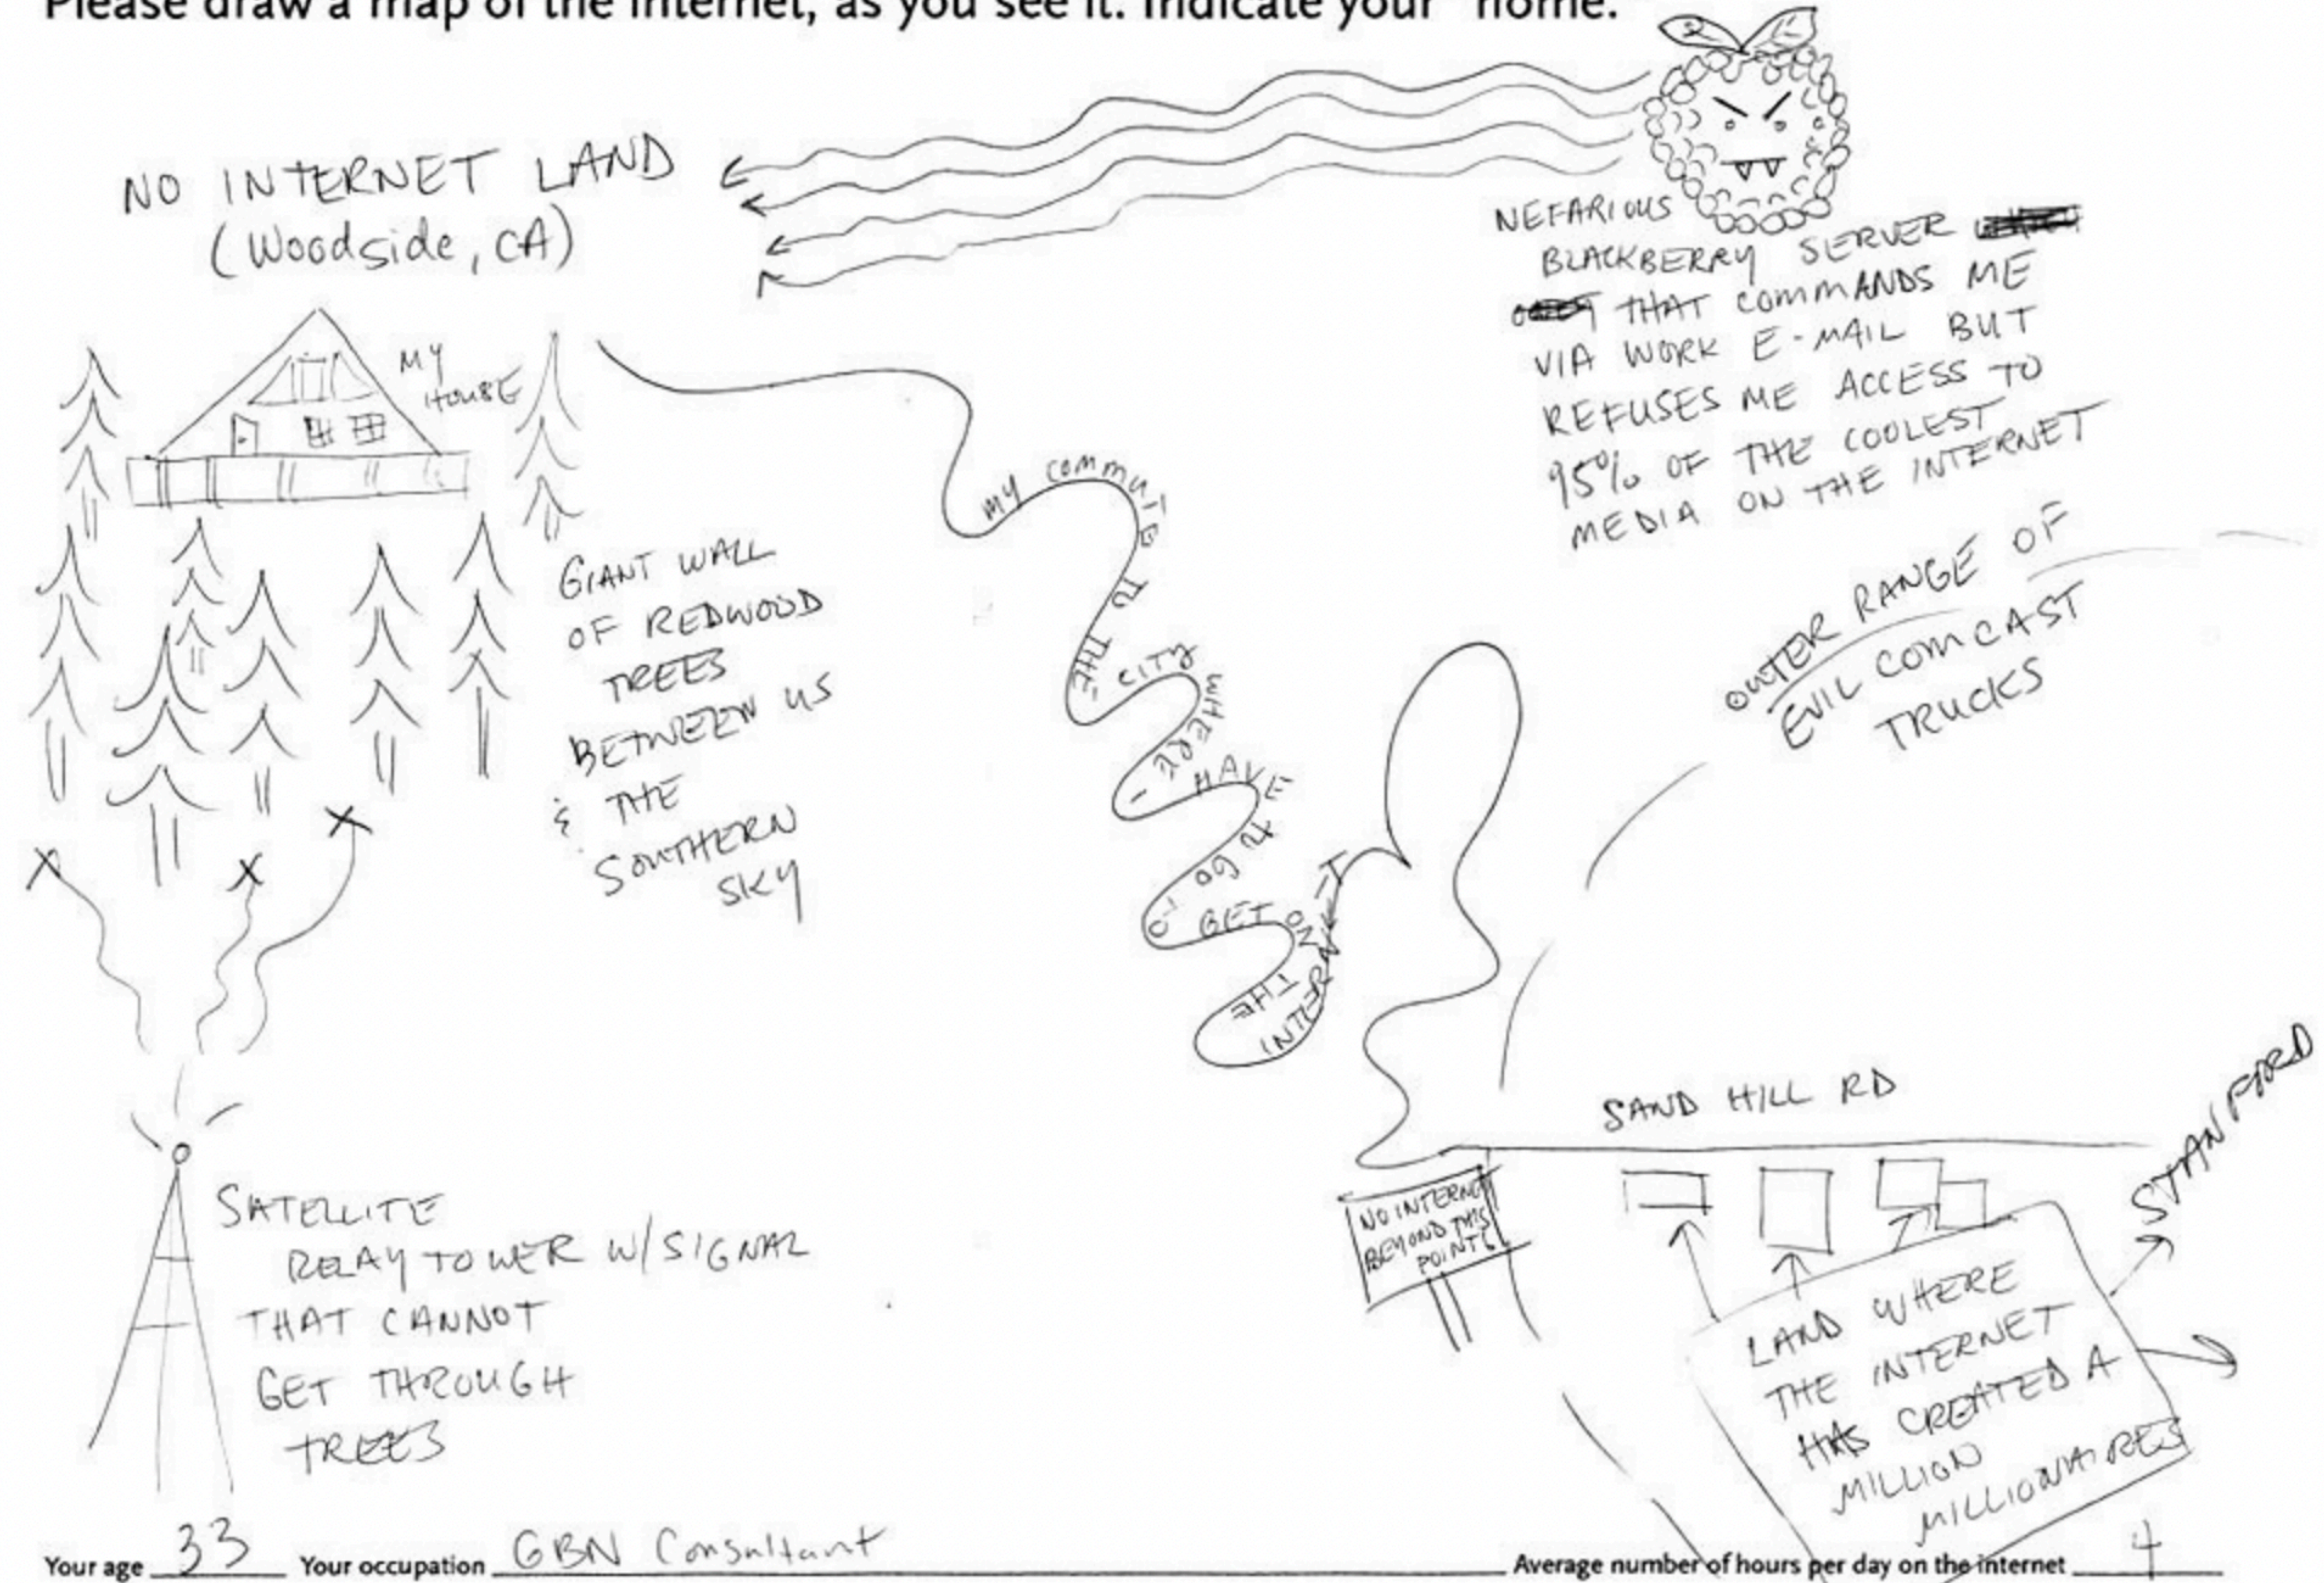


Towards an interactive visualization of the Internet

Maxime Mouchet — LIP6, Sorbonne Université


THE INTERNET MAPPING PROJECT

Please draw a map of the internet, as you see it. Indicate your "home."



THE INTERNET MAPPING PROJECT

Please draw a map of the internet, as you see it. Indicate your "home."




Your age _____ Your occupation _____ Average number of hours per day on the internet _____

Return to Kevin Kelly (kk@kk.org, or 149 Amapola Ave., Pacifica, CA 94004) For results see www.kk.org.

Kevin Kelly "The Internet Mapping Project"

THE INTERNET MAPPING PROJECT

Please draw a map of the internet, as you see it. Indicate your "home."



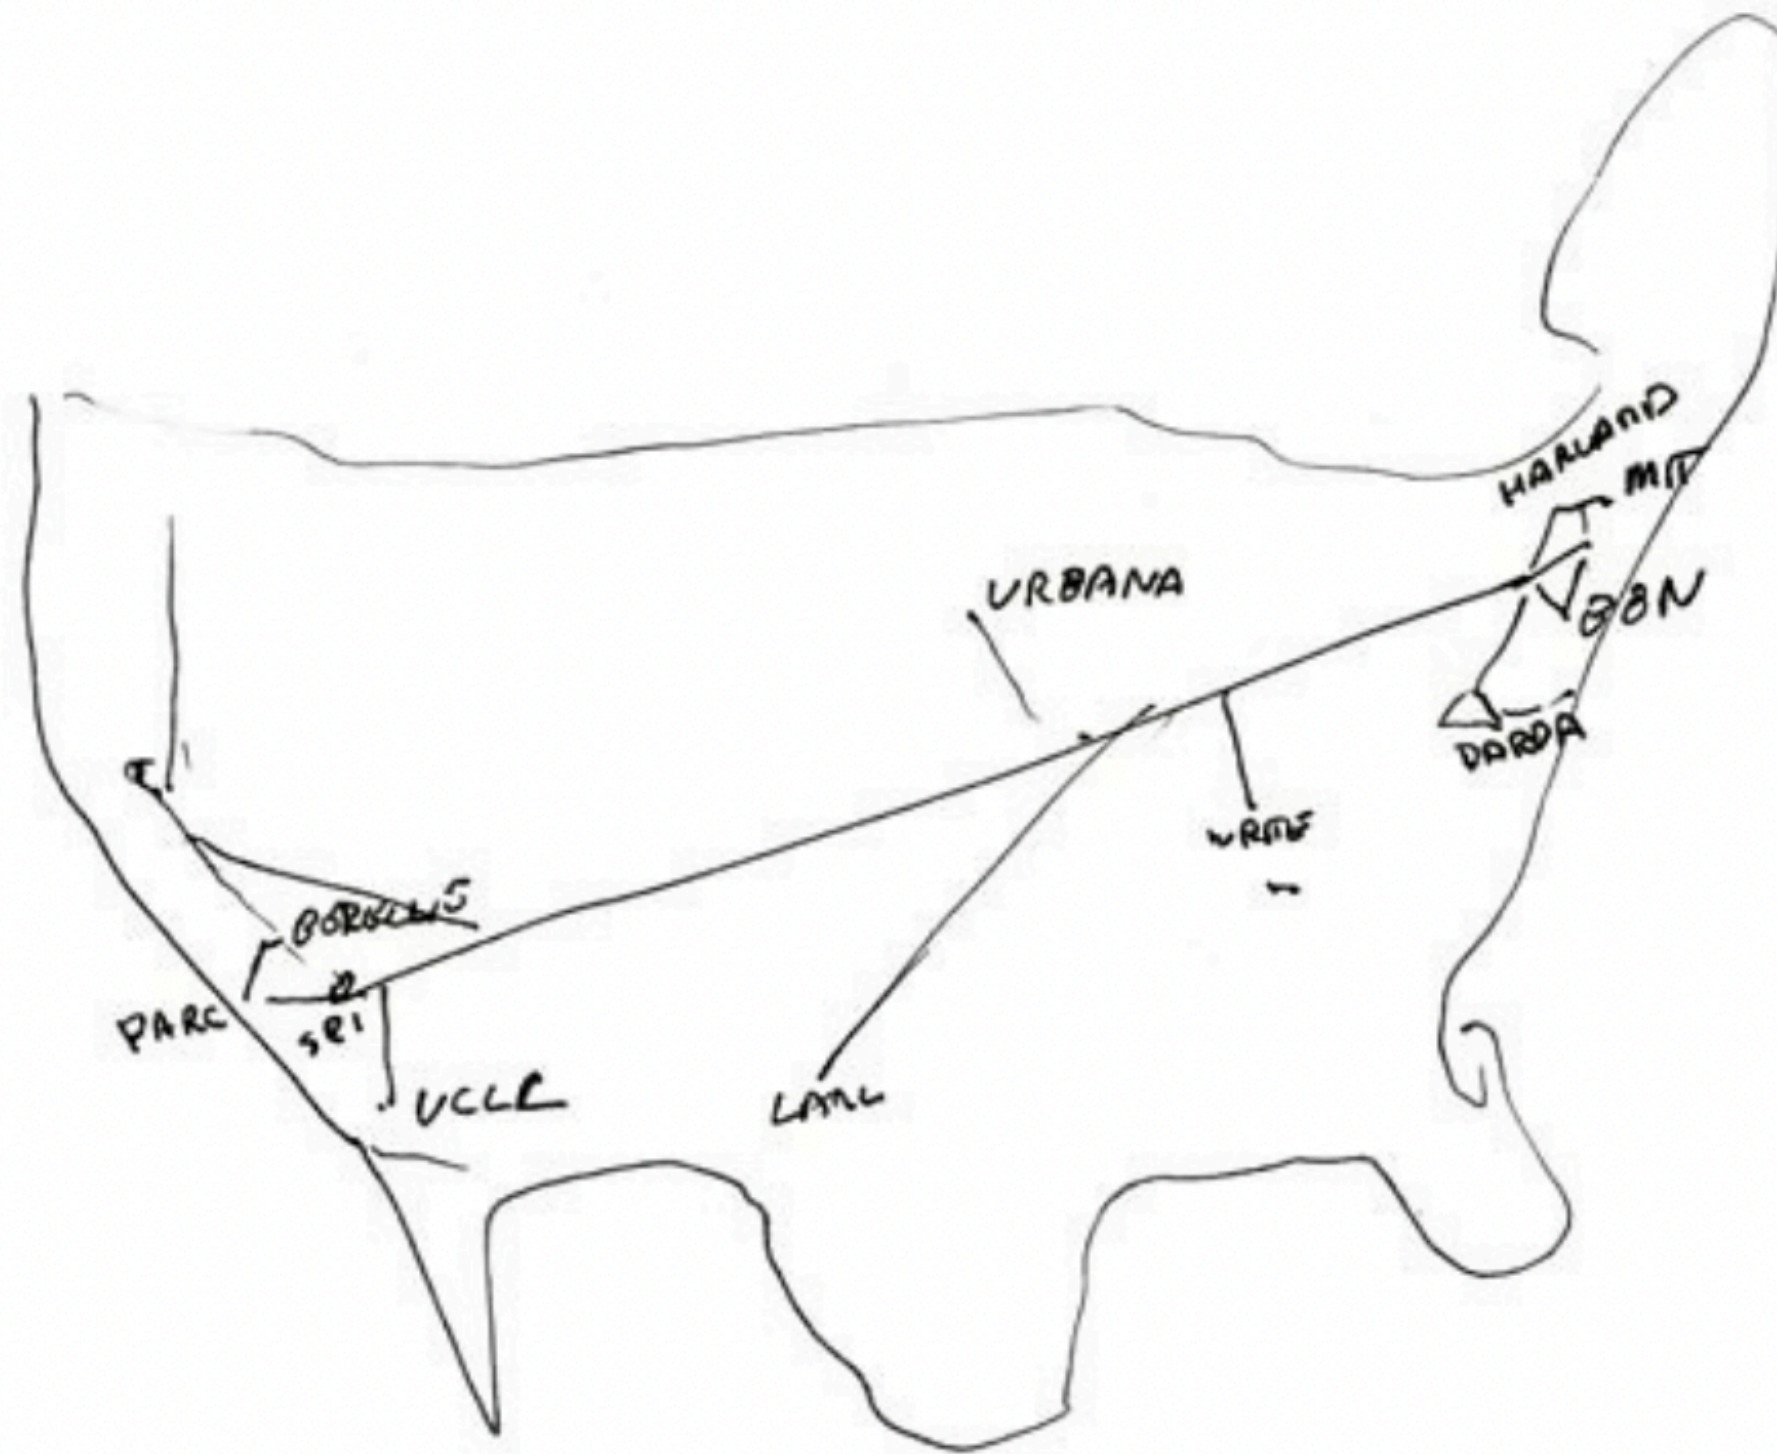
Your age 13 Your occupation School Average number of hours per day on the internet 1 hour on weekends

Return to Kevin Kelly (kk@kk.org, or 149 Amapola Ave., Pacifica, CA 94004) For results see www.kk.org.

Kevin Kelly "The Internet Mapping Project"

THE INTERNET MAPPING PROJECT

Please draw a map of the internet, as you see it. Indicate your "home."



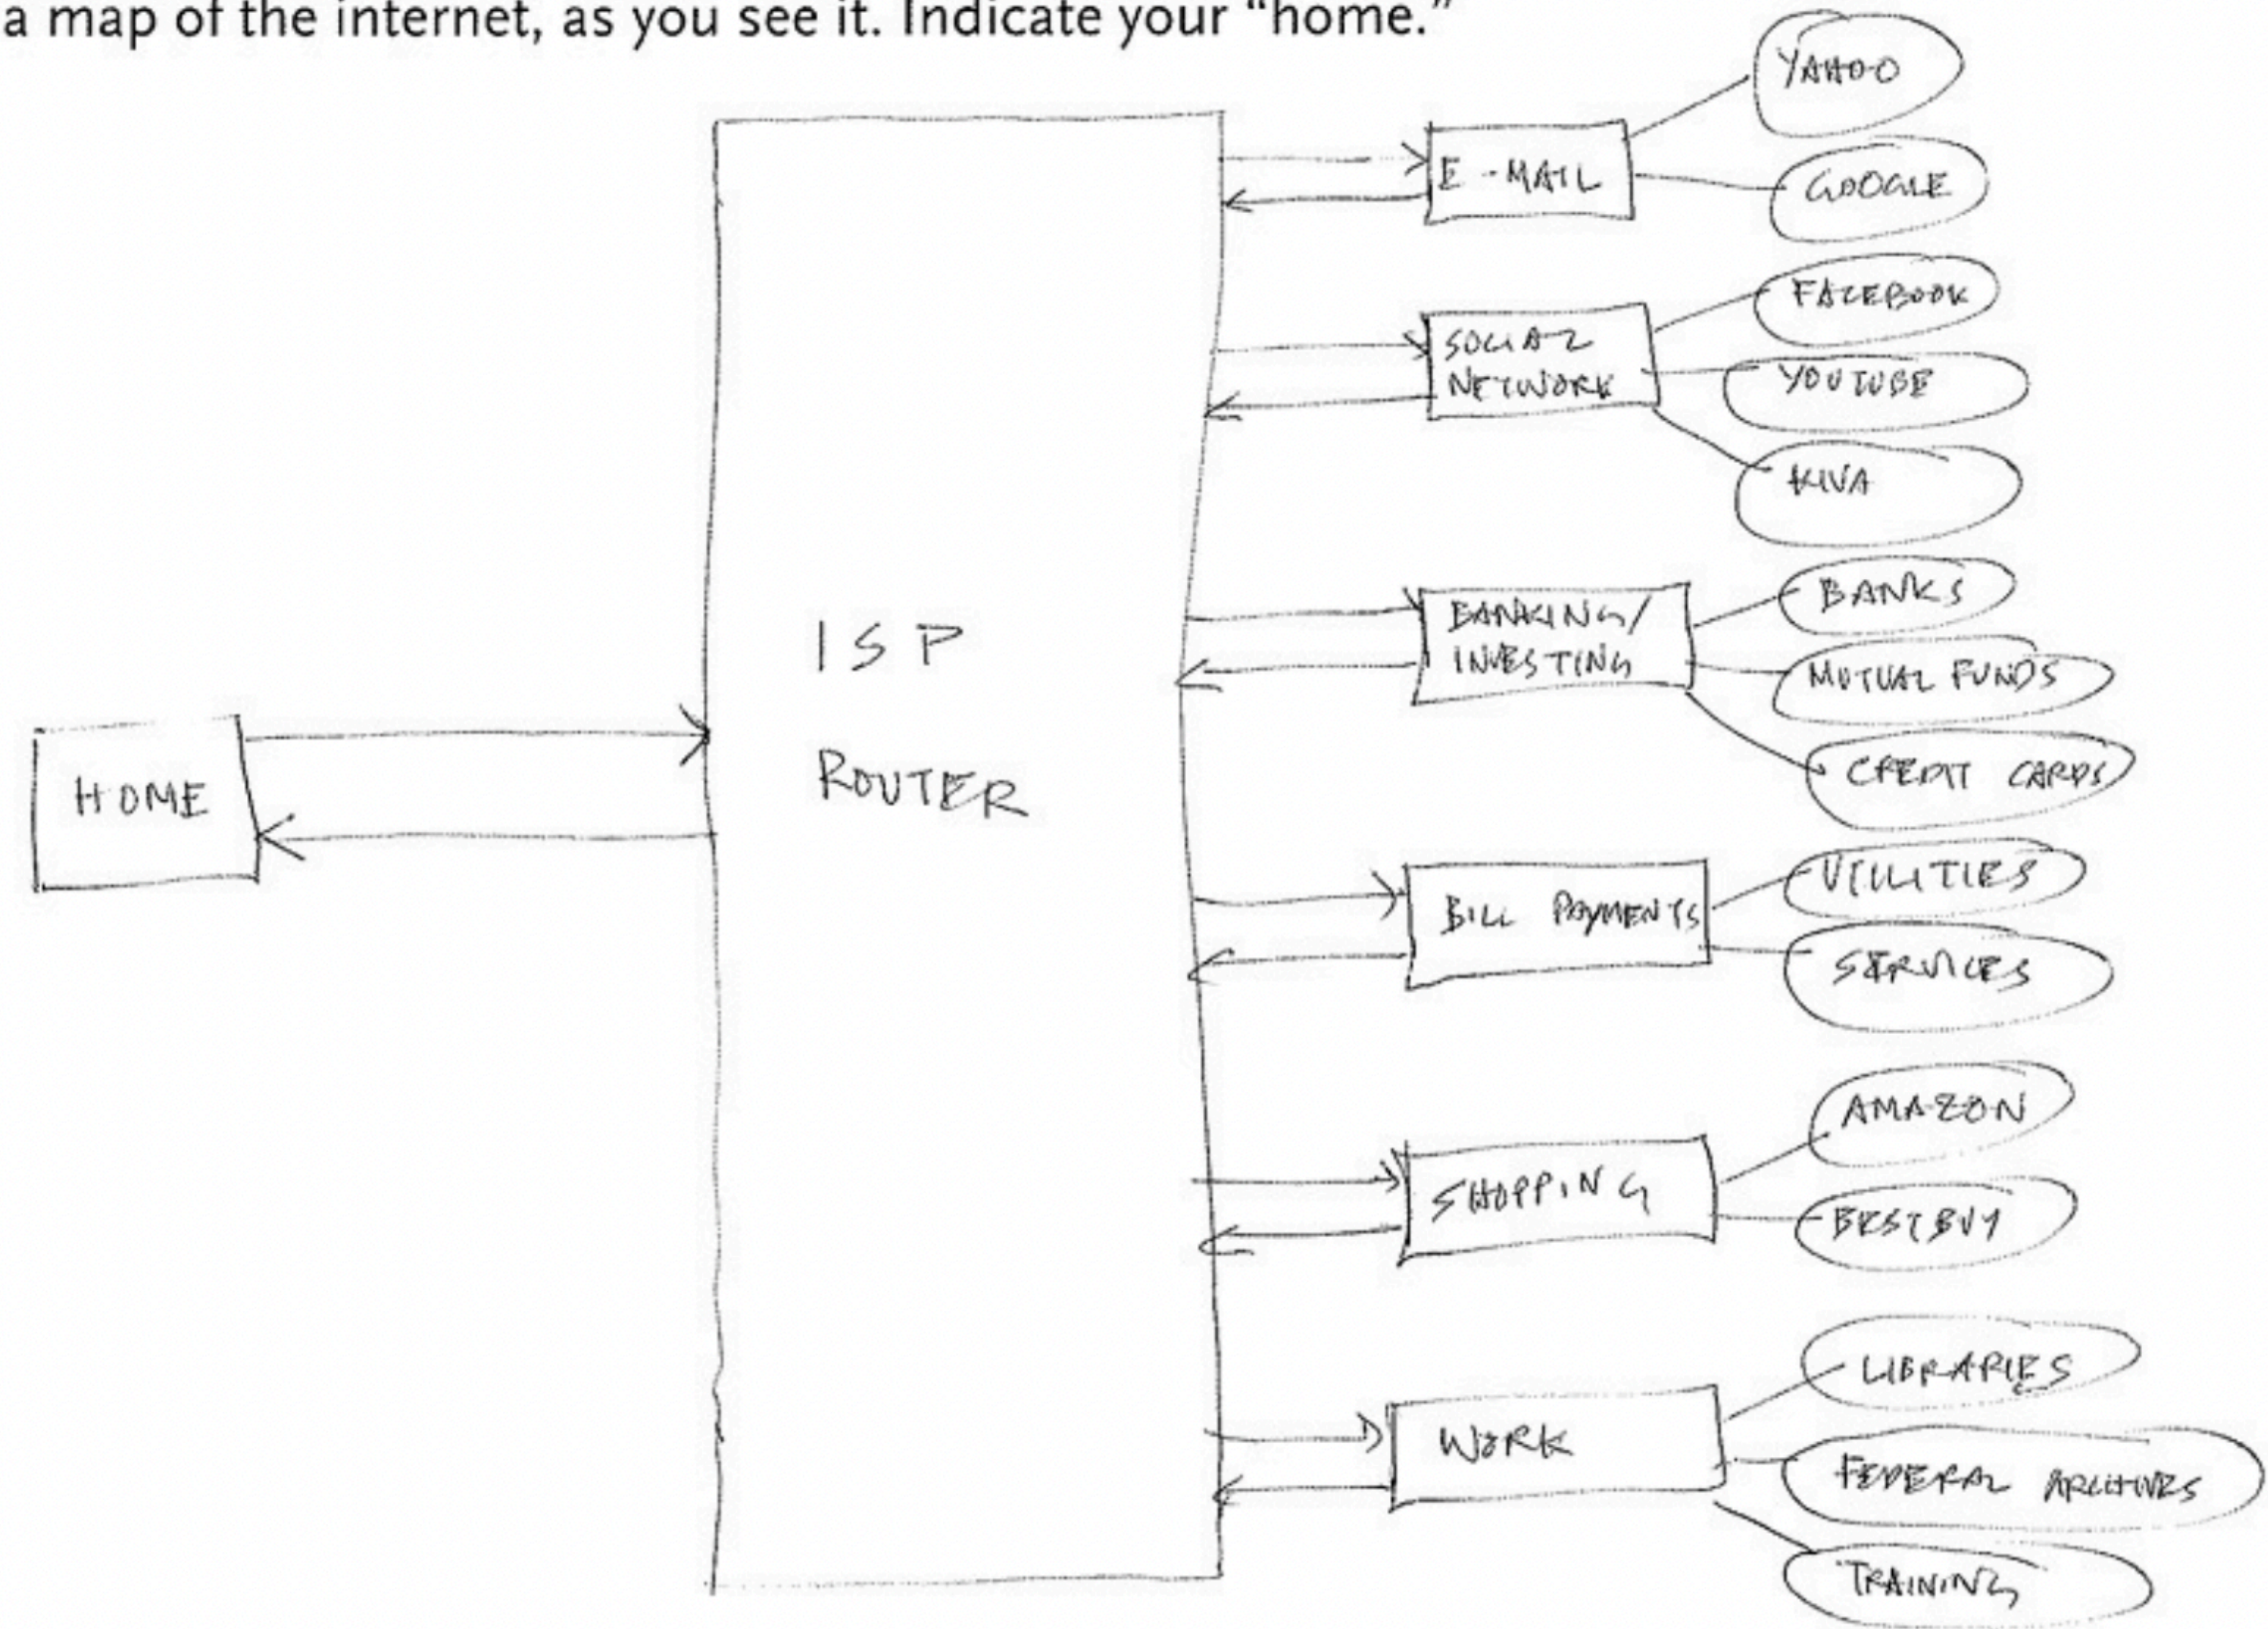
Your age 53 Your occupation Inventor Average number of hours per day on the internet _____

Return to Kevin Kelly (kk@kk.org, or 149 Amapola Ave., Pacifica, CA 94004) For results see www.kk.org.

Kevin Kelly "The Internet Mapping Project"

THE INTERNET MAPPING PROJECT

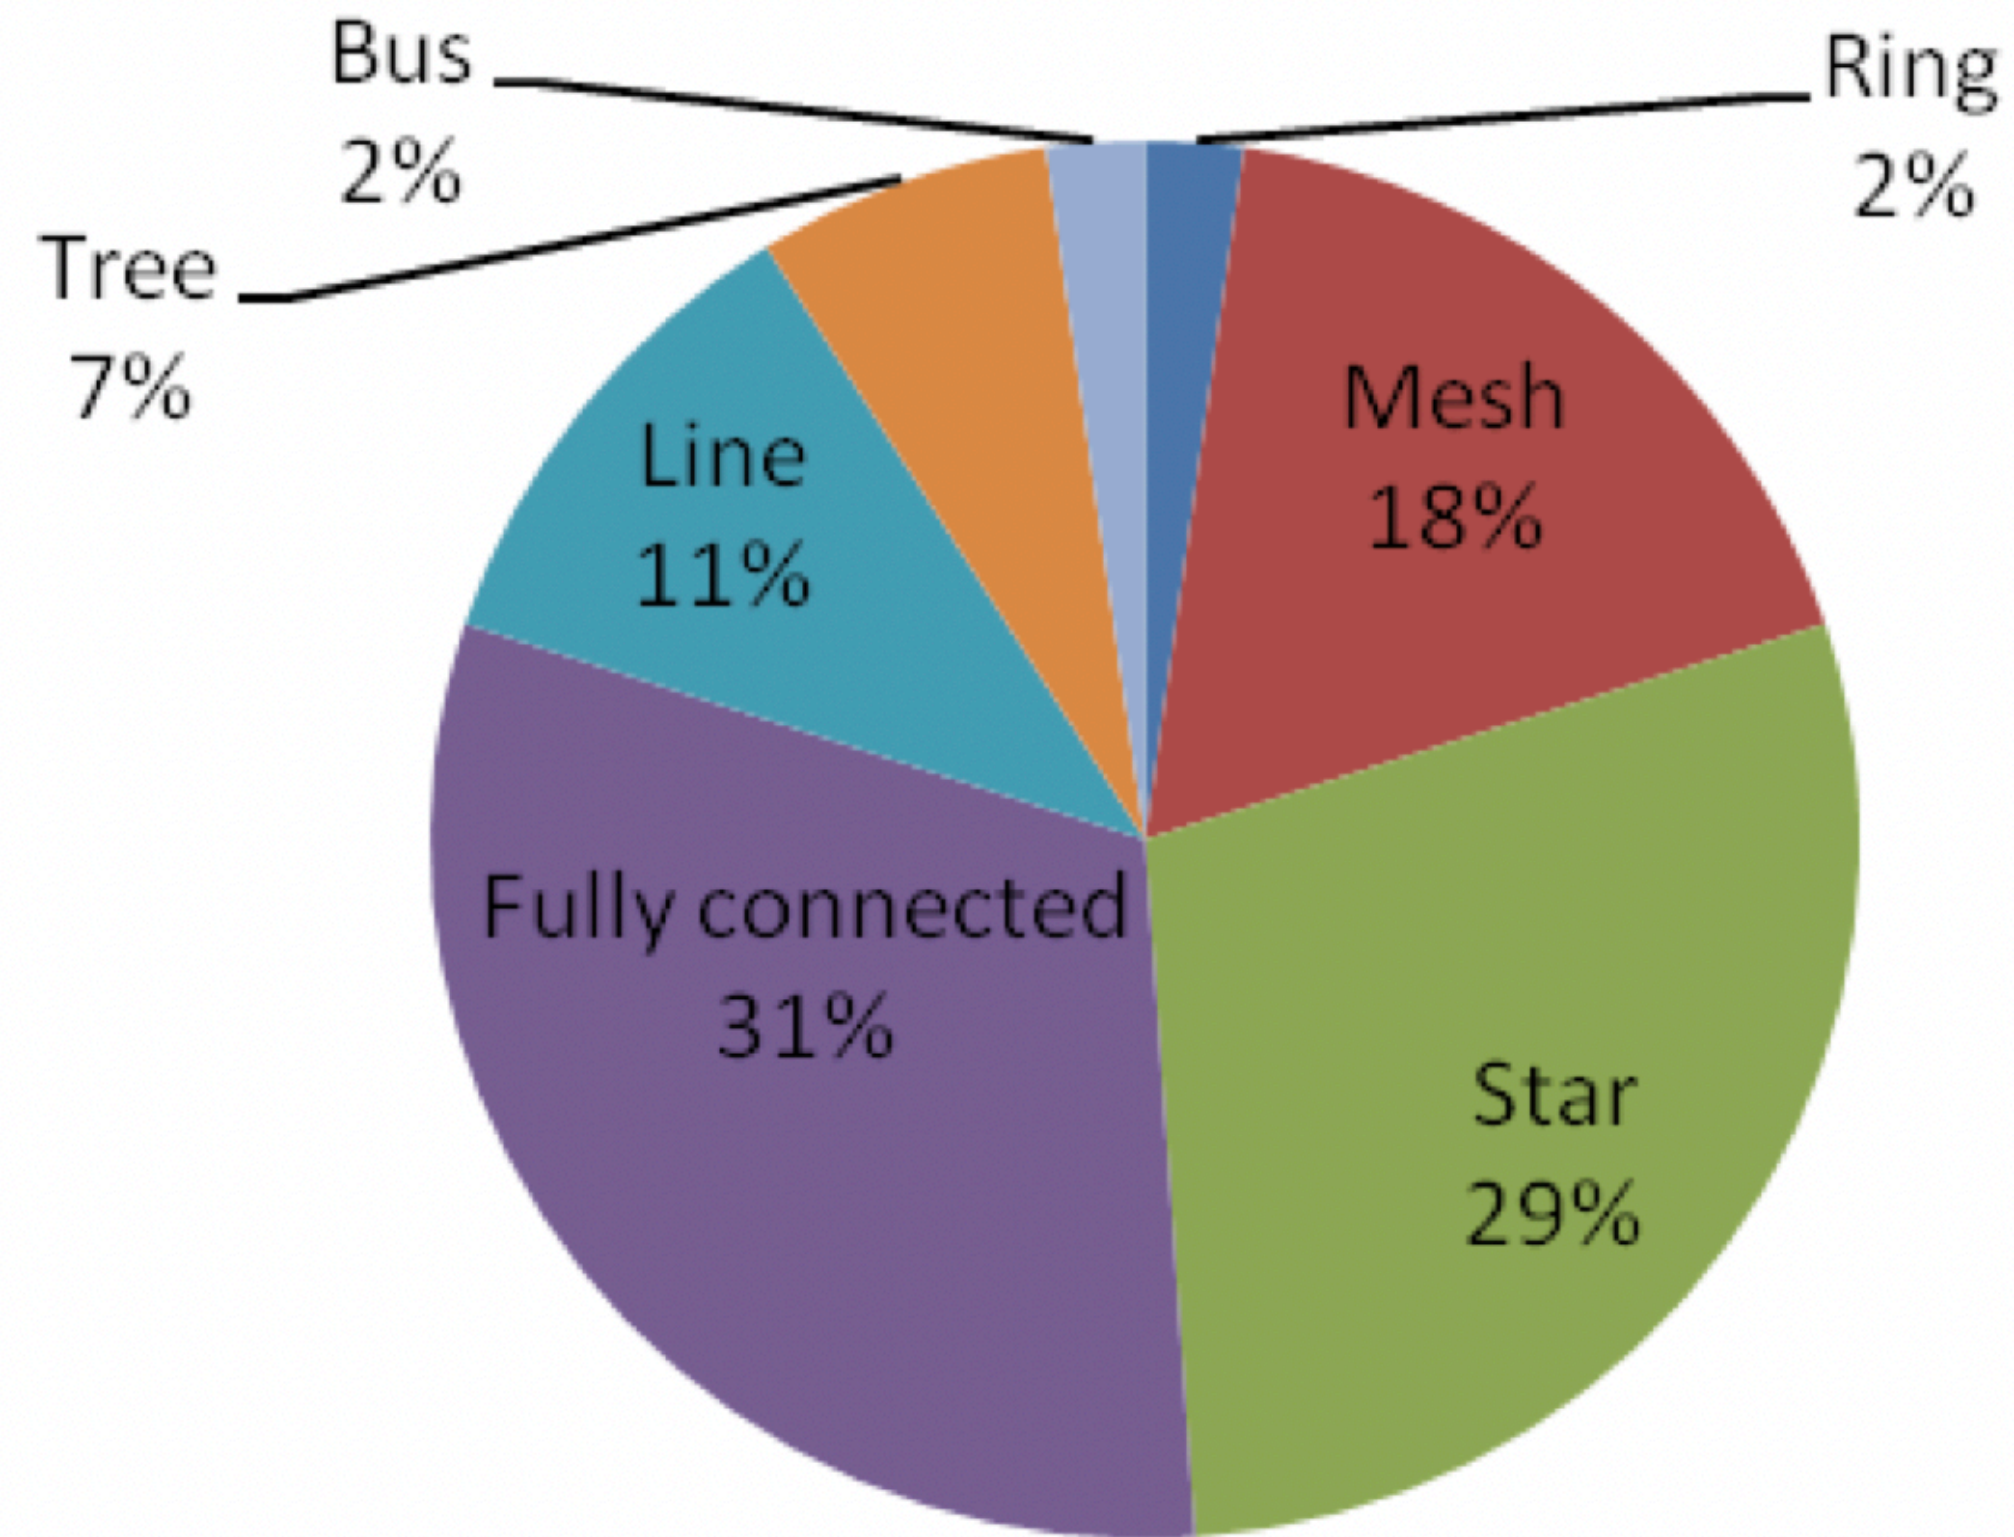
Please draw a map of the internet, as you see it. Indicate your "home."



Your age 39 Your occupation ENGINEER

Average number of hours per day on the internet 12

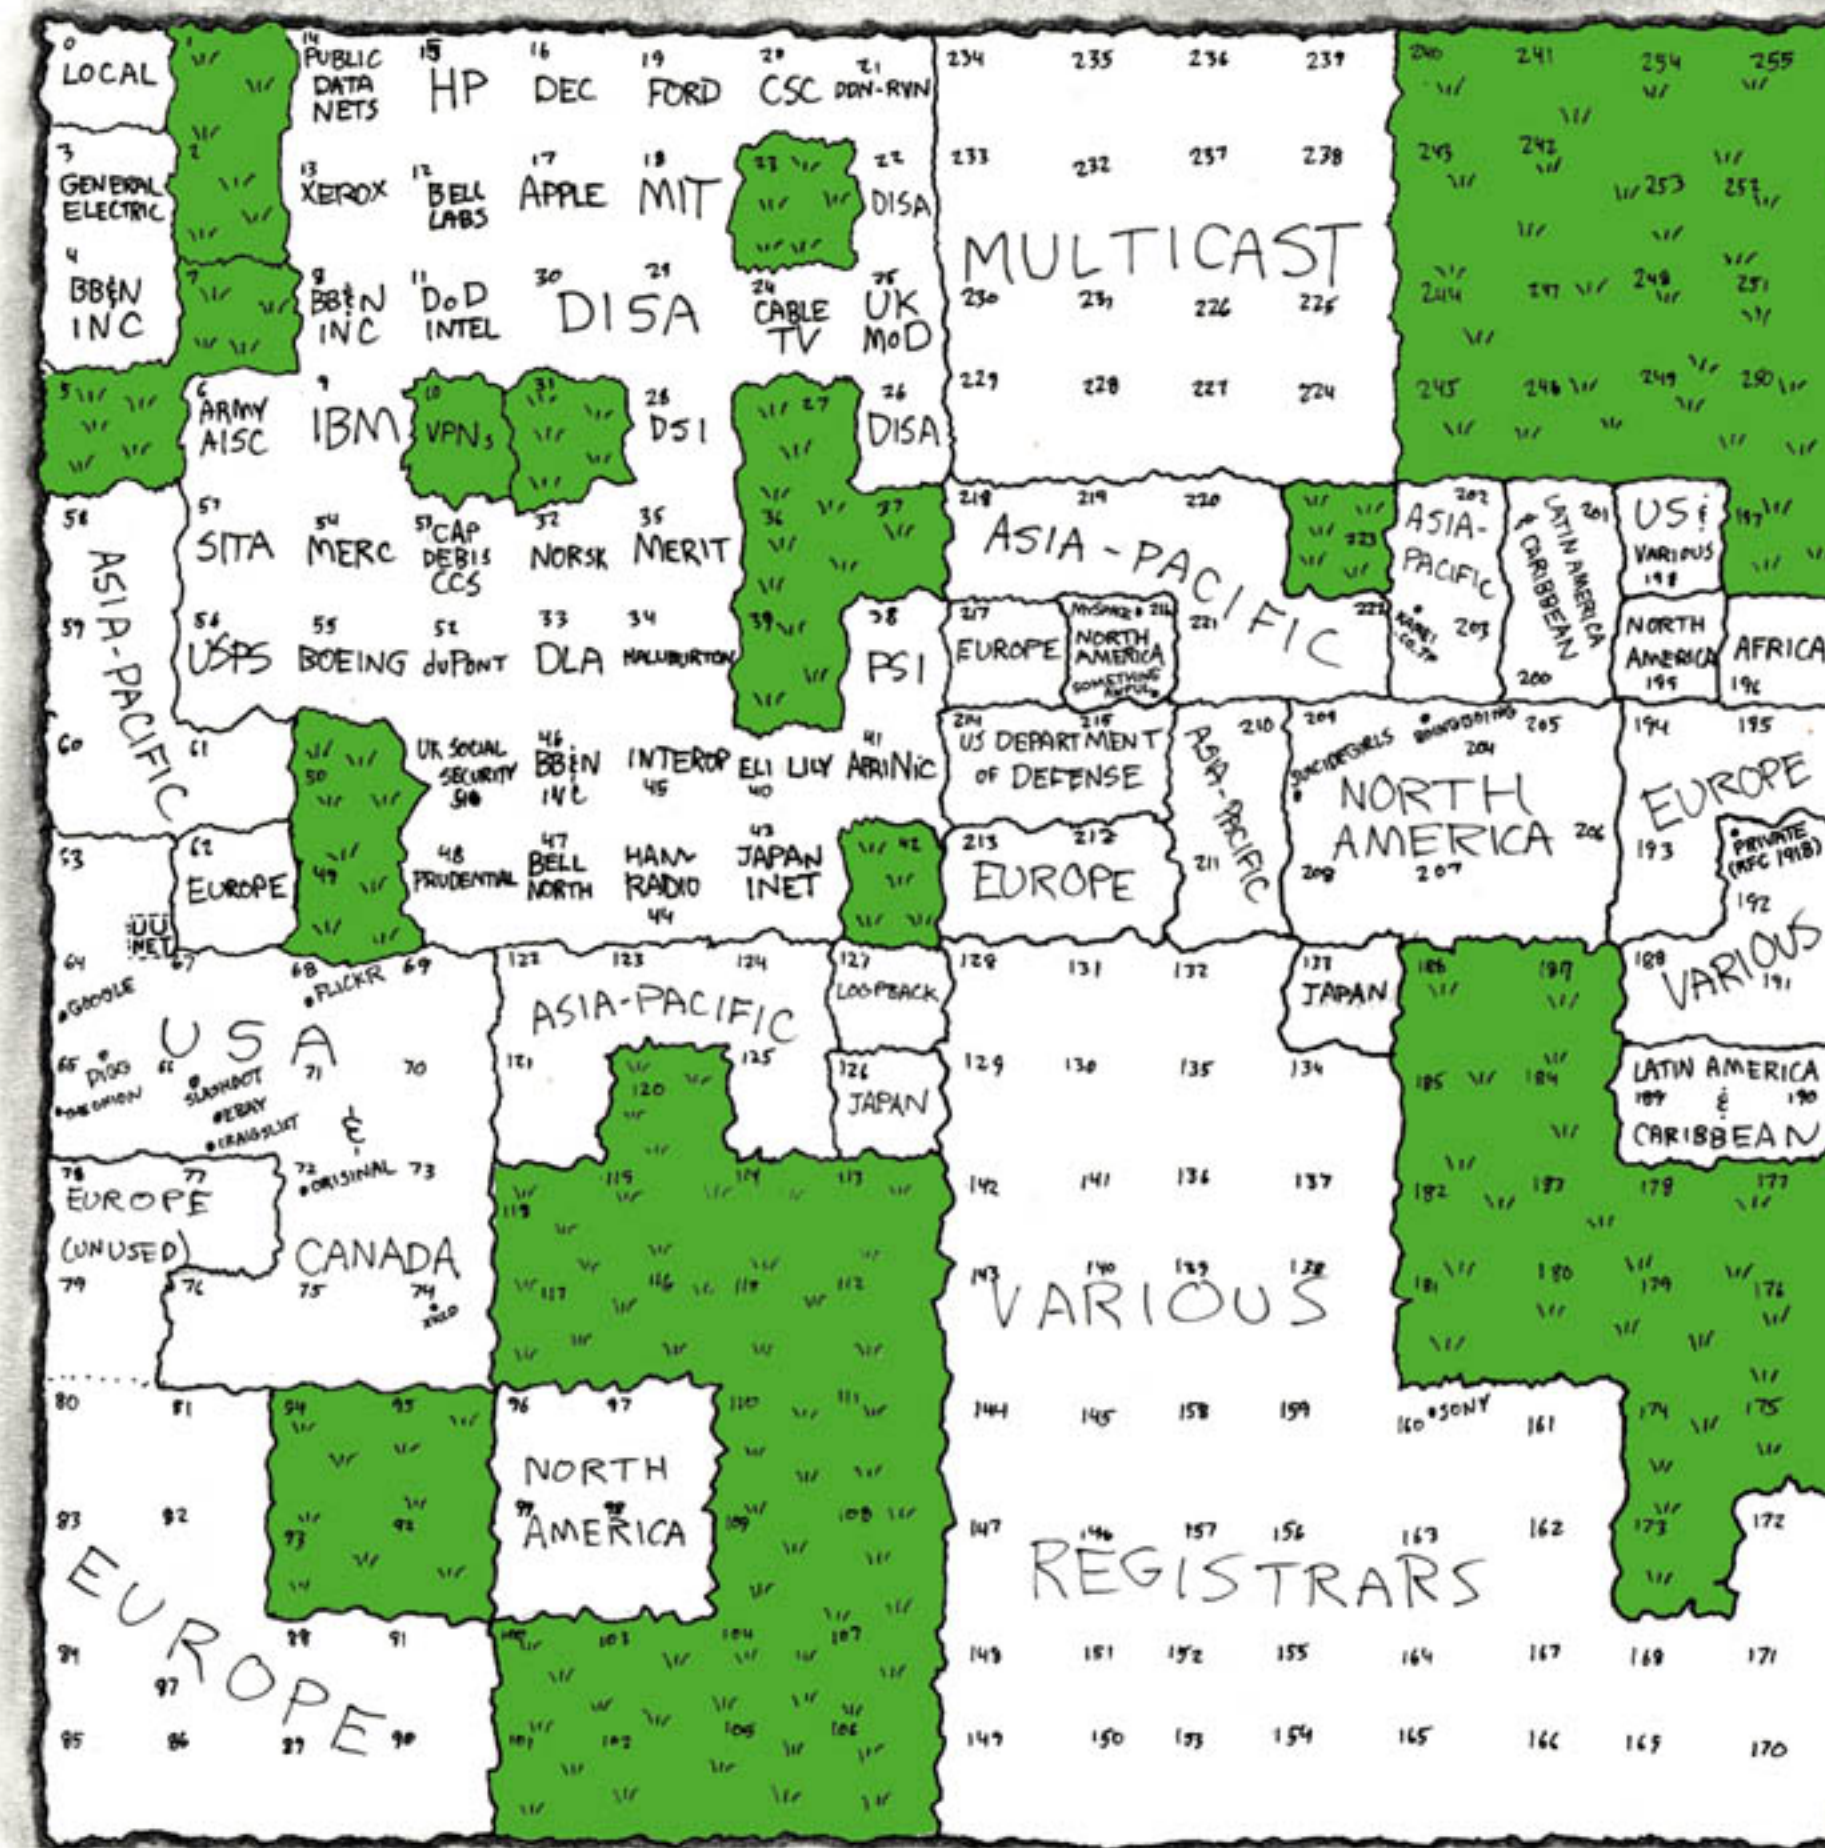
Network Topologies



Graphic 1


MAP OF THE INTERNET


THE IPv4 SPACE, 2006

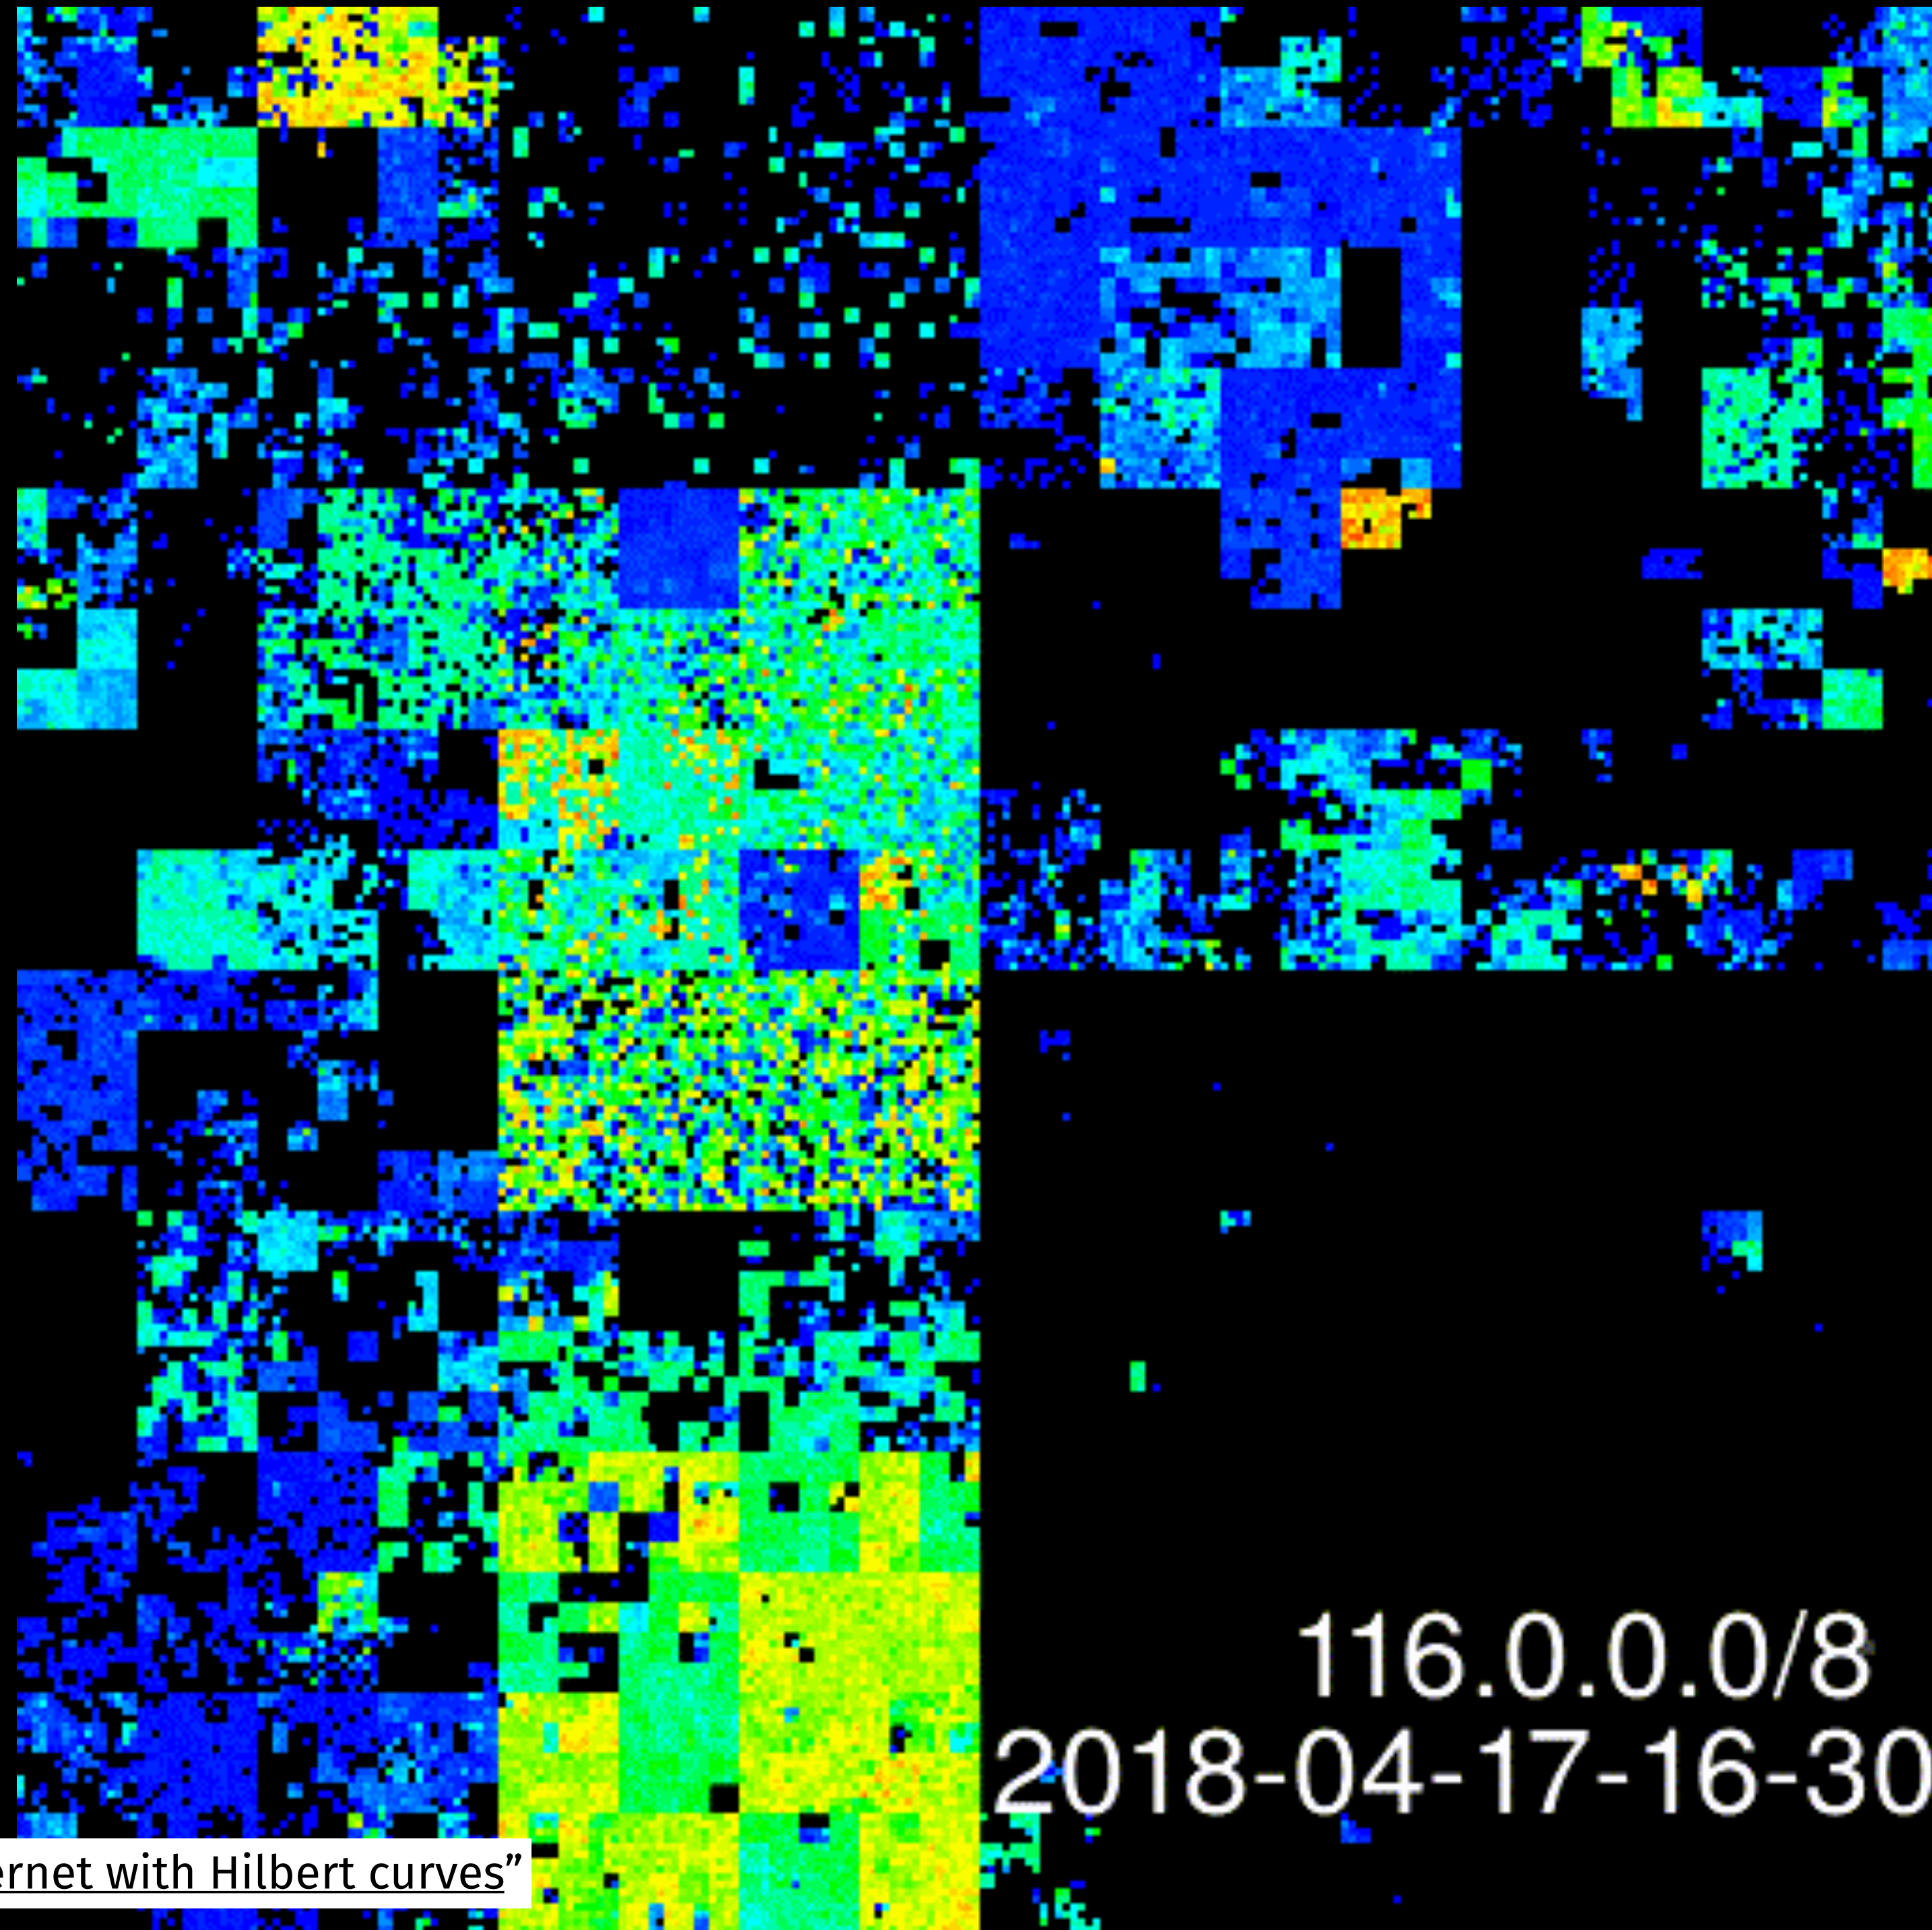


THIS CHART SHOWS THE IP ADDRESS SPACE ON A PLANE USING A FRACTAL MAPPING WHICH PRESERVES GROUPING-- ANY CONSECUTIVE STRING OF IPs WILL TRANSLATE TO A SINGLE COMPACT, CONTIGUOUS REGION ON THE MAP. EACH OF THE 256 NUMBERED BLOCKS REPRESENTS ONE /8 SUBNET (CONTAINING ALL IPs THAT START WITH THAT NUMBER). THE UPPER LEFT SECTION SHOWS THE BLOCKS SOLD DIRECTLY TO CORPORATIONS AND GOVERNMENTS IN THE 1990's BEFORE THE RIRs TOOK OVER ALLOCATION.

0	1	14	15	16	19	→
3	2	13	12	17	18	
4	7	8	11			
5	6	9	10			



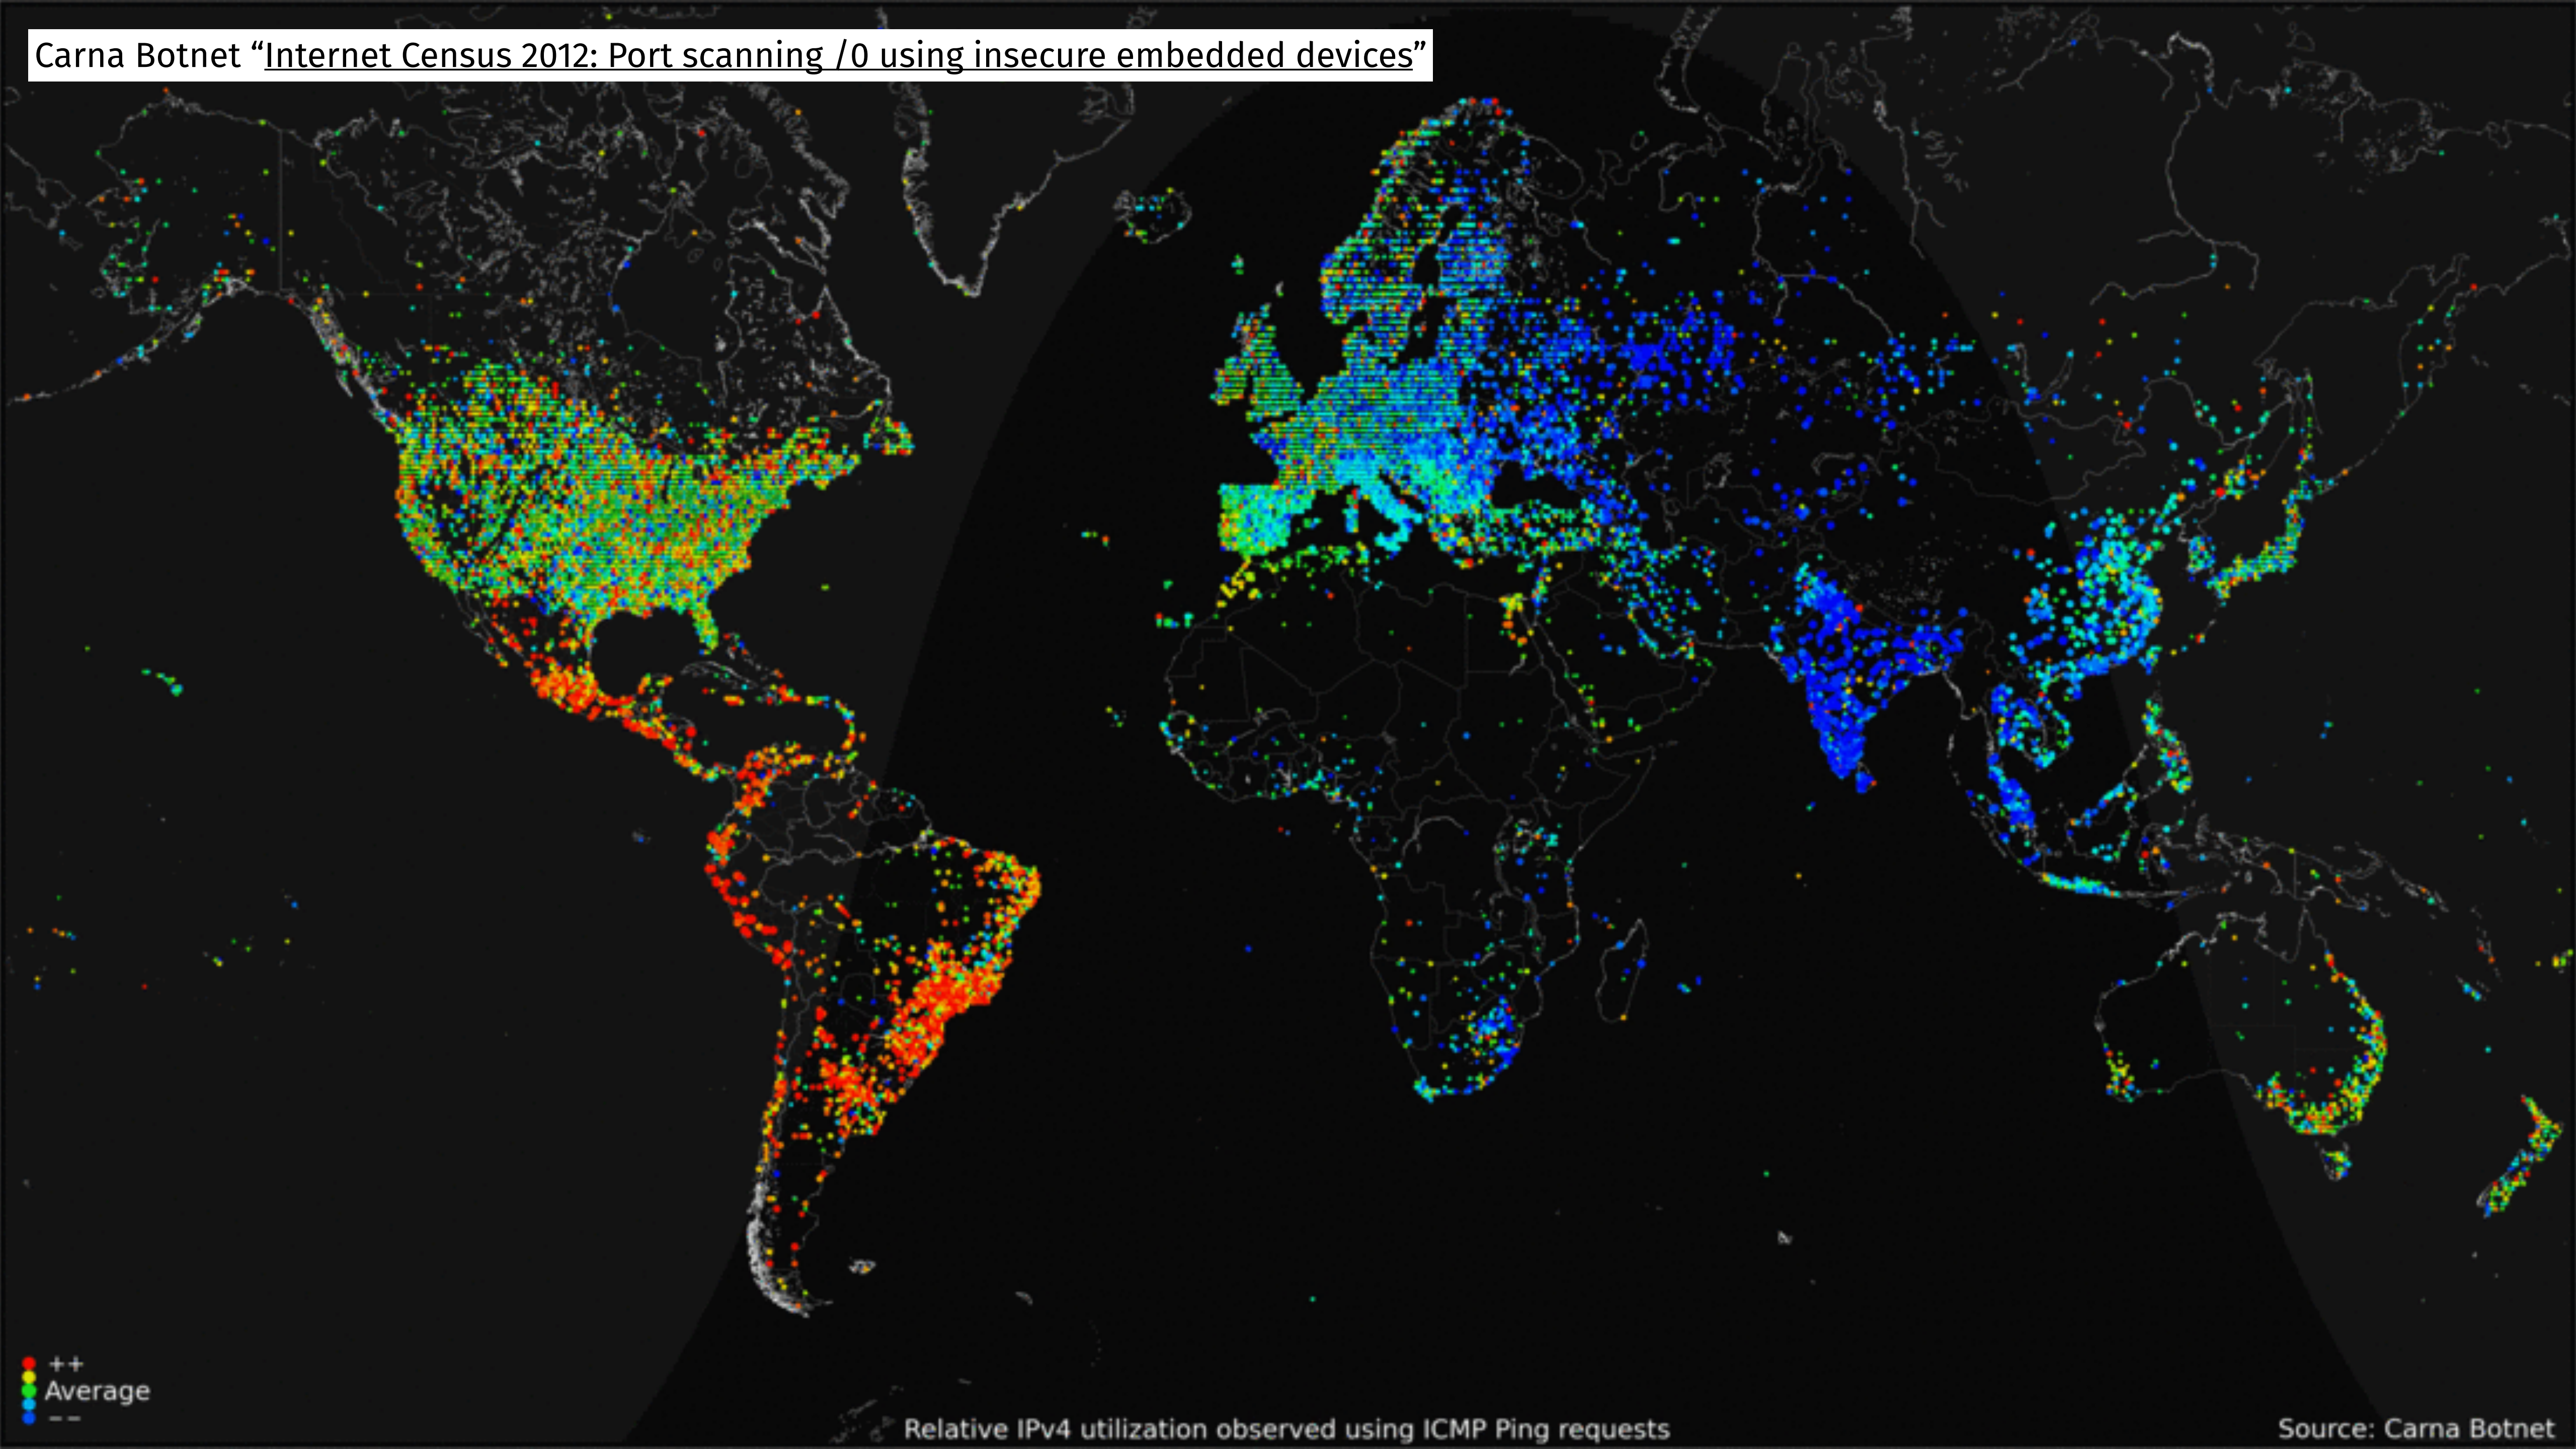
 = UNALLOCATED BLOCK



116.0.0.0/8
2018-04-17-16-30

Ben Cox "Mapping the whole internet with Hilbert curves"

Carna Botnet "Internet Census 2012: Port scanning /0 using insecure embedded devices"



++
Average
--

Relative IPv4 utilization observed using ICMP Ping requests

Source: Carna Botnet

Submarine Cable Map 2021

Submarine Cable 101

HOW MANY KILOMETERS OF CABLE ARE THERE?
We believe there are over 1.3 million kilometers of submarine cables in service globally. Some cables are quite short, like the 131 kilometer CelticConnect cable between Ireland and the United Kingdom. In contrast, others are incredibly long, such as the 20,000 kilometer Asia America Gateway cable.

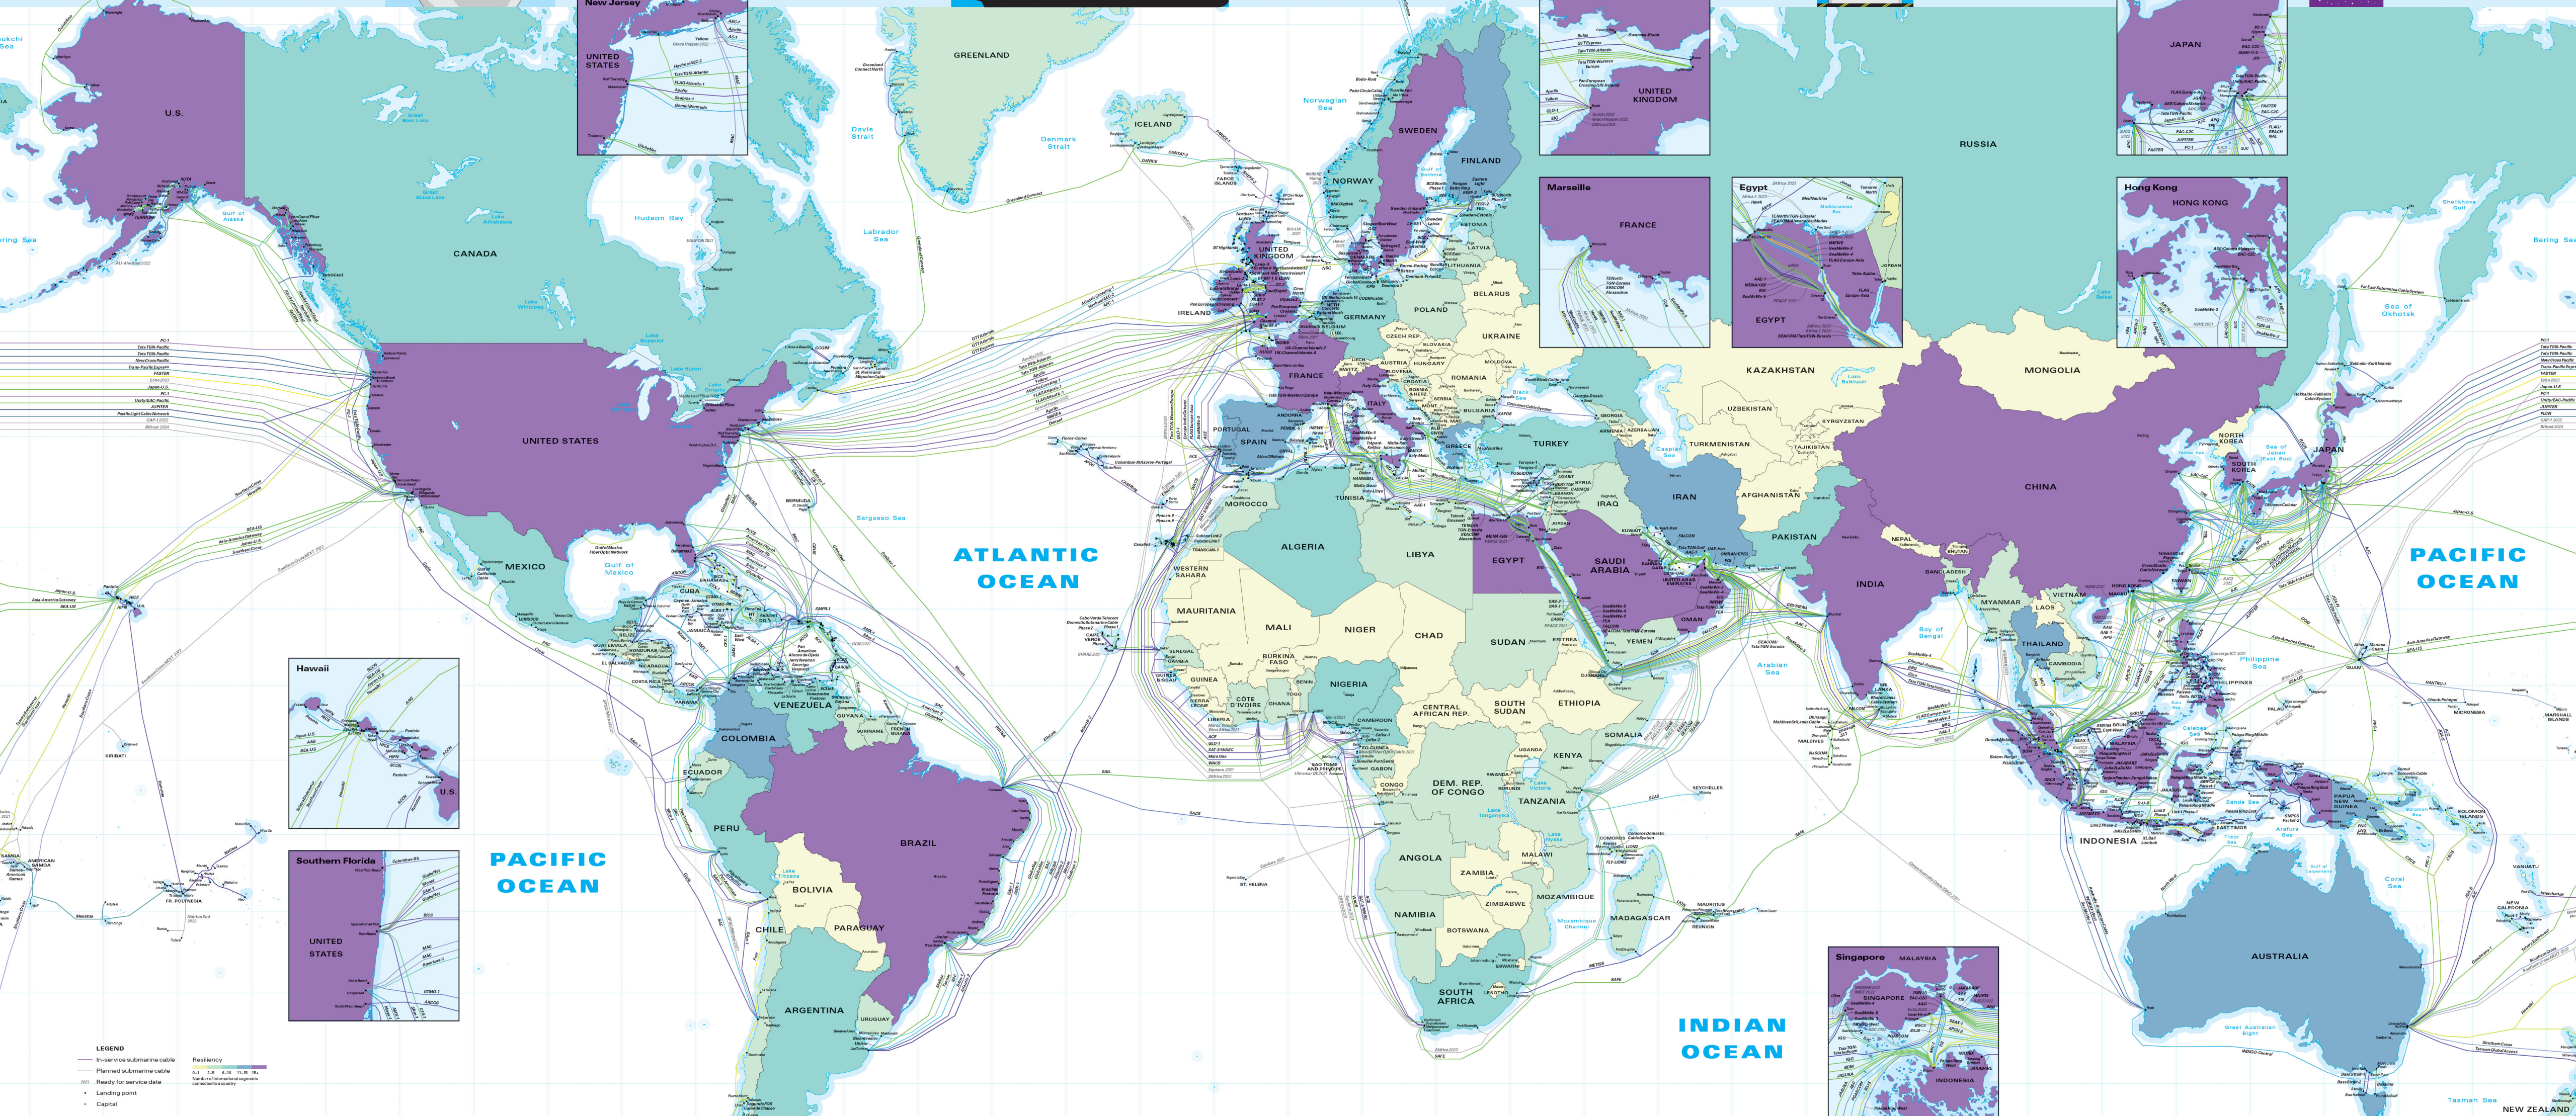
DO THE CABLES ACTUALLY LIE ON THE BOTTOM OF THE OCEAN FLOOR?
Yes, cables go all the way down. Nearer to the shore cables are buried under the seabed for protection, but in the deep sea they are laid directly on the ocean floor.

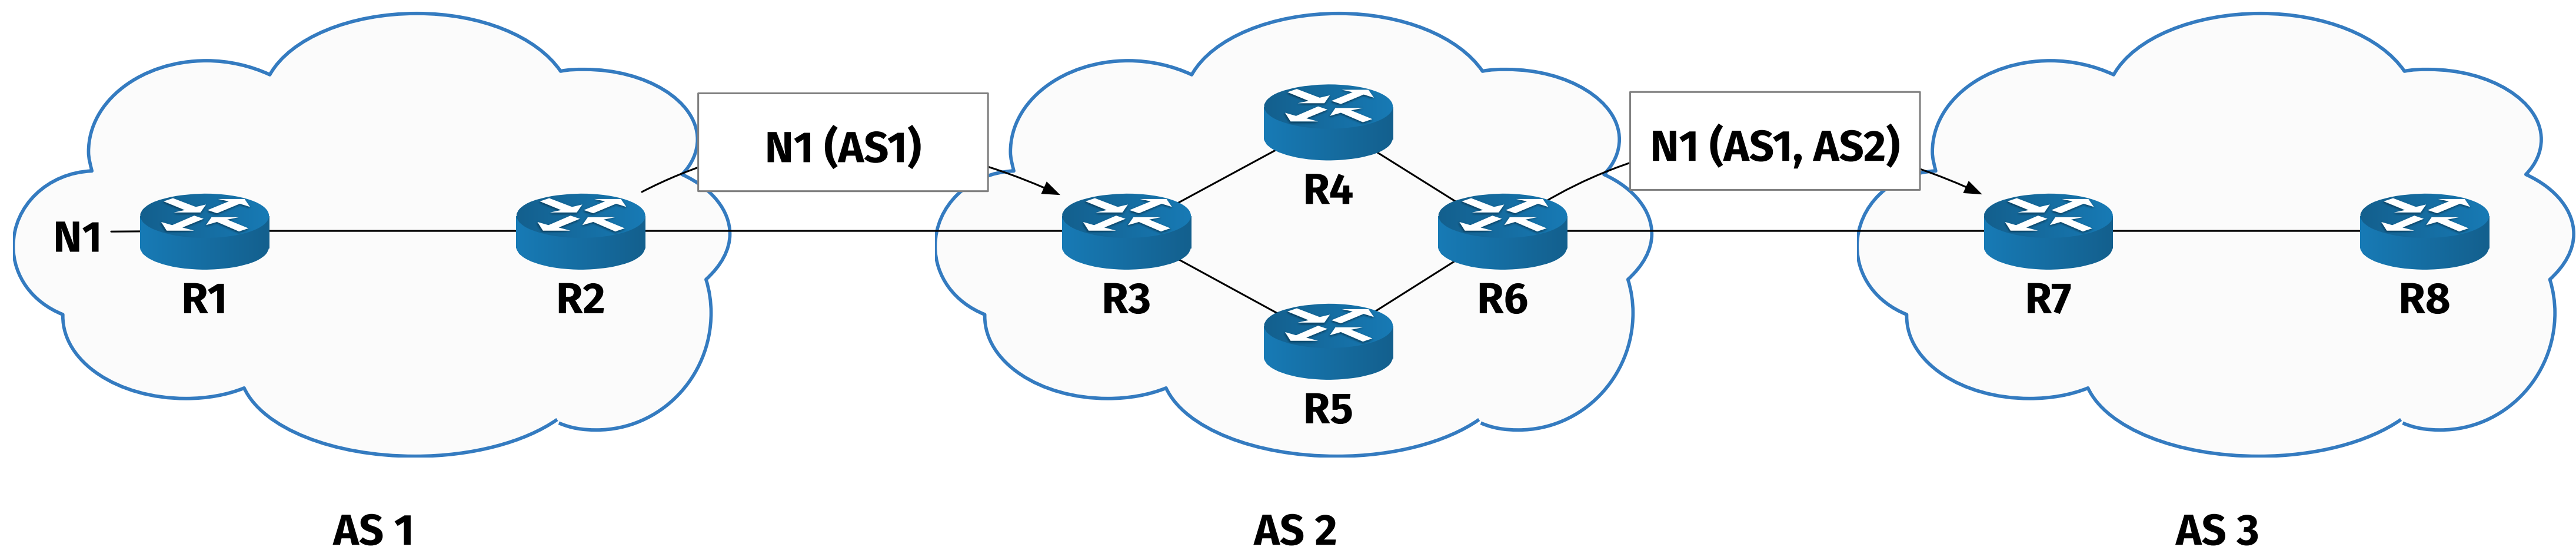
HOW MUCH INFORMATION CAN A CABLE CARRY?
Cable capacities vary a lot. Typically, newer cables are capable of carrying more data than cables laid 15 years ago. For example, GTT Atlantic (2001) is capable of carrying 23 Tbps, while the new MAREA cable is capable of carrying 224 Tbps.

DON'T THESE CABLES EVER BREAK?
Cable faults are common. On average, there are over 100 each year.


I'VE HEARD THAT SHARKS ARE KNOWN FOR BITING CABLES. IS THAT TRUE?
According to data from the International Submarine Cable Protection Committee, fish bites is a category that includes sharks accounted for zero cable faults since 2007. The majority of damage to submarine cables comes from human activity, primarily fishing and anchoring, not sharks.

WHY DON'T COMPANIES USE SATELLITES INSTEAD?
Cables can carry far more data at far less cost than satellites. It's hard to know exactly how much of all international traffic is still carried via satellite, but it's very small. Statistics released by U.S. Federal Communications Commission indicate that satellites account for just 0.37% of all U.S. international capacity.





CAIDA'S IPV4 AS CORE GRAPH JANUARY 2020








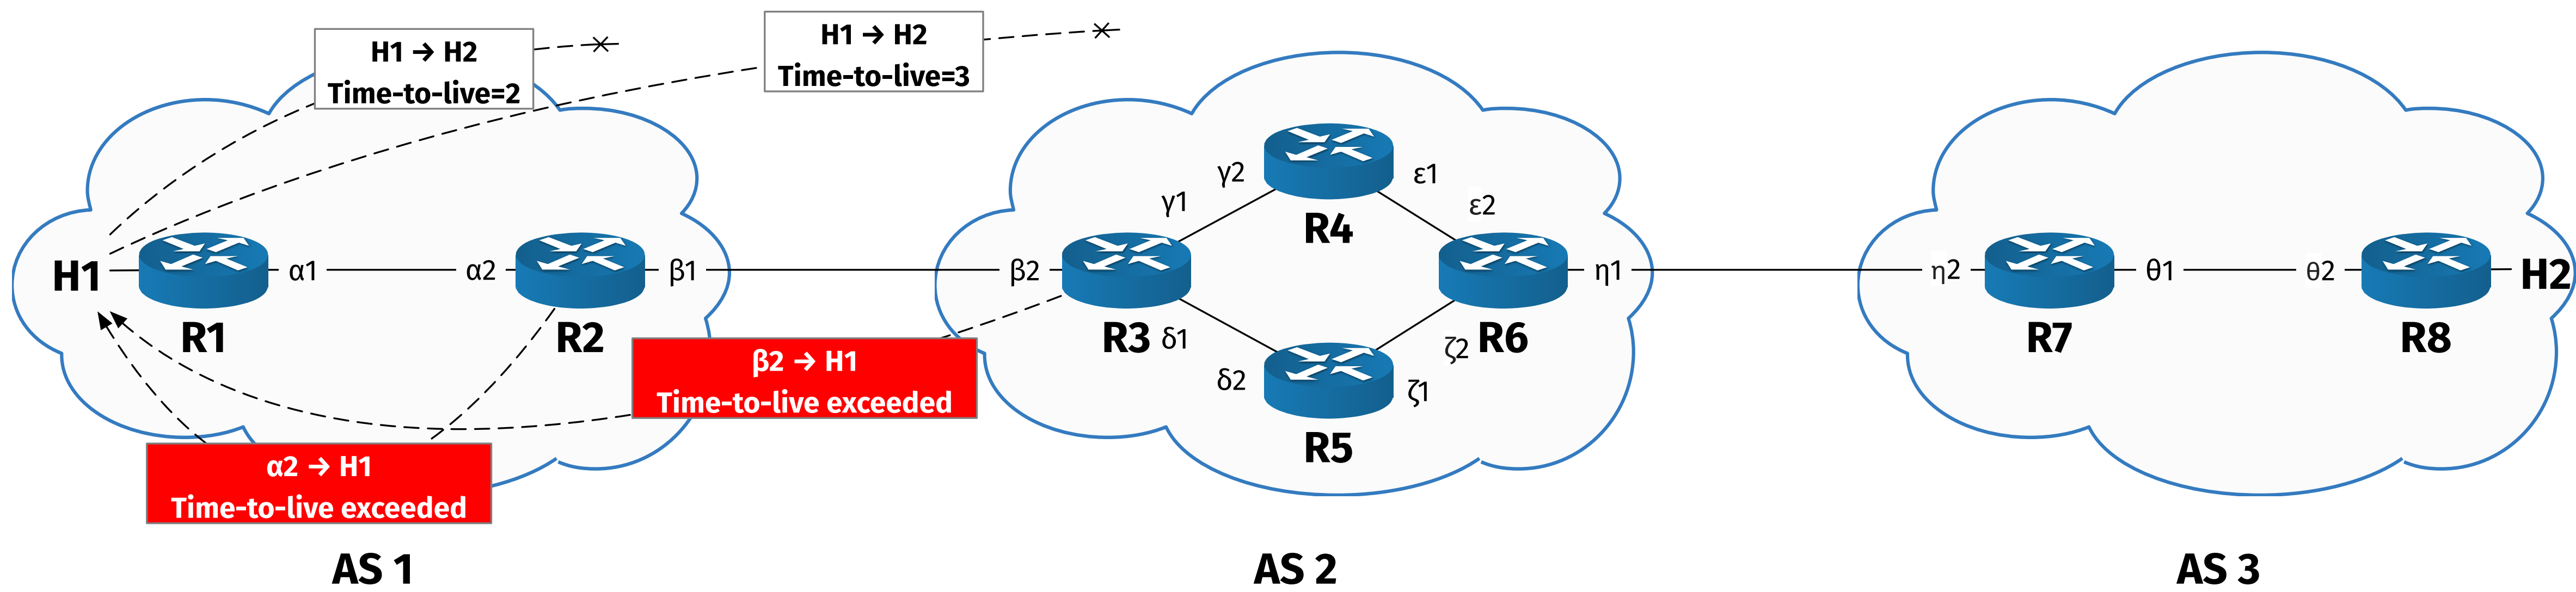
Figure 1. Successive iterations of the layout. Level one vertices (red circles) are placed randomly on a sphere around the root node (black circle). The system is allowed to iterate through time satisfying attractive and repulsive forces until at rest. Level two nodes (blue circles) are placed randomly on spheres directed away from the current layout. Again, the system is allowed to evolve through time till at rest. This process is iterated for the entire graph.



Barrett Lyon / The Opte Project "The Internet 2010"



Barrett Lyon / The Opte Project "The Iranian Internet Shutdown of Nov 2019"






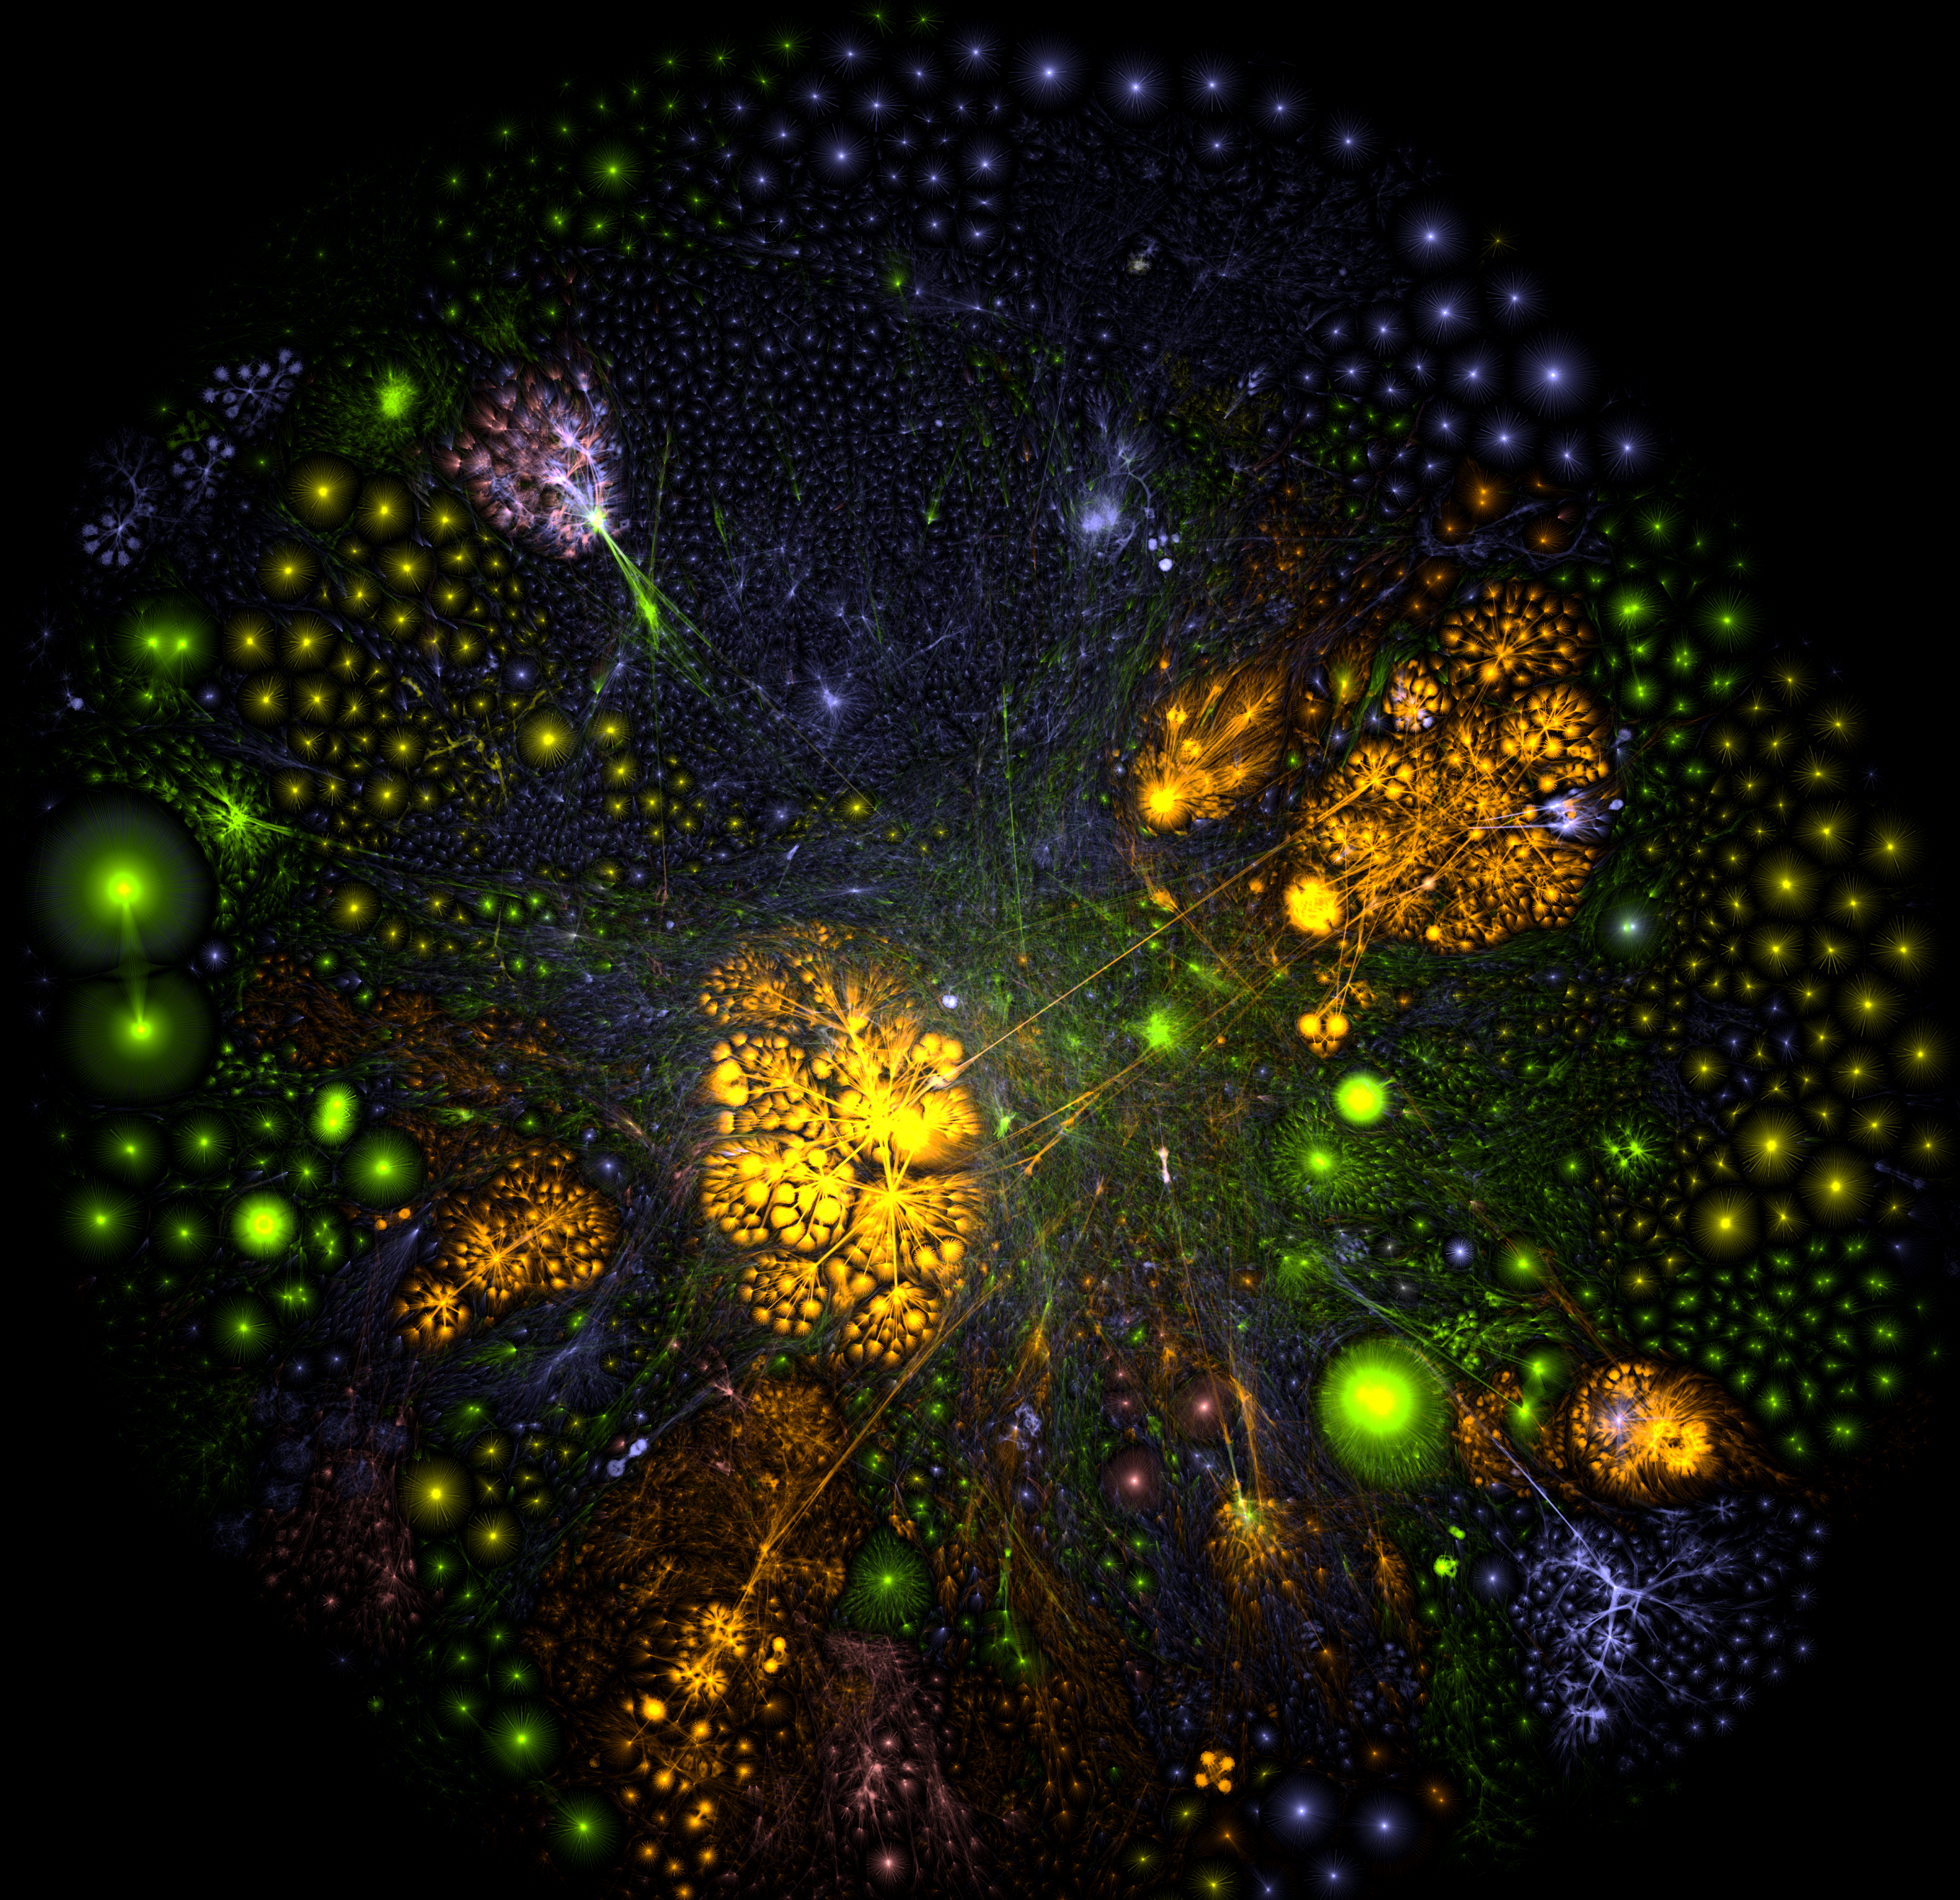


Figure 1: "Peacock-on-the-windshield" map from data taken in September 1998. The circled network is `cw.net`. Color versions of all figures are available at <http://www.cheswick.com/ches/map/mapfigs/>







Barrett Lyon / The Opte Project "The Internet 2003"




Exhaustive multipath measurement from the LIP6 in July 2021



Exhaustive multipath measurement from the LIP6 in February 2022 — Google (AS15169)




Exhaustive multipath measurement from the LIP6 in February 2022 — CHINA169-Backbone (AS4837)



Exhaustive multipath measurement from the LIP6 in February 2022 — Deutsche Telekom AG (AS3320)

Adding interactivity

- Static images are fast to draw
- (Very) high-resolution images are inconvenient (memory requirements)
- How to show millions of edges in real-time?
- How to add information (e.g. AS name or link endpoints) without overloading the screen?
- Already solved for web maps (e.g. Google Maps)
- Reuse free & open-source cartography tools: MapLibre GL and Tippecanoe



Data pipeline

1. Compute links from traceroute measurements
 1. Exclude anonymous nodes
 2. Exclude prefixes with *reply amplification* or *routing loops*
 3. Exclude *unlikely* links
2. Extract the largest connected component
3. Compute the LGL layout
4. Augment the data with AS, RIR and facility information
5. Generate the tile set by dropping densest features at each zoom level

Demo

Future works

- Animate interactive graphs over time
- Use RTT (round-trip time) information
 - Delay-based layout?
 - Color edges by RTT from source
- Geo-based layout
- Router-level graph
- Include data from multiple vantage points
- Include other data sources (CAIDA Ark, M-Lab, RIPE Atlas)
- Highlight paths from the source

Thanks!