Curvature-based Analysis of Network Connectivity in Private Backbone Infrastructures

Accepted at **ACM Sigmetrics 22'**

Loqman Salamatian (Columbia University)
Scott Anderson (University of Wisconsin-Madison)
Joshua Mathews (University of Wisconsin-Madison)
Paul Barford (University of Wisconsin-Madison)
Walter Willinger (Niksun Inc.)
Mark Crovella (Boston University)

Cloud Providers are investing billions $ in their backbones

TECH | KEYWORDS: CHRISTOPHER MIMS

Google, Amazon, Meta and Microsoft Weave a Fiber-Optic Web of Power

The four tech giants increasingly dominate the internet's critical cable infrastructure

COMPANIES > AMAZON

AWS Will Be Google and Facebook's Neighbor on the New US-Europe Submarine Cable

TELECOM

Google to build 3 submarine cables to battle Microsoft, AWS for cloud customers

By Sean Buckley • Jan 16, 2018 10:19am

Microsoft, Facebook and Telxius complete the highest-capacity subsea cable to cross the Atlantic

People and organizations rely on global networks every day to provide access to internet and cloud technology. Those systems enable tasks both simple and complex, from uploading photos and searching webpages to conducting banking transactions and managing air-travel logistics. Most people are aware of their daily dependency on the internet, but few understand the critical role played by the subsea networks spanning the planet in providing that connectivity.

Press release

Share

AWS Announces Global Expansion of AWS Local Zones

February 16, 2022 at 6:00 PM EST

Building off the successful launch of AWS Local Zones in 16 US cities, over 30 new AWS Local Zones will deliver single-digit millisecond latency performance at the edge of the cloud to hundreds of millions of people worldwide

Netflix, Couchbase, Supercell, and FOX Corporation among thousands of customers using AWS Local Zones

Cloud Providers are investing billions $ in their backbones

Every cloud deployment is unique

Differences in infrastructure arises in AWS, Azure and GCP through:

- constraints created by geography
- historical infrastructure
- unique strategies and different business goals

Cloud Providers are investing billions $ in their backbones

Every cloud deployment is unique

How to map the topology: traditional approaches

- Apply bag of heuristics to infer router-level interconnectivity from traceroute measurements
- Geolocate routers with rDNS hints and generate physical topology.

Cloud Providers are investing billions $ in their backbones

Every cloud deployment is unique

How to map the topology: traditional approaches

- Apply bag of heuristics to infer router-level interconnectivity from traceroute measurements
- Geolocate routers with rDNS hints and generate physical topology.

Already described in **Rocketfuel (Spring *et al.* *2002*)** more than 20 years ago!

Cloud Providers are investing billions $ in their backbones

Every cloud deployment is unique

How to map the topology: traditional approaches

- Apply bag of heuristics to infer router-level interconnectivity from traceroute measurements
- Geolocate routers with rDNS hints and generate physical topology.

Already described in **Rocketfuel (Spring *et al.* *2002)*** more than 20 years ago!

Why does this not work for CP?

Why does this not work for CP?

Historical challenges

Modern-day challenges

Dependency on the willingness of cloud providers and ISPs to collaborate

Goals

For the academic researcher

- Getting a clearer view of hardly observable topologies that carry a huge portion of the traffic.
- Comparing and contrasting network properties of the different cloud providers.
- Perspective on a post-traceroute world.

For the network engineer at Google

- Discovering what part of my topology is visible from an outsider perspective
- Monitoring your network with a new visualization of tools

Starting point:

Imagine that we have access to devices to schedule measurements and we know where they are (physically) located

Starting point:

Imagine that we have access to devices to schedule measurements and we know where they are (physically) located

Problem:

We are interested in discovering how they are connected to each other

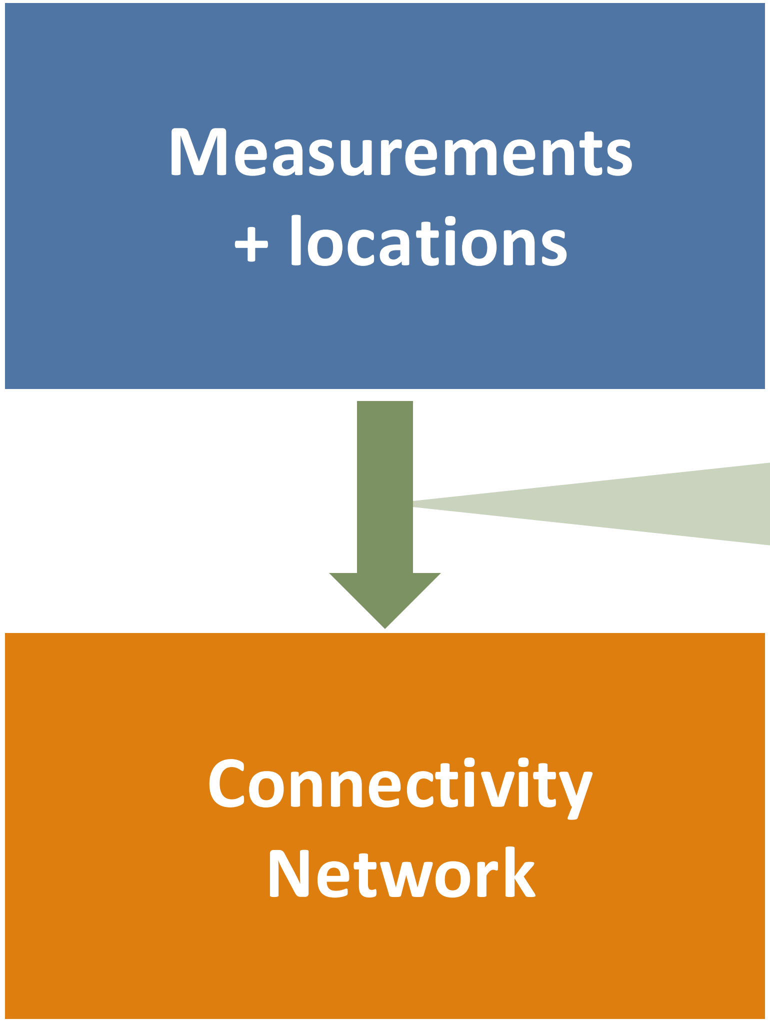

Two key ideas:

1. Transferring complexity from measurement to modeling: cast the problem as an instance of manifold learning.

2. Geography as part of the topological representation of the Internet: constrain the manifold to be on the world map.

Onto the Methodology

1) Measuring the topology

2) Creating graph through filtrations

3) Computing Ricci curvature

4) Analysis of the graphs

5) Merging everything into a single manifold visual

What is residual latency?

Deterministic latency

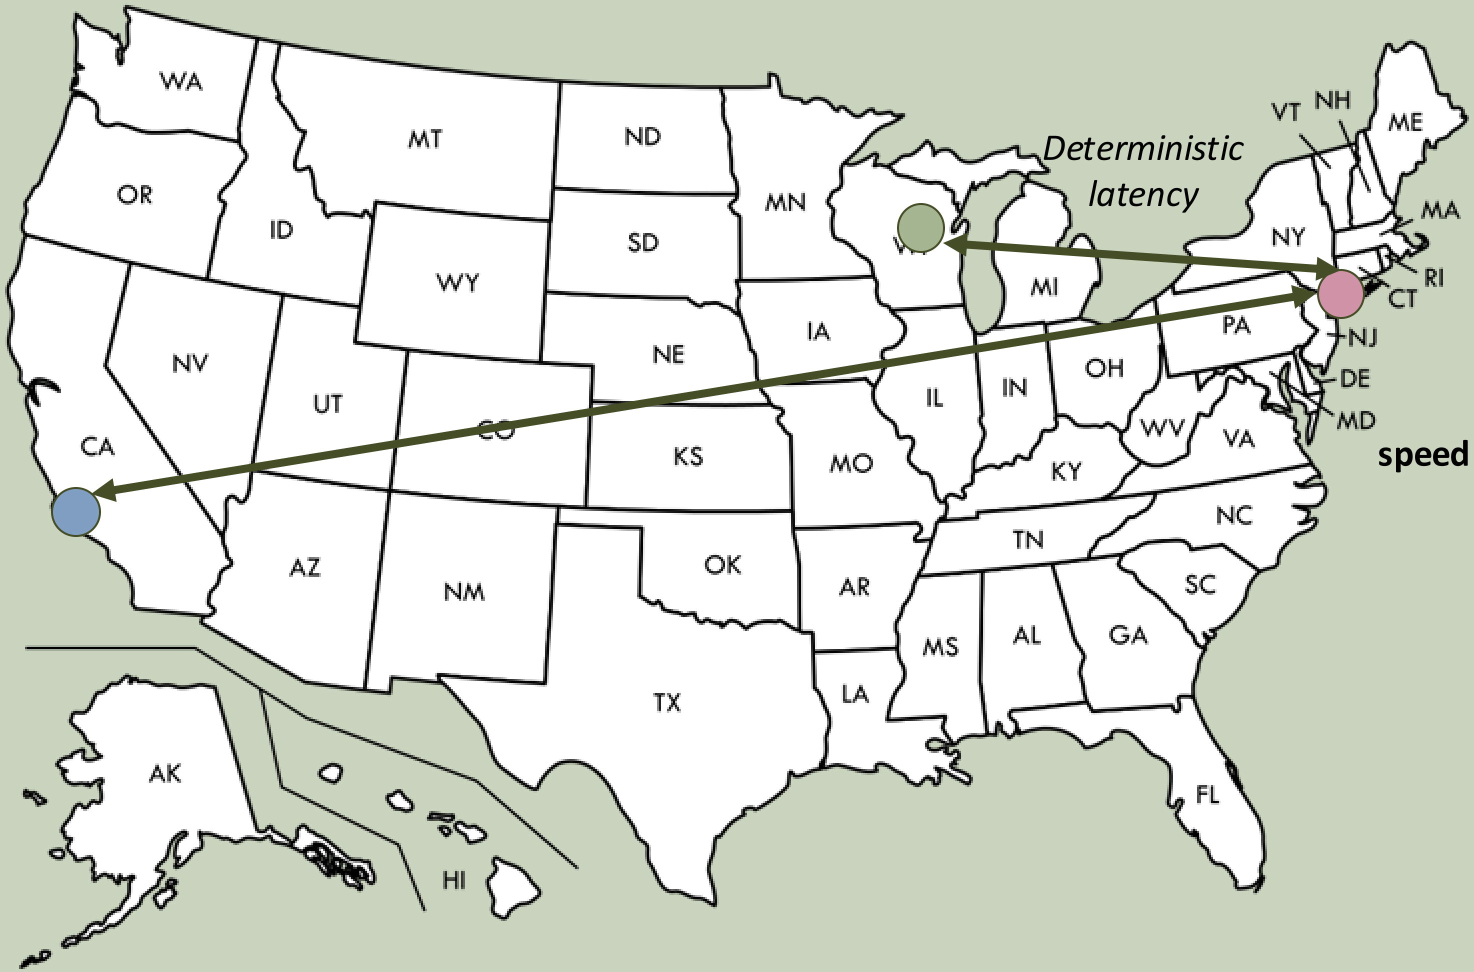

speed in fiber = ⅔ * speed of light in vacuum

COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK

What is residual latency?

We schedule ping measurements

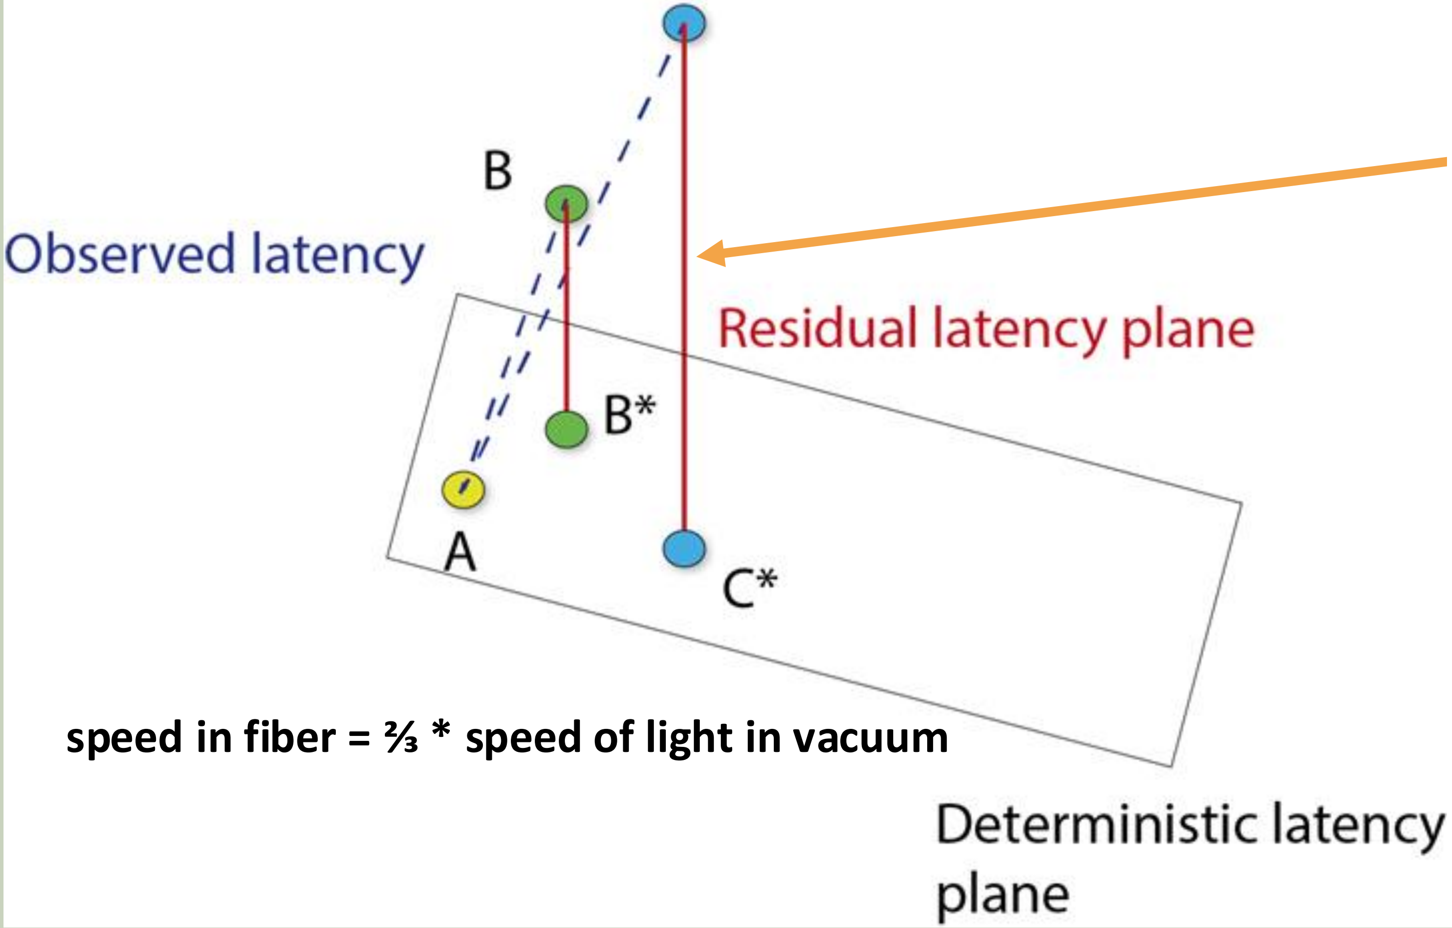

Queueing delays, transmission delays, traffic engineering, ***absence of straight path…***

speed in fiber = ⅔ * speed of light in vacuum

What is residual latency?

We schedule ping measurements : alternate view

Sampling through a long period of time

$$X(A,B) = min_{t \in \mathbb{R}} (X_t)(A,B)$$

Real shortest latency

minimize queueing delays

minimize transmission delays

What is residual latency?

We schedule ping measurements : alternate view

Sampling through a long period of time

Building a *distance* matrix from the residual latency

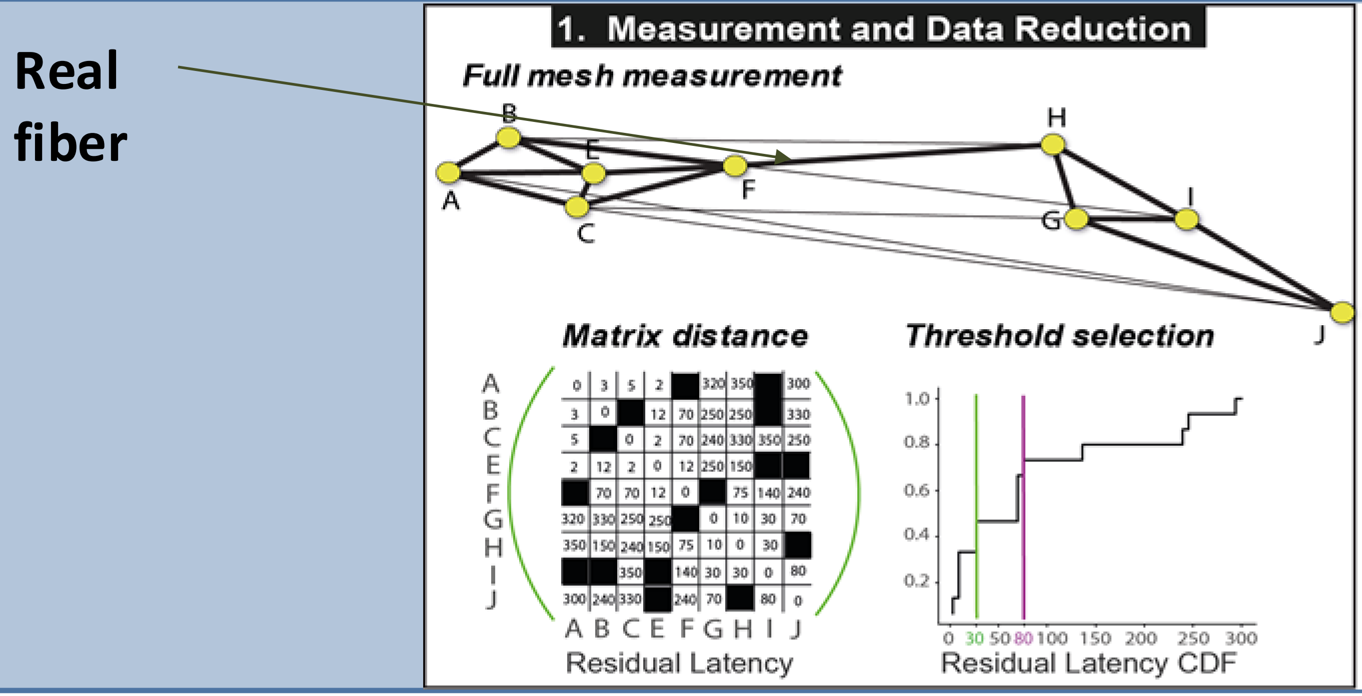

COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Building graphs

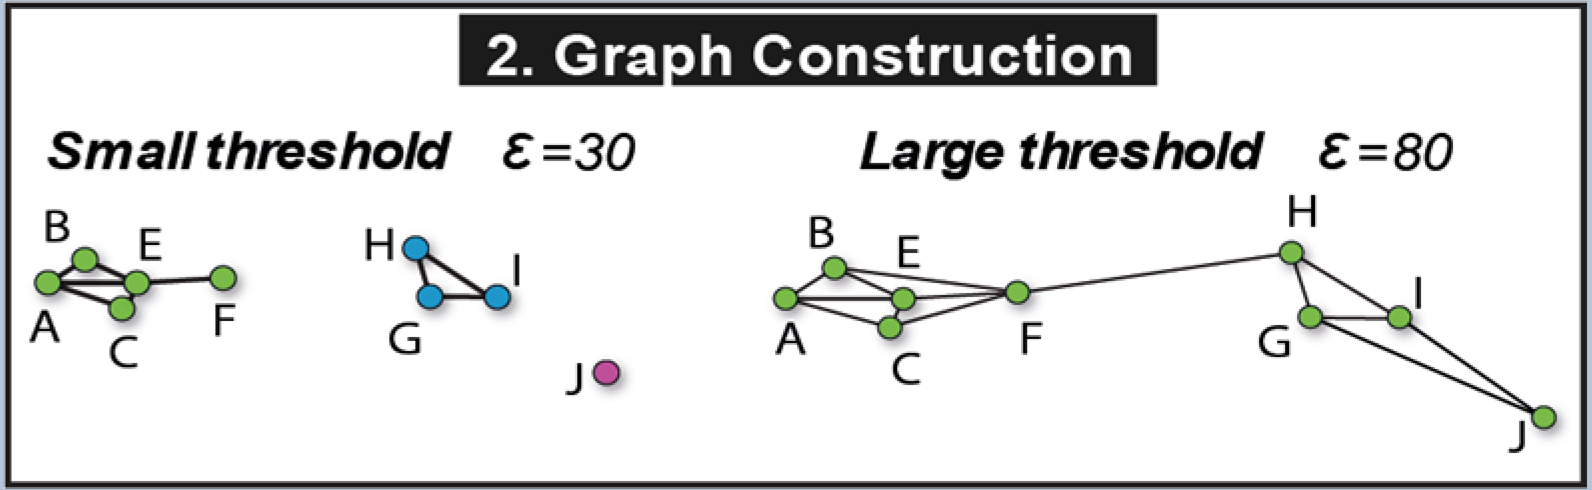

Inspired by Algebraic Topology filtrations

Building graphs

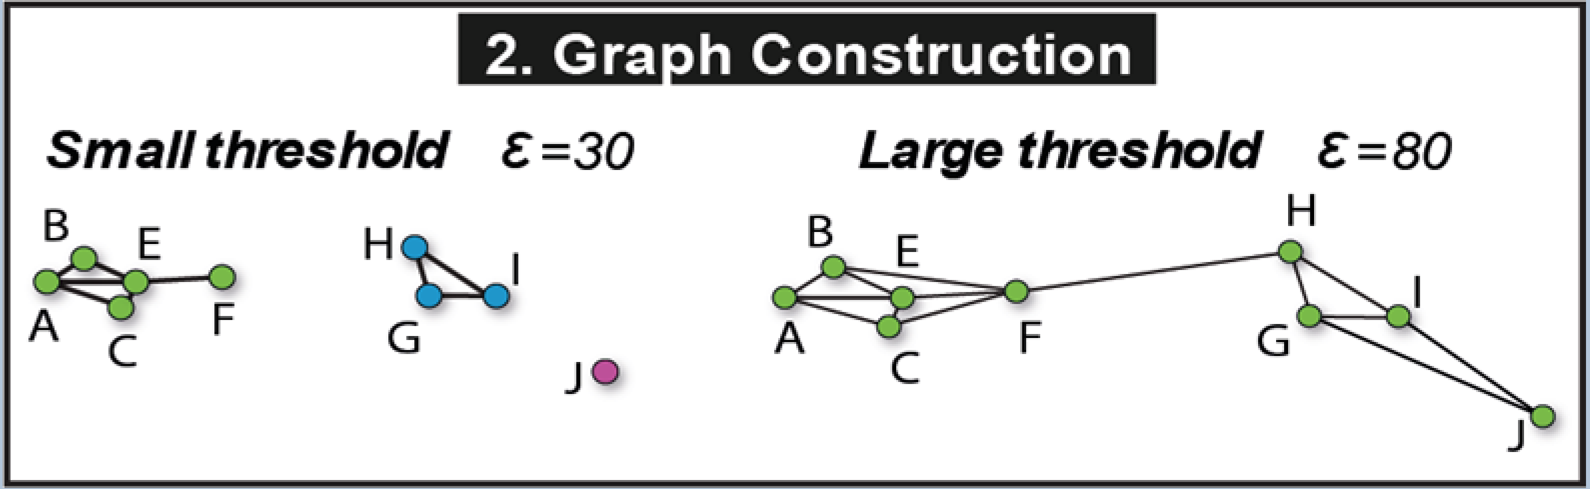

Inspired by Algebraic Topology filtrations

How to detect when the real topology has been uncovered?

Geometrization of graphs

Idea: Leveraging geometric ideas to discover shapes through the discrete Olivier-Ricci curvature

Optimal Transport Recaps

Evaluate the cost of transferring a specific distribution μ of mass on the nodes of a graph to another distribution ν

Two distributions: $\sum_x \mu(x) = 1 \quad \sum_x \nu(x) = 1$

A distance : $d(x, y)$ or $c(x, y)$ is the cost of transport one unit of mass, i.e. Euclidean distance, Manhattan distance

Optimal transport: $\theta^*(\mu, \nu) = \arg\min_\theta \sum_{x,y \in V} \theta(x, y) d(x, y),$ constrained to

$$\sum_{x \in V} \theta(x, y) = \nu(y) \quad \text{for all } y \in V.$$

$$\sum_{y \in V} \theta(x, y) = \mu(x) \quad \text{for all } x \in V$$

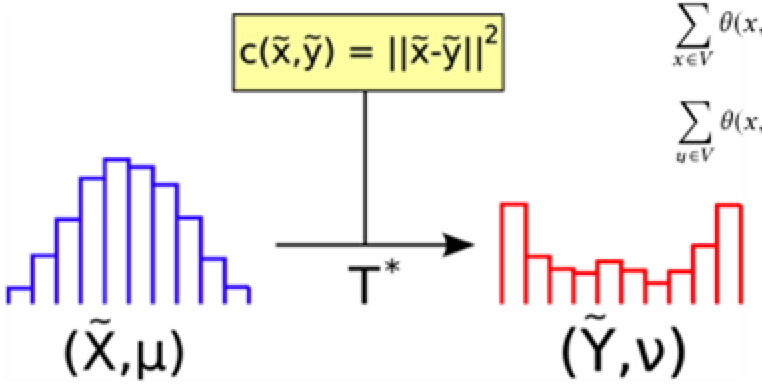

X

A worked out example

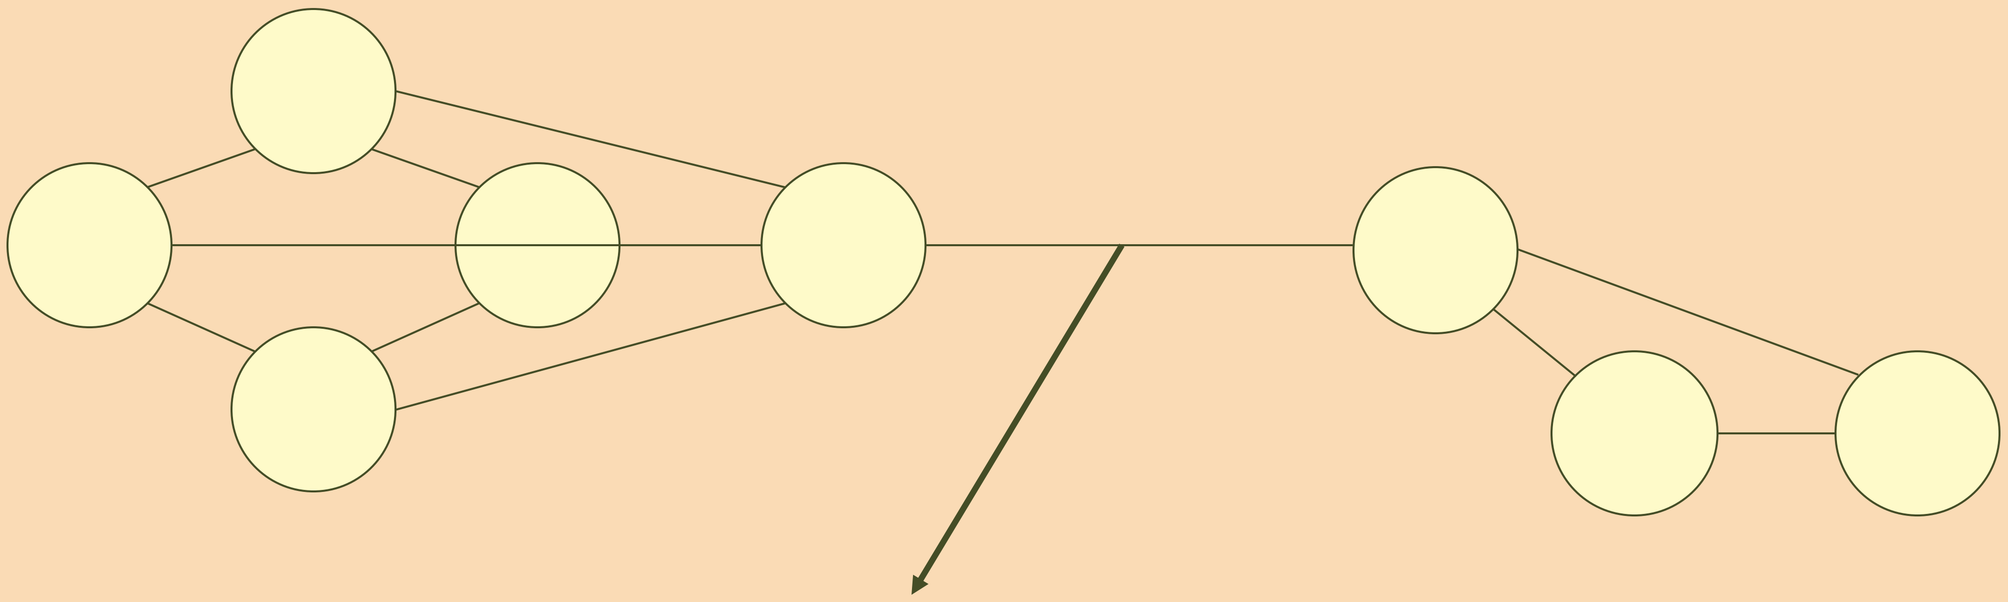

We are interested in computing the optimal transport of this link neighbors

Put dirt on the green nodes and empty boxes on the blue ones

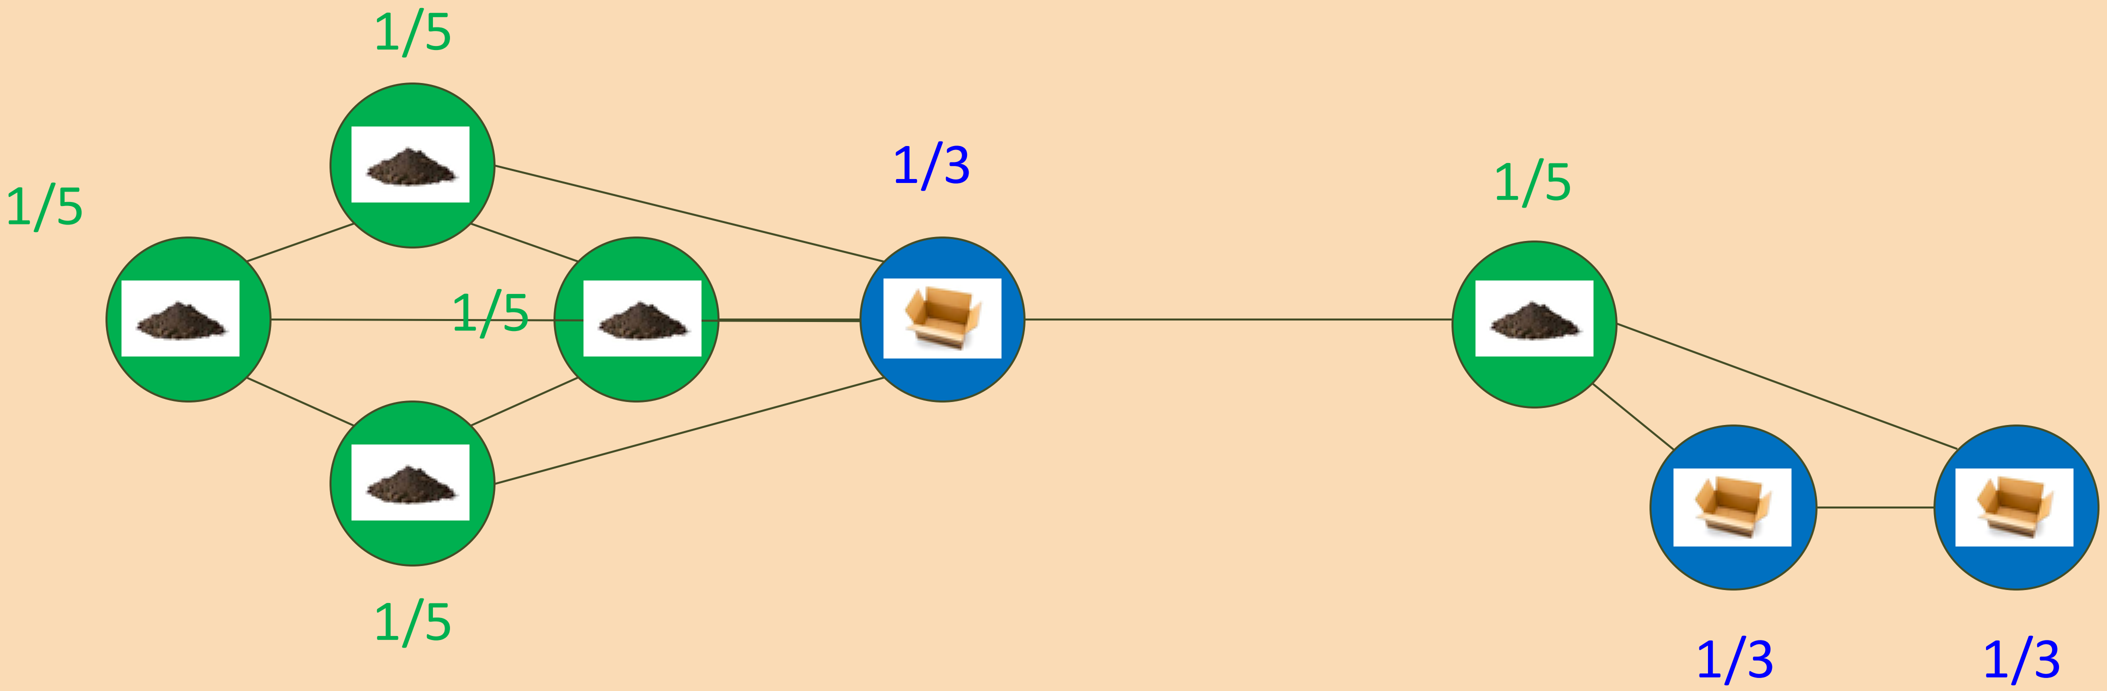

Move the weight from upper node to an empty box

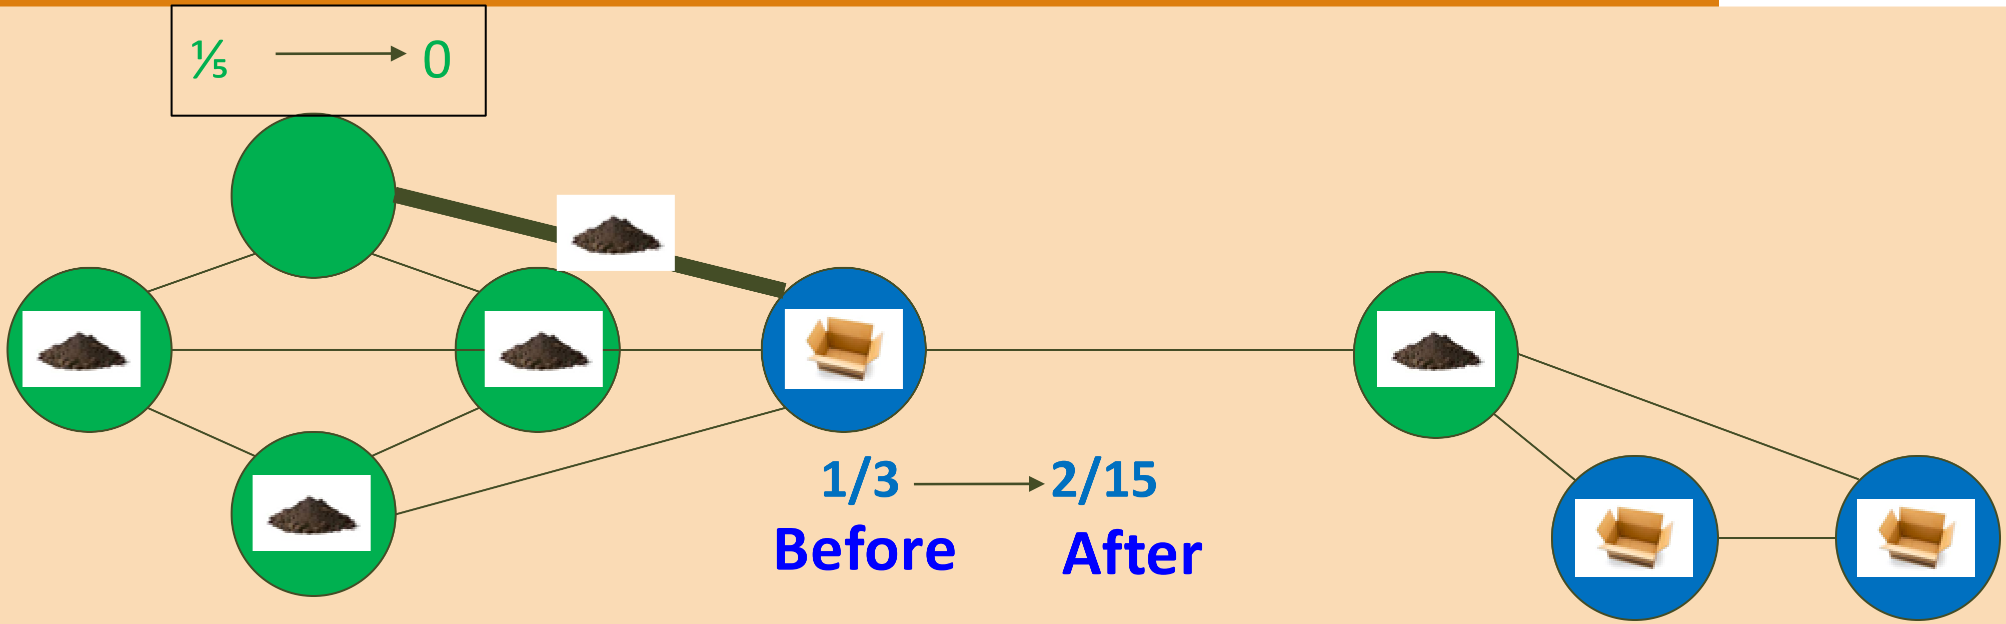

Filling a box by moving a parcel of dirt

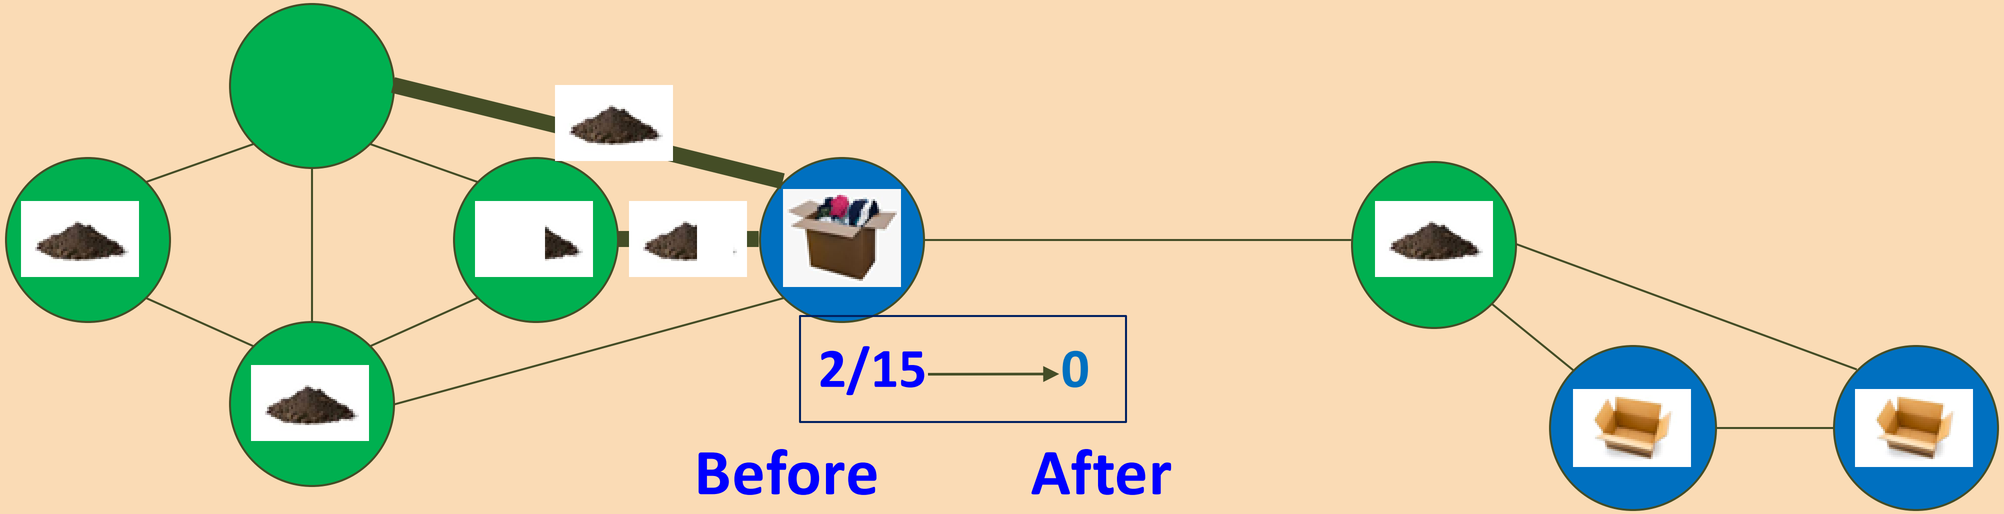

The edge at the middle is a bottleneck

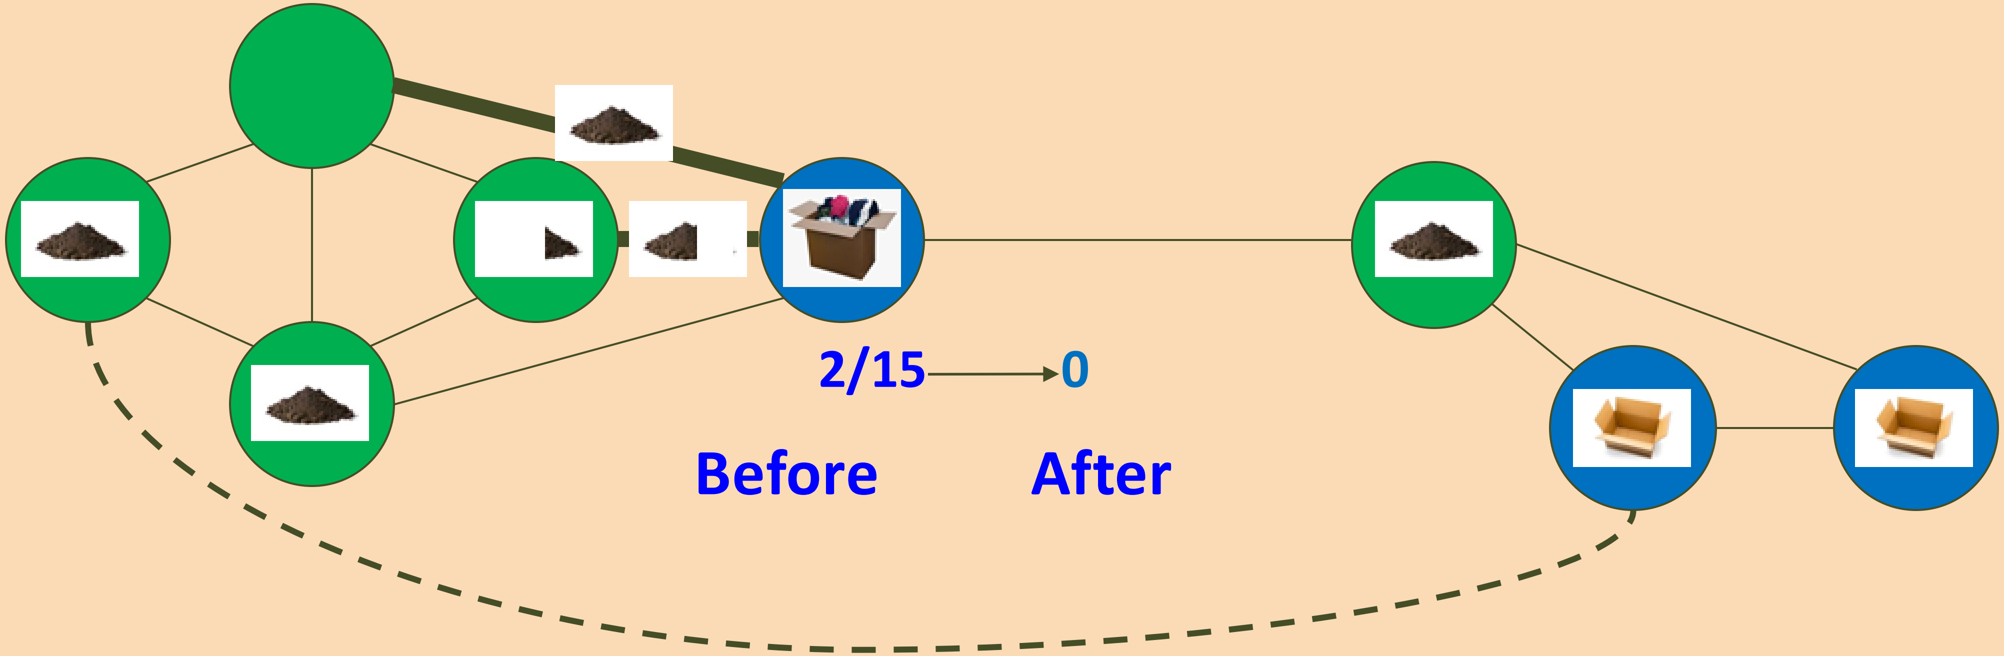

Only far boxes available!

Identifying the edges that have carried a lot of dirt

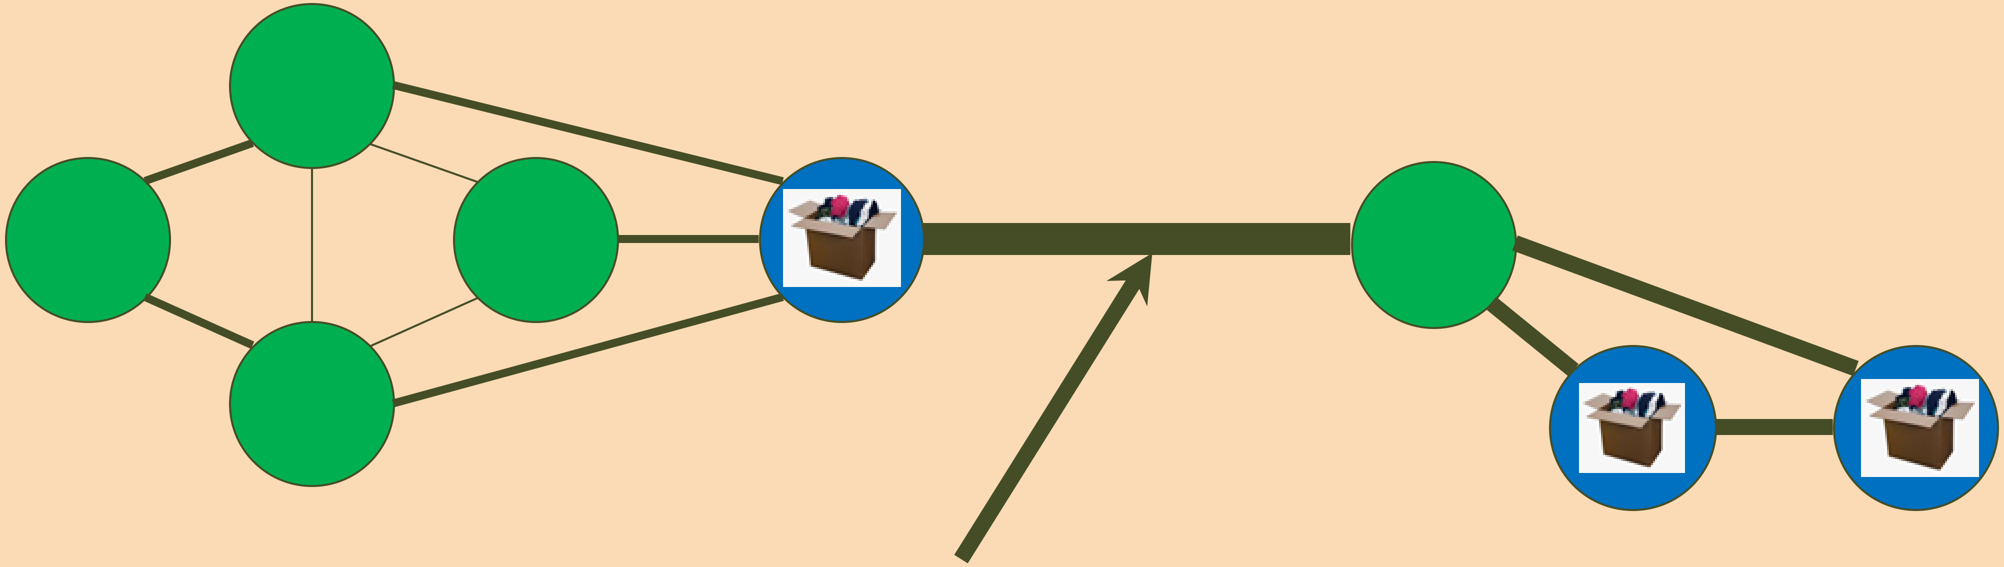

Lot of the dirts cross this edge!

Continuous and discrete curvature

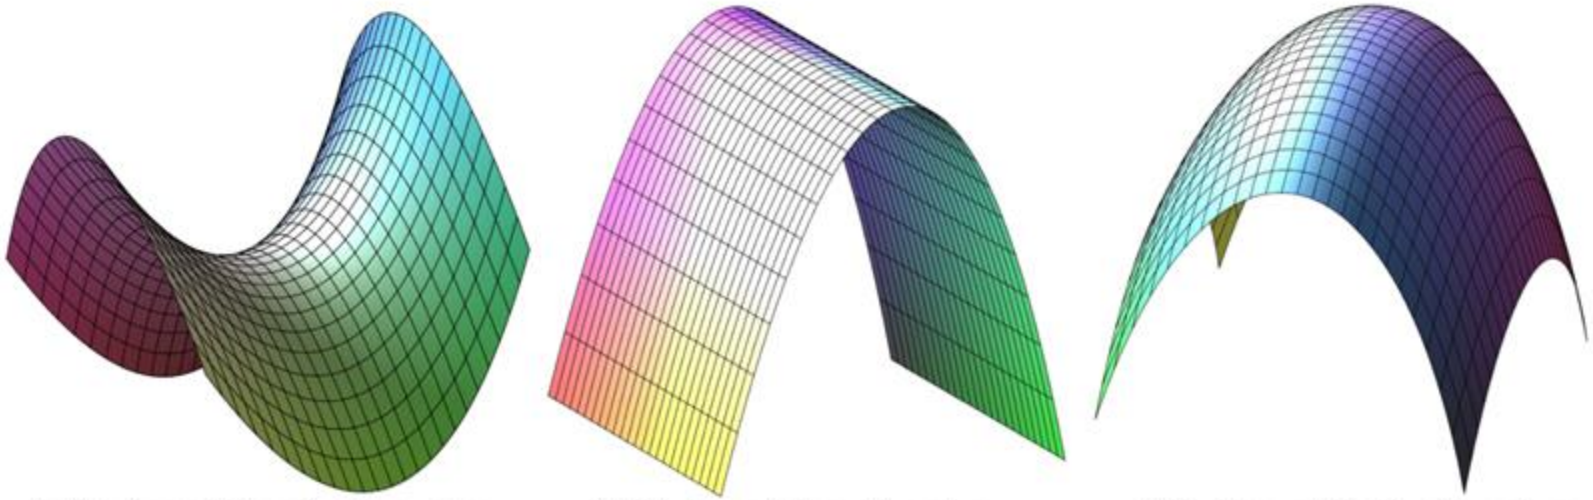

(a) Surface of Negative Curvature (b) Surface of Zero Curvature (c) Surface of Positive Curvature

Negative Curvature

Distribution μ

Distribution ν

Optimal mass displacement

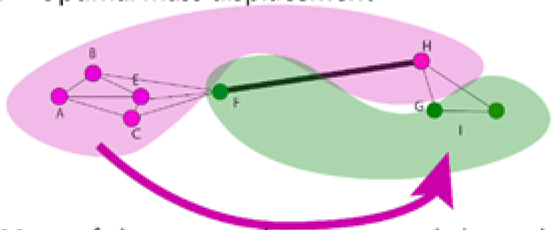

Most of the masses have to travel through the edge (FH)

Neutral Curvature

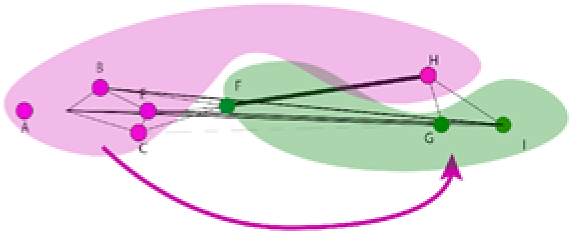

The neigborhoods are well connected to each other, (FH) is only transiting its own mass

Positive Curvature

Distribution μ and ν

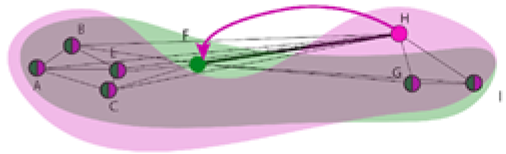

Most of the masses are mingled. The only mass that has to be moved is the one from F to H

Negative Ricci curvature matters

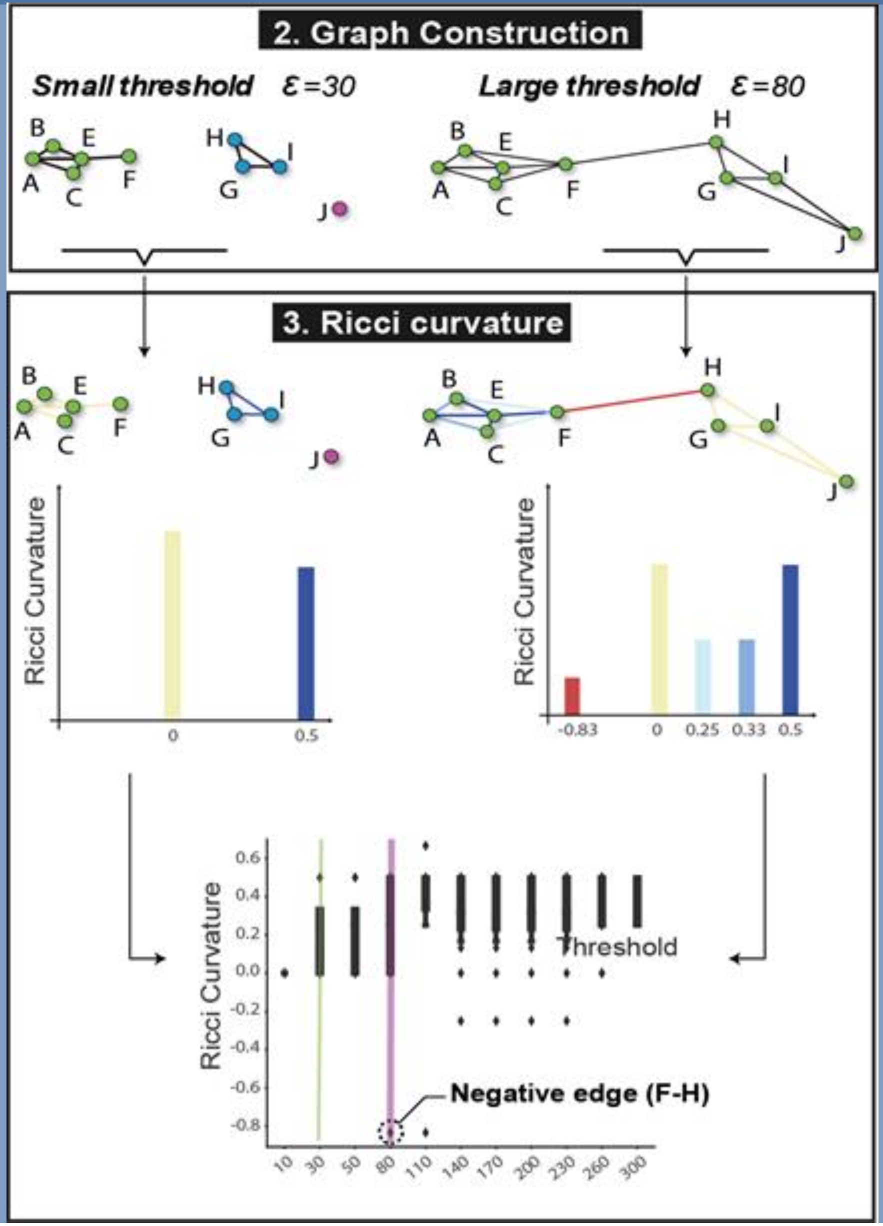

Finding interesting threshold and edges

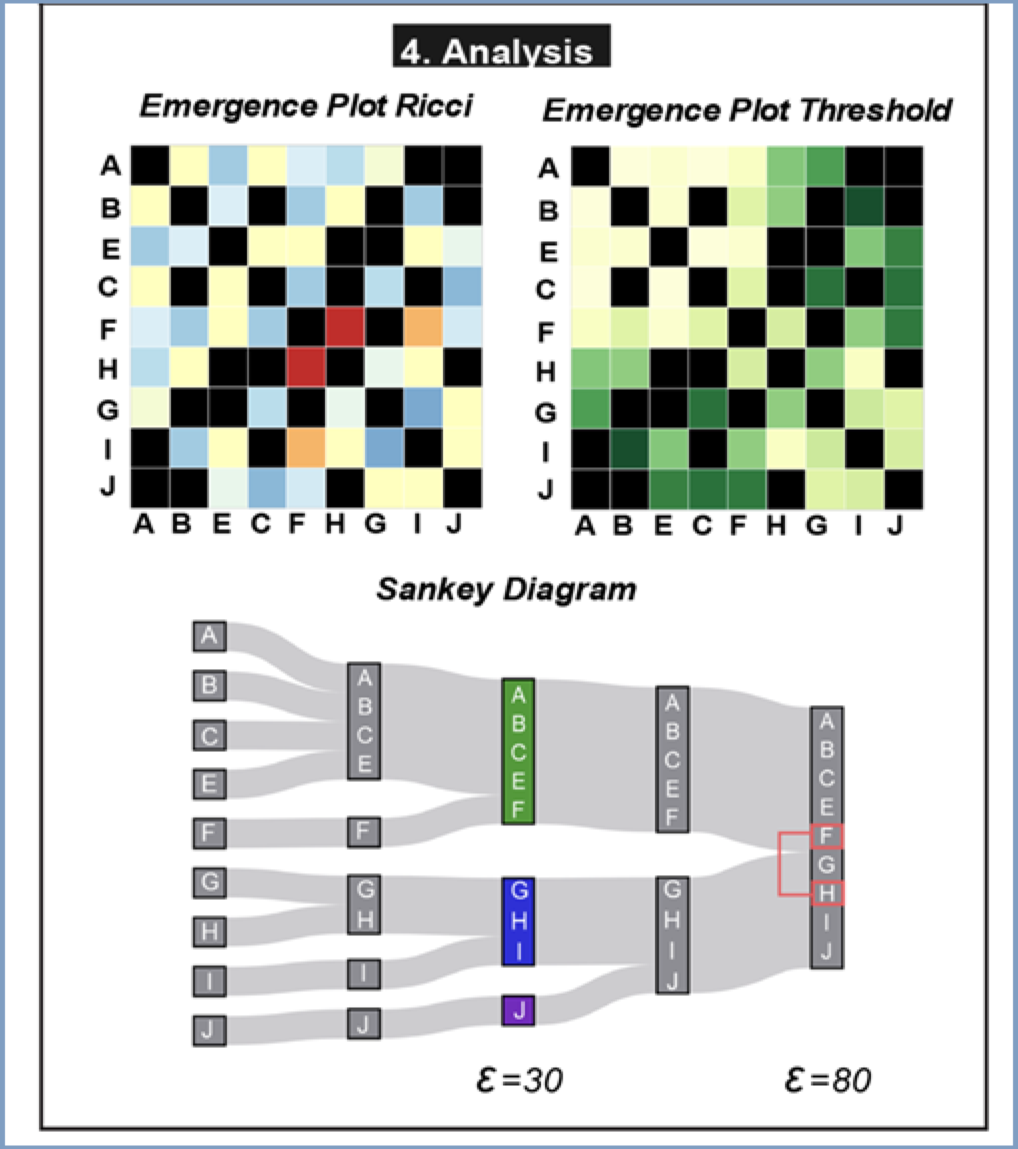

COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK

5) Merging everything into a single manifold visual

Linking back to the continuous view

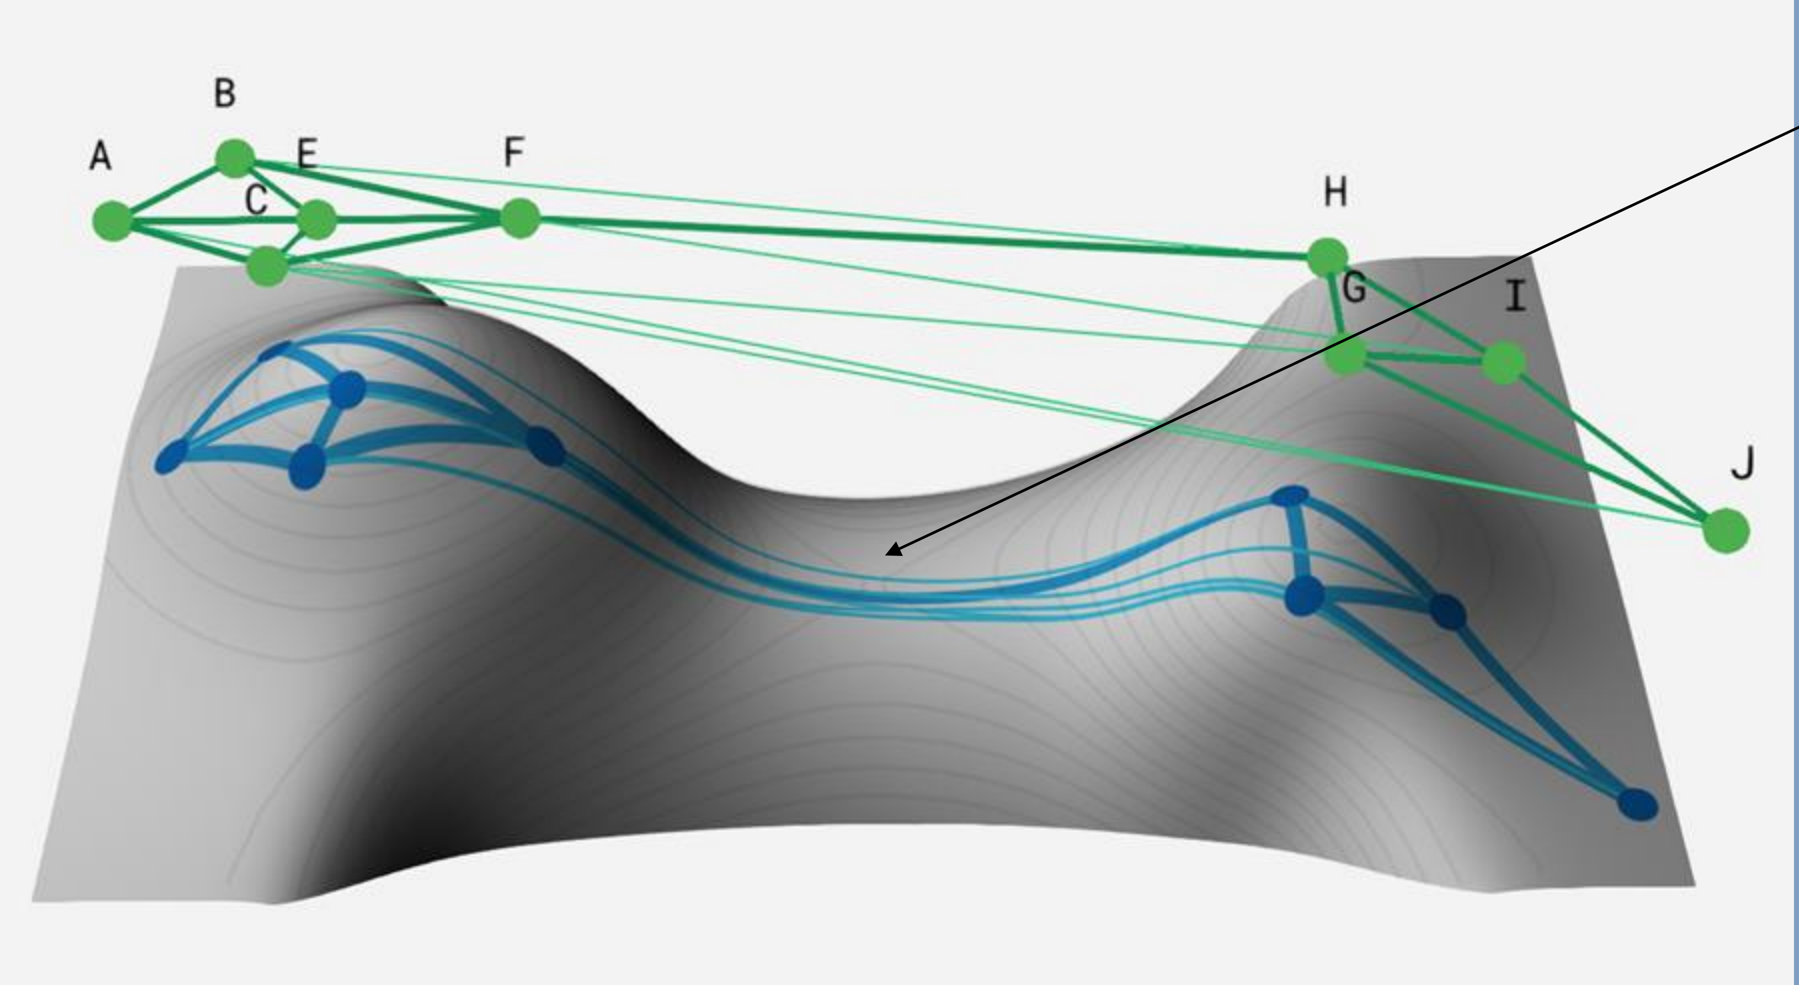

Geodesic fits physical topology

This view exists for each threshold

Onto the Application

1) Data Collection

2) Building the graphs

3) Analyzing the manifold

4) Validating the manifold

5) Sensitivity analysis

We schedule ping measurements from the different CPs

From the Cloud

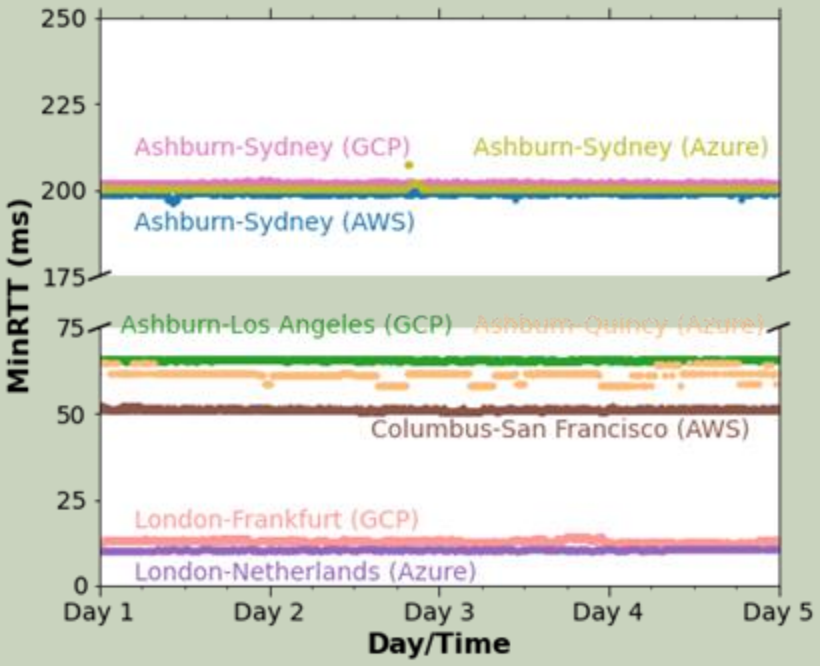

Public Internet in comparison

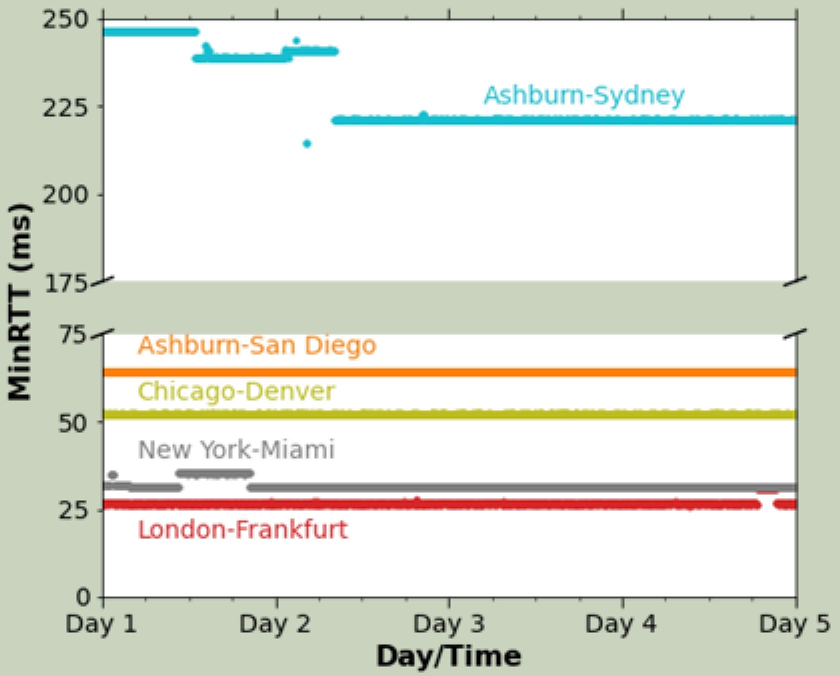

Identifying interesting thresholds

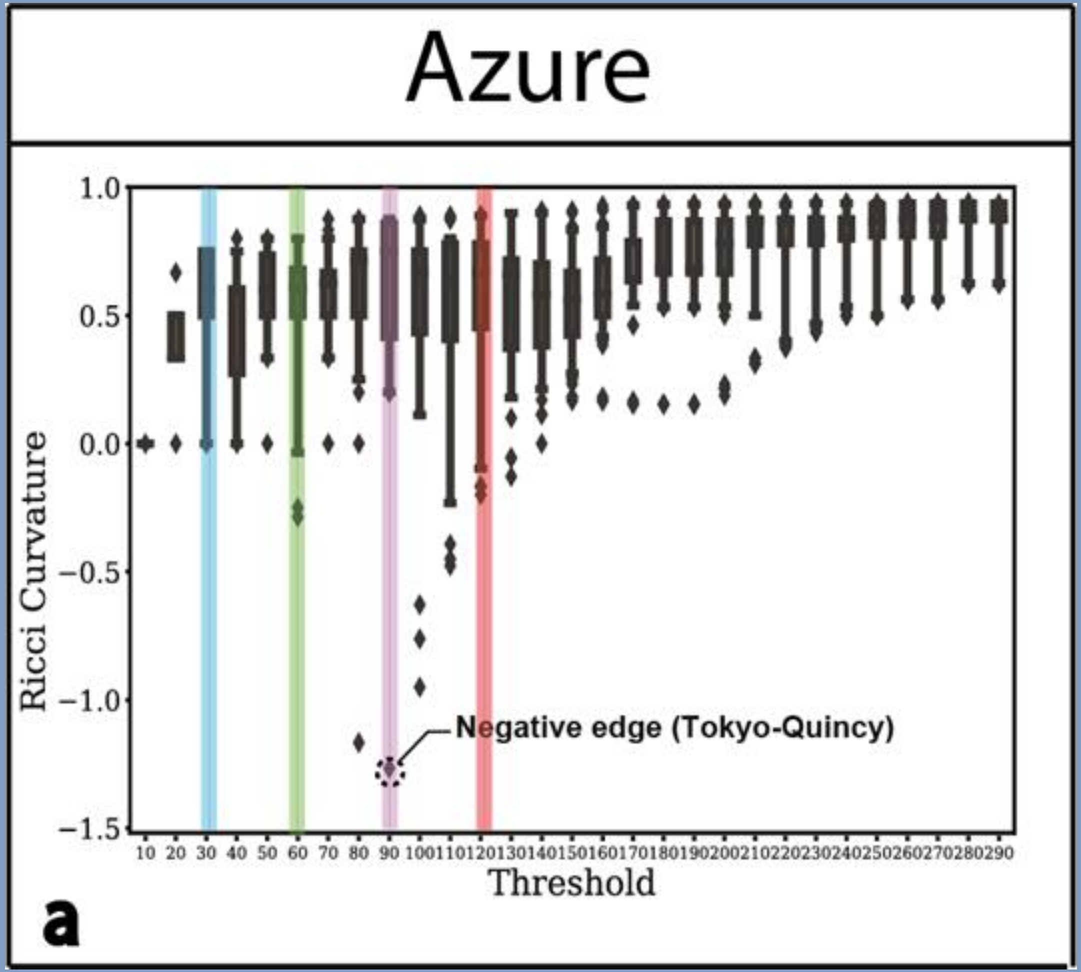

Building the graphs

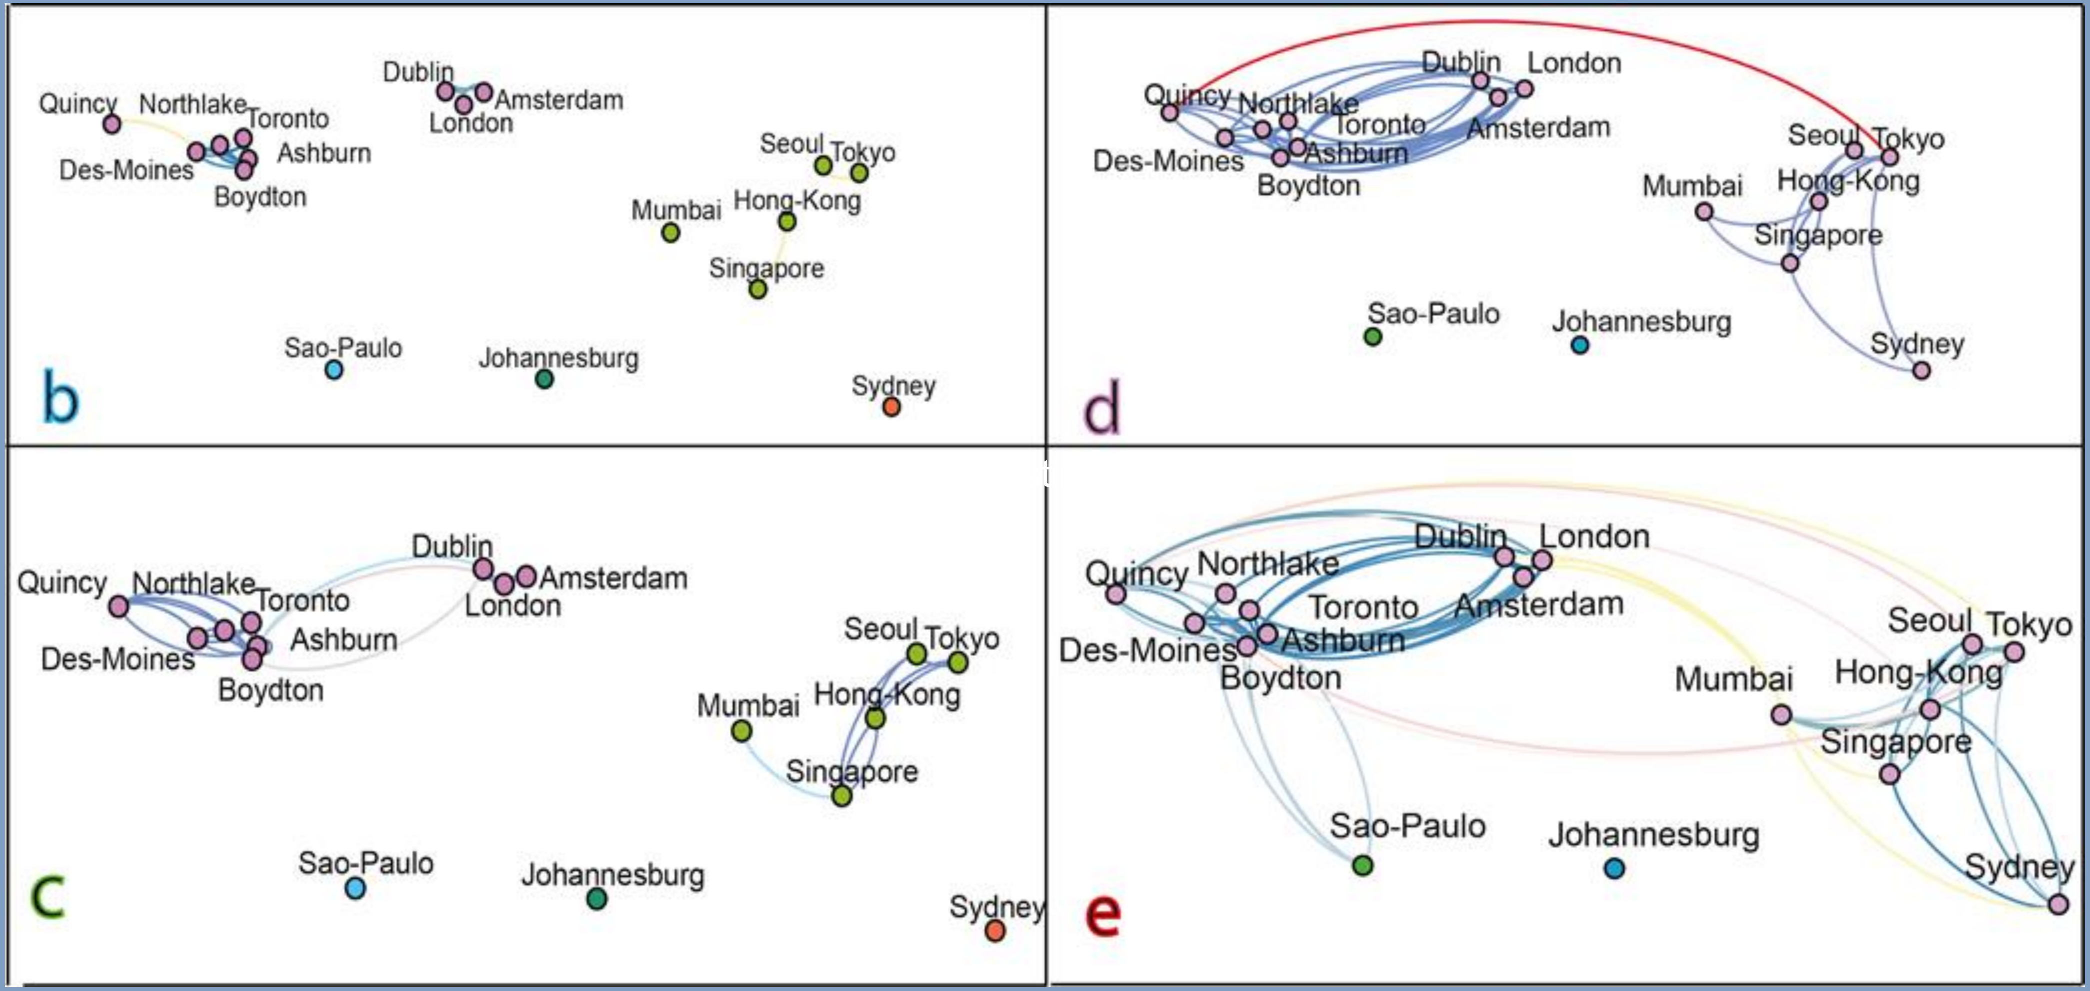

Aggregating the observation through the heatmaps

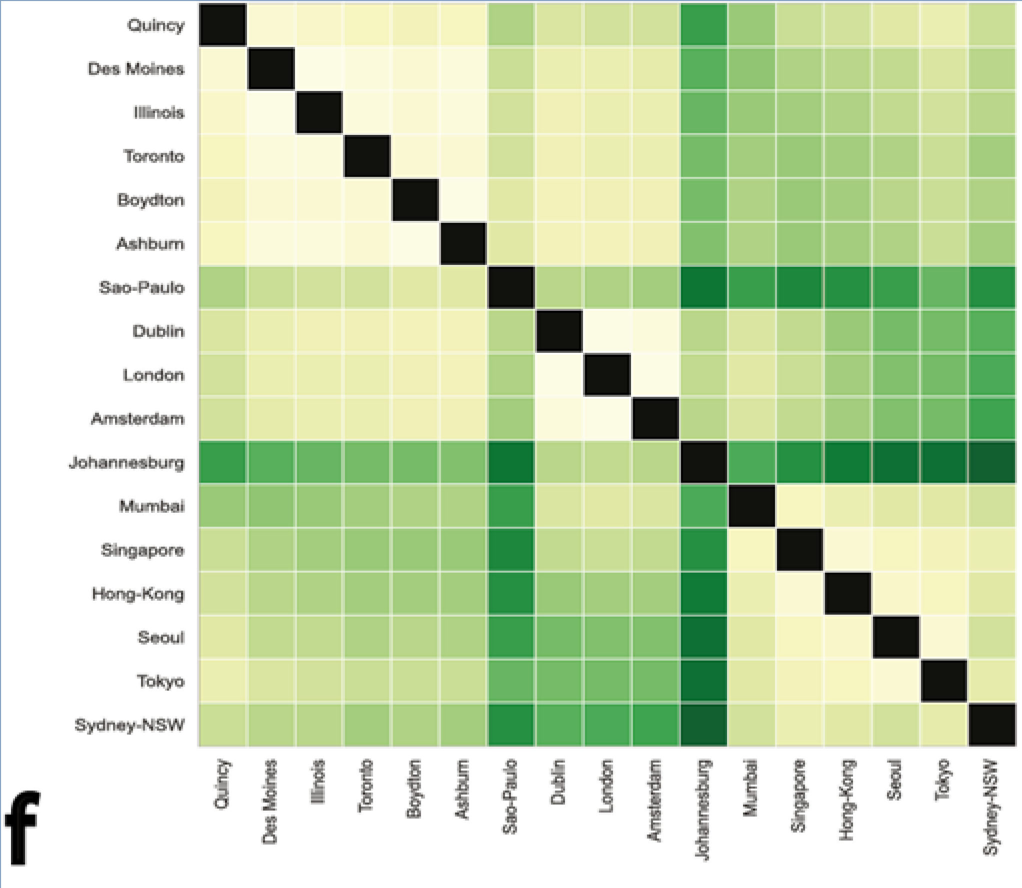

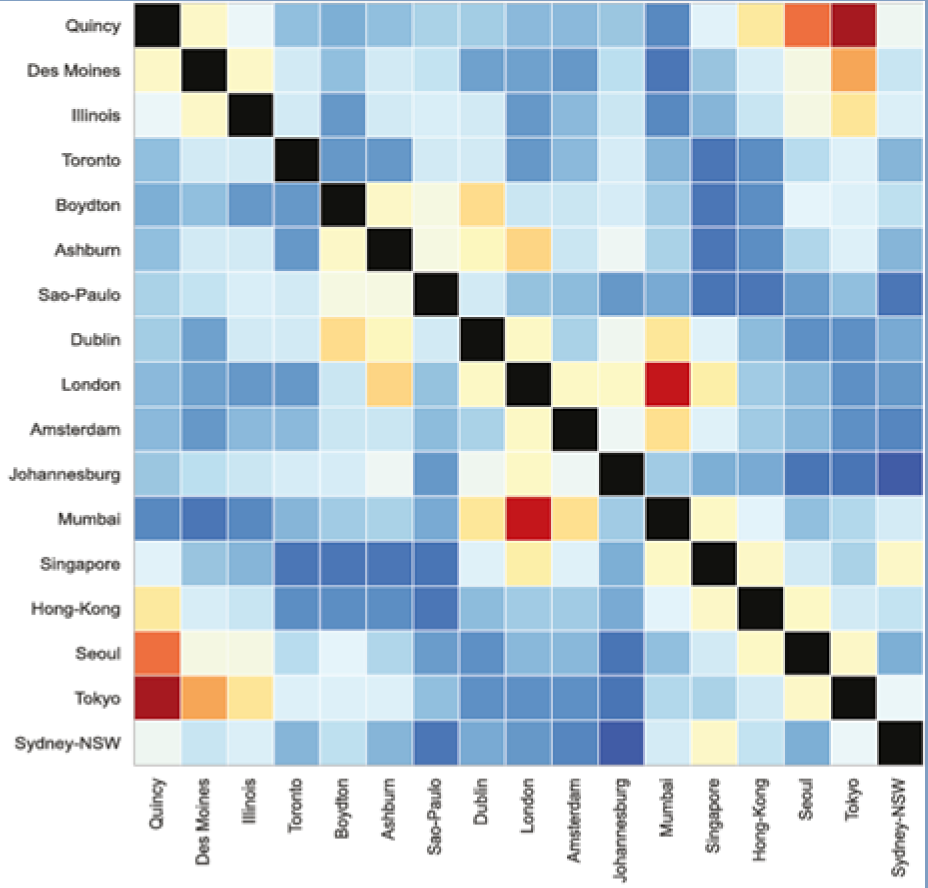

How different islands of connectivity emerge and disappear

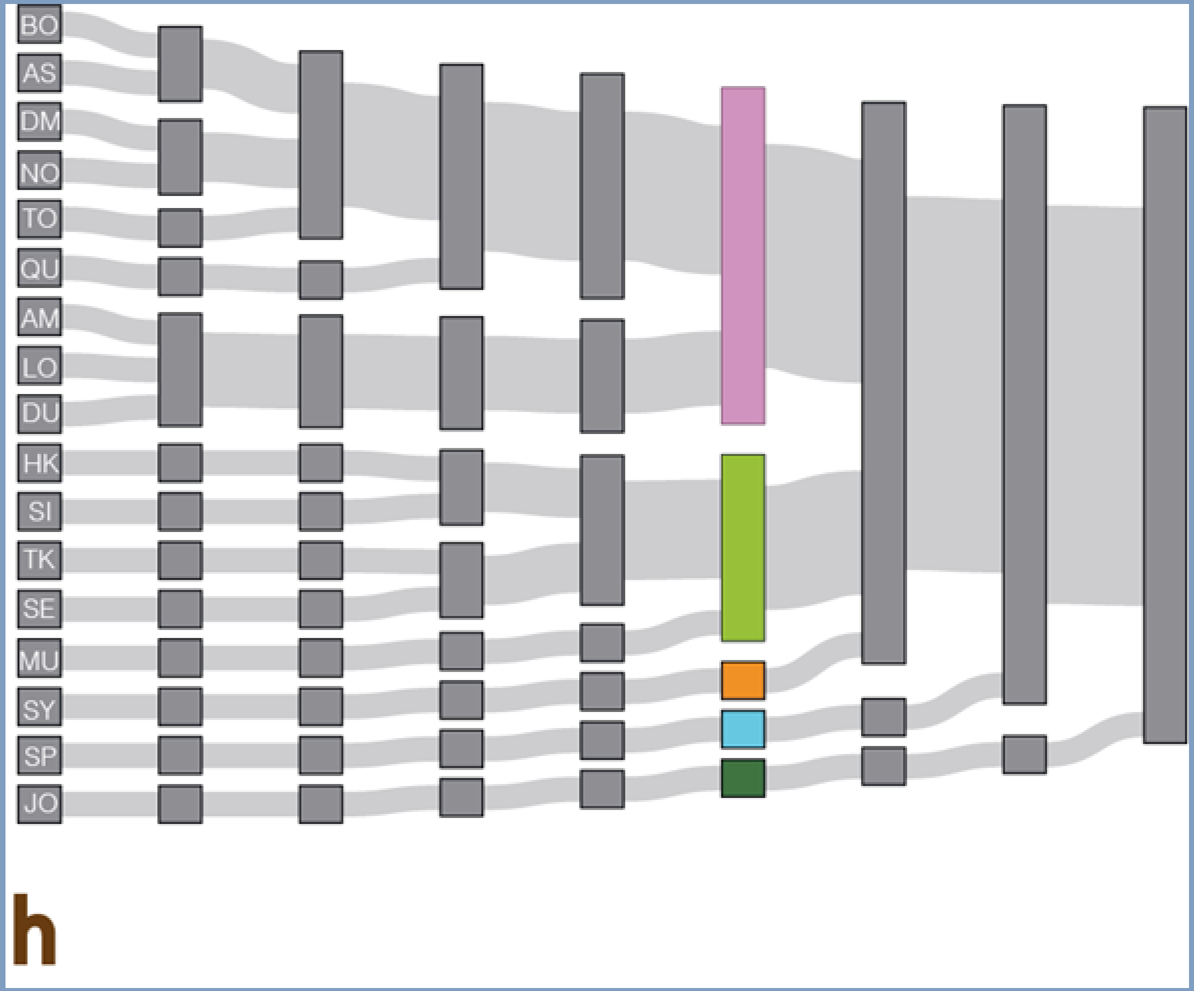

Building a single manifold view

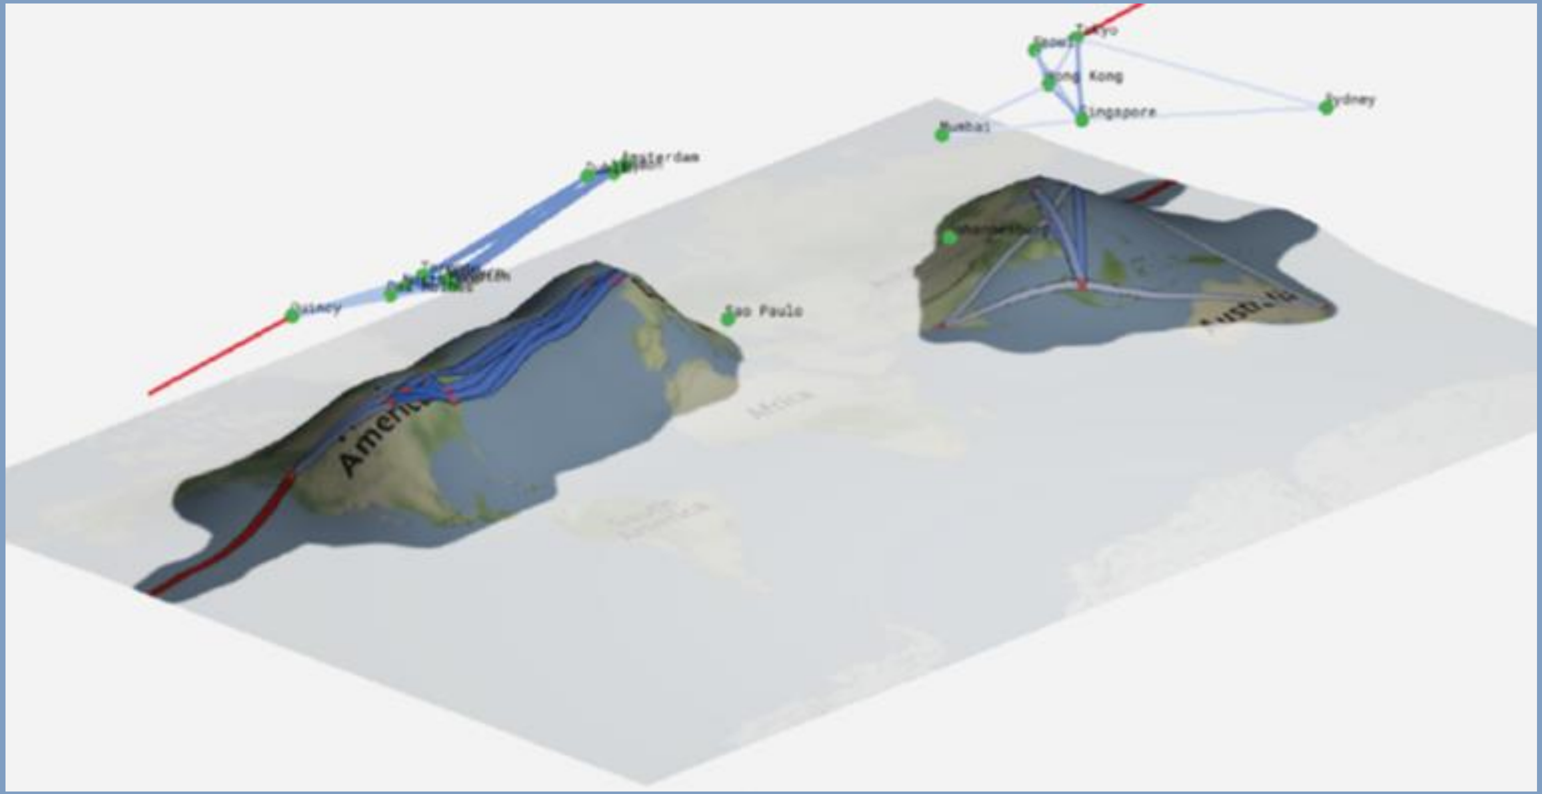

Do we recover the cables according to Microsoft?

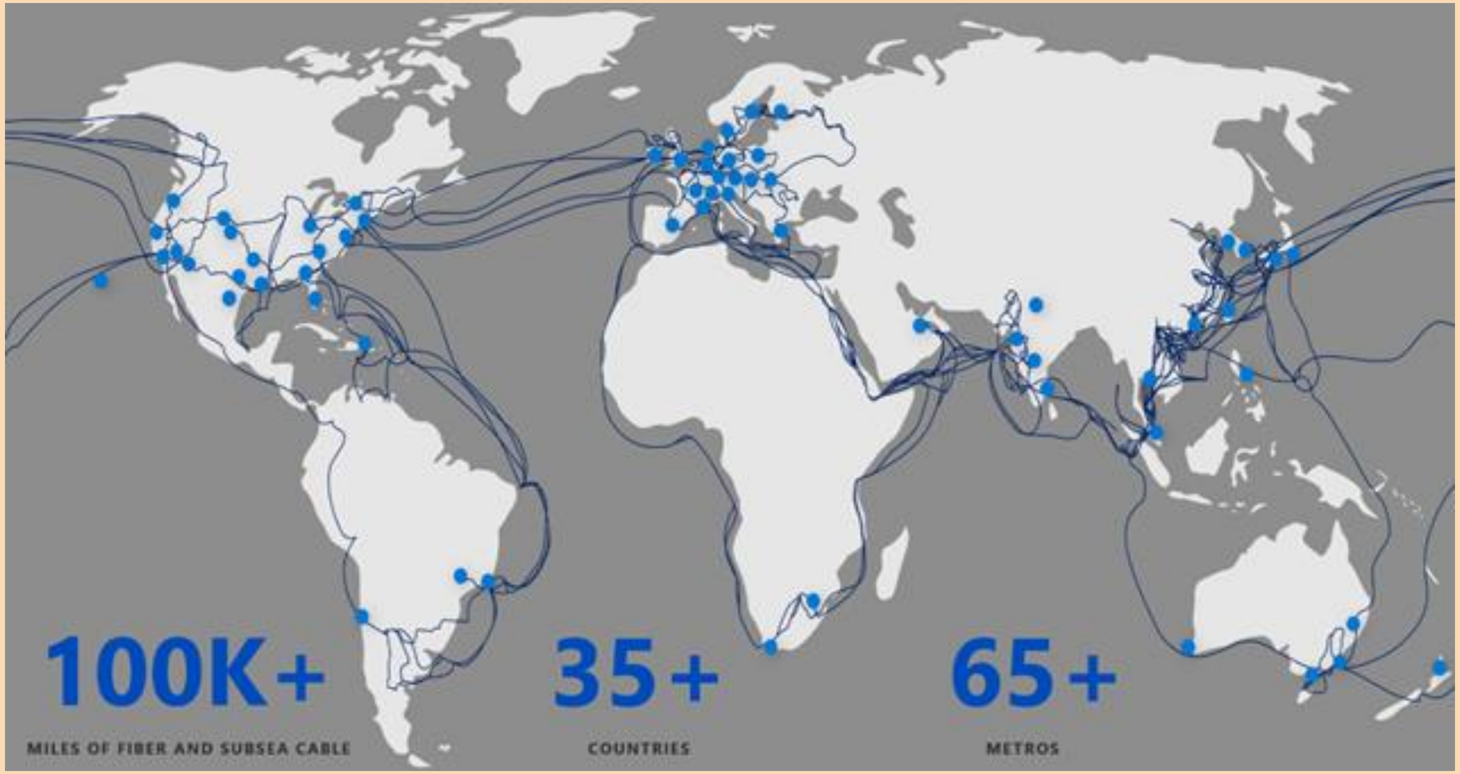

https://www.microsoft.com/en-us/research/project/project-arno-cloudification-of-telecom-network-infrastructure/

Comparing cloud providers

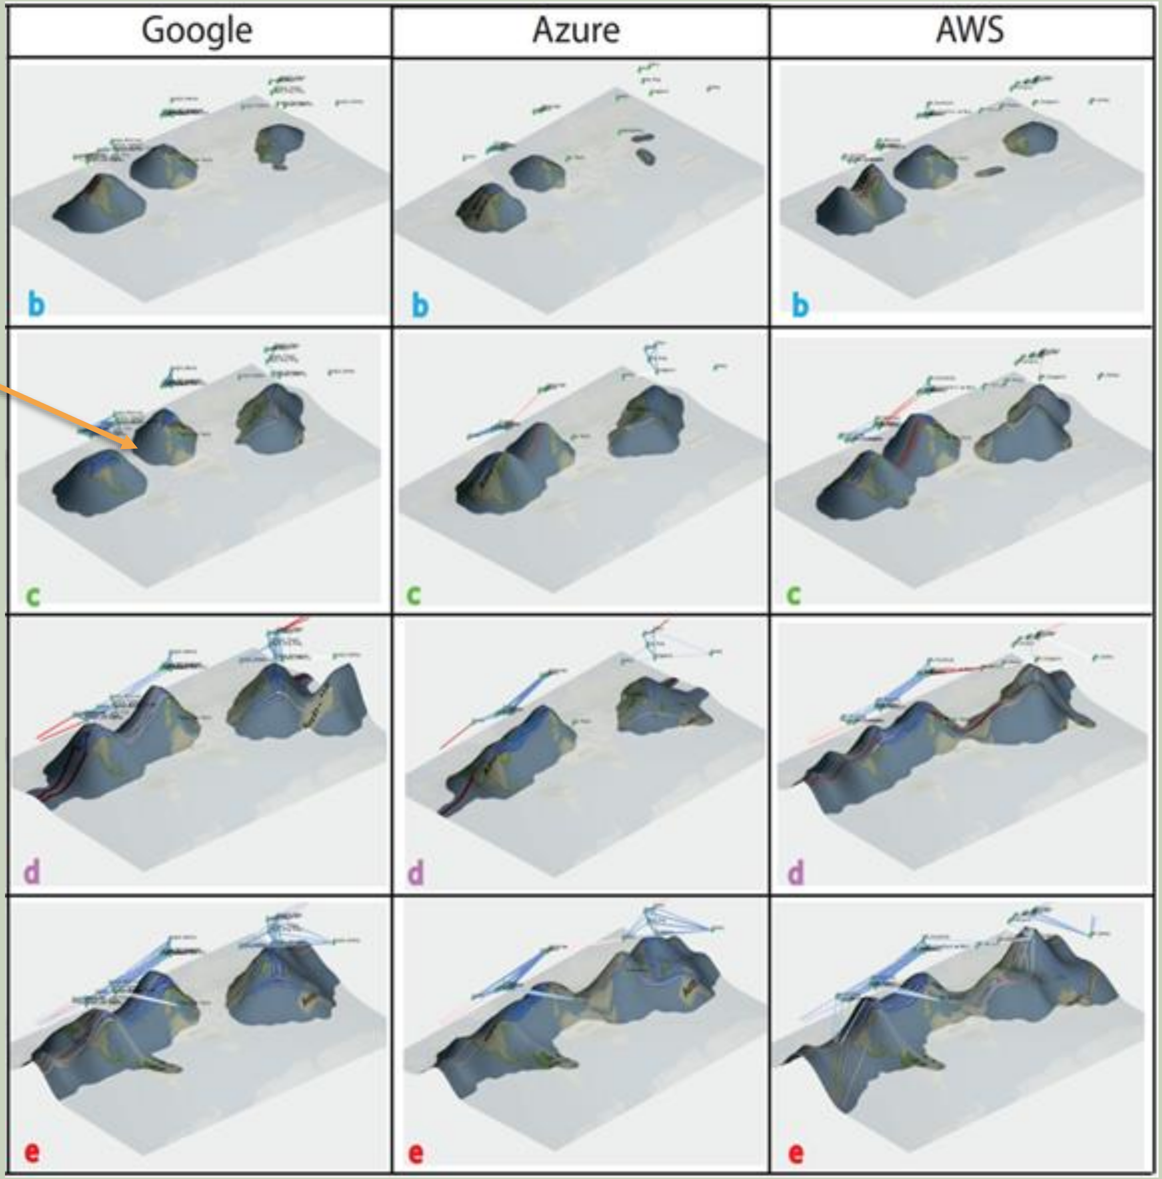

Google has not percolated between Europe and US

COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK

Comparing cloud providers

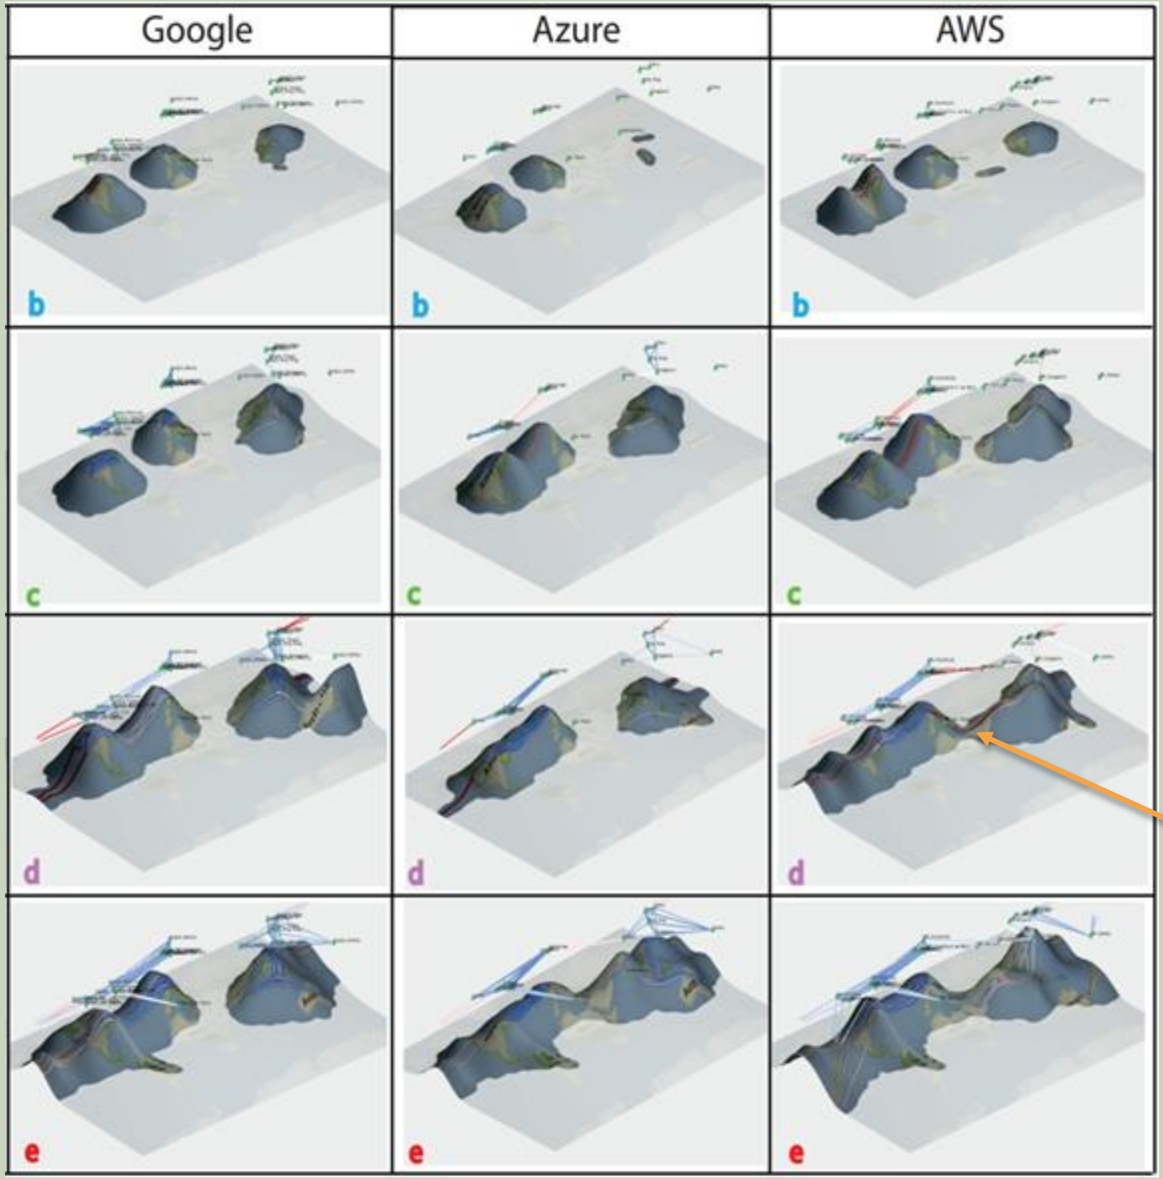

AWS is the only cloud that has a direct path between Asia and Europe

Comparing cloud providers

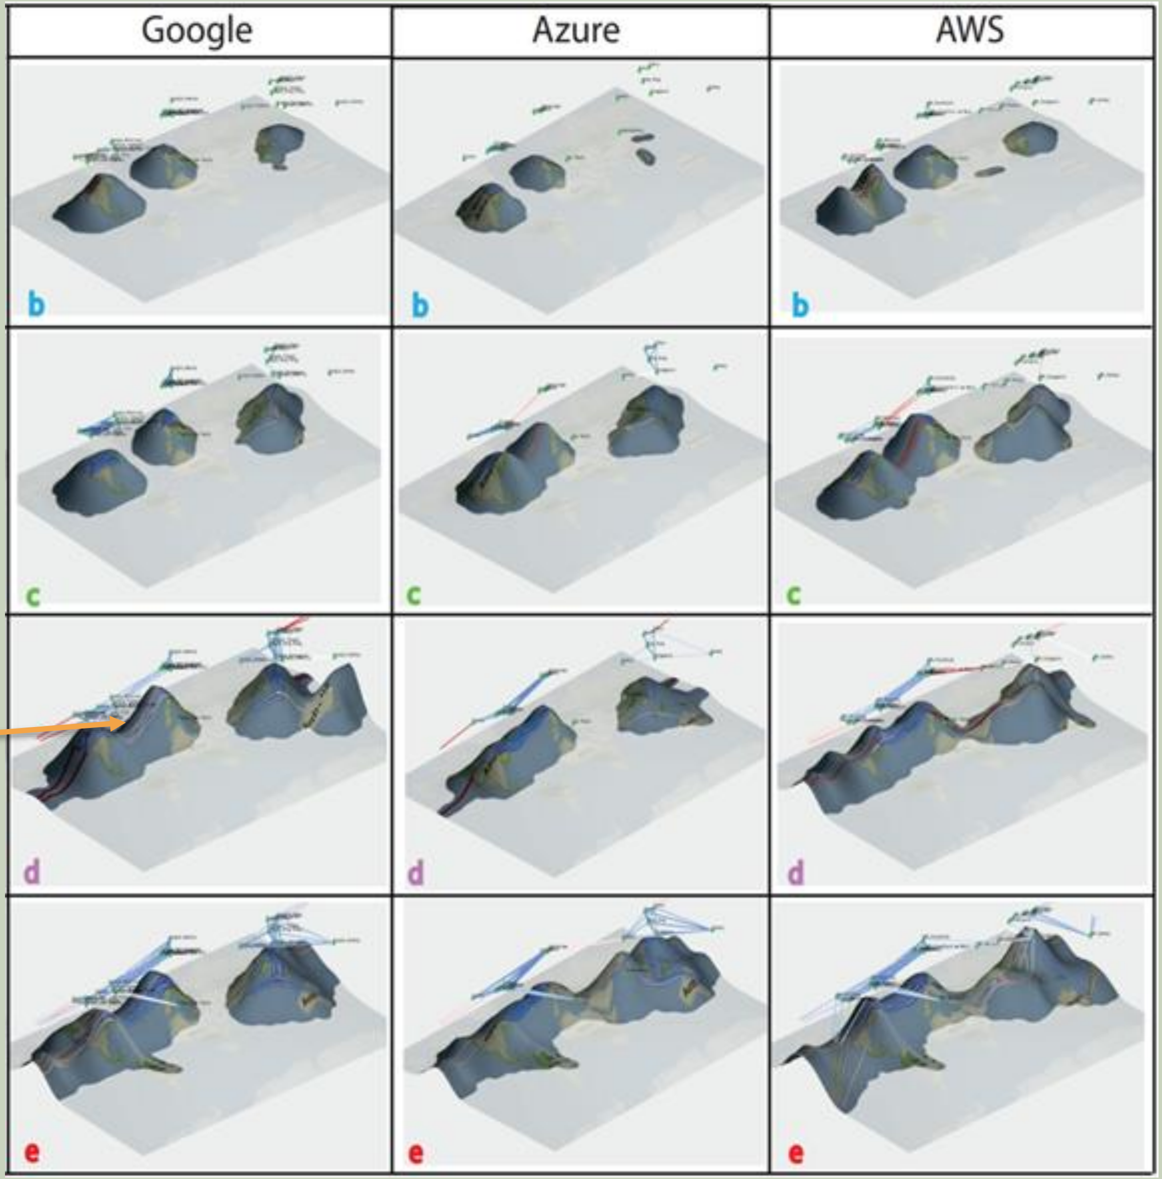

Google has the most positive curvature between NA and EU

Comparing cloud providers

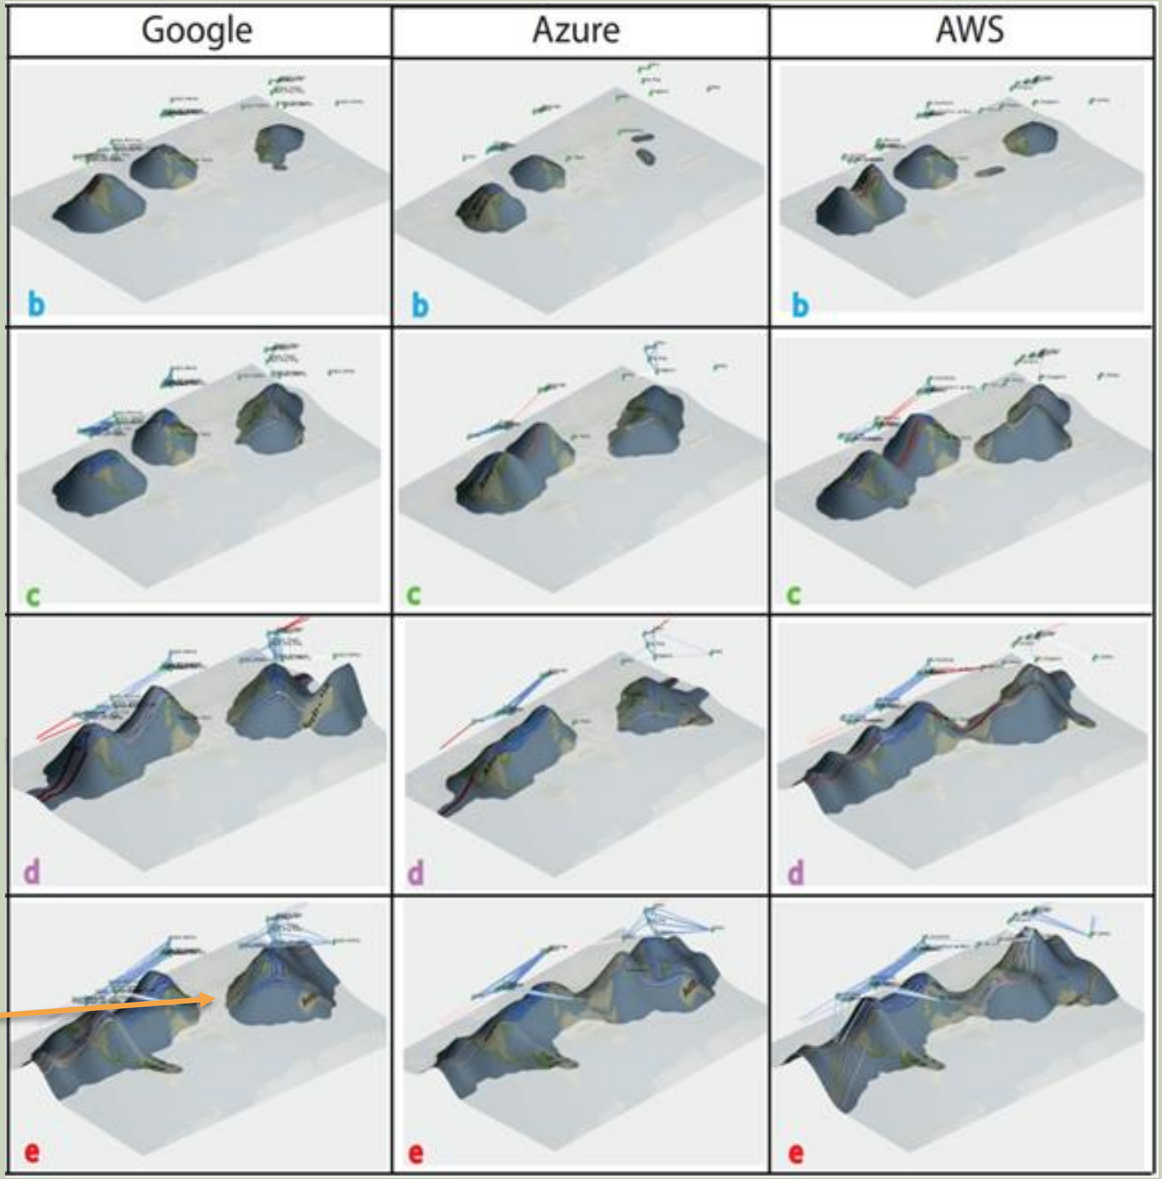

Google is not connecting Asia to Europe directly.

The geodesic on the manifold predicts the latency better than the great circle distance the GCD

Cloud	City A	City B	Latency	δ_{GCD}	δ_{Geo}	d_{RD}	d_{GCD}
Amazon	Sydney	Tokyo	104.6	-28.5	-1.2	10184	8147
	Ashburn	Sao Paulo	118.3	-12.2	3.0	9580	7664
	Dublin	Tokyo	202.0	39.0	-11.3	11980	9584
	Paris	Singapore	155.5	-27.0	-4.0	13415	10732
	Mumbai	Paris	104	-14.5	-6.0	8761	7009
Google	Osaka	London	221	55.4	5.7	12205	9499
	Mumbai	London	333	211.4	42.3	9247	7191
	Las Vegas	Sydney	142	-74.6	-48.7	16578	14320
	Seoul	Jakarta	106	14.5	40.2	6421	5293
	Portland	Frankfurt	124	-22.3	-42.4	9872	7834
Azure	Sydney	Tokyo	106.7	-26.3	-0.1	10184	8147
	Ashburn	Sao Paulo	117.8	-12.4	10.1	9580	7664
	Dublin	Tokyo	233.5	70.5	10.2	11980	9584
	Seoul	Amsterdam	220.5	75.7	25.6	11263	8516
	Singapore	Des-Moines	194.2	-59.5	-56.8	17872	14926

Selecting route distance over great circle distance

Right of way distance (or route distance) as a more appropriate metric to encapsulate topography

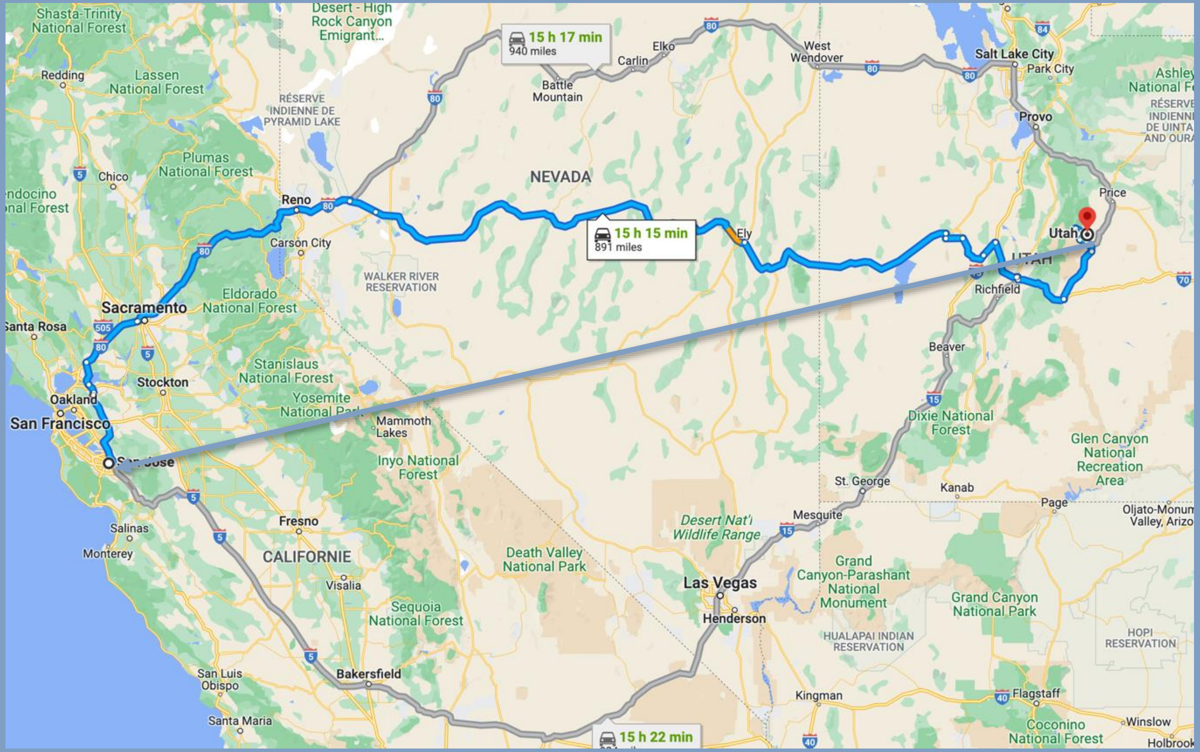

Selecting routing distance over great circle distance

Right of way distance (or route distance) as a more appropriate metric to encapsulate topography

GCD is sufficient for our purpose

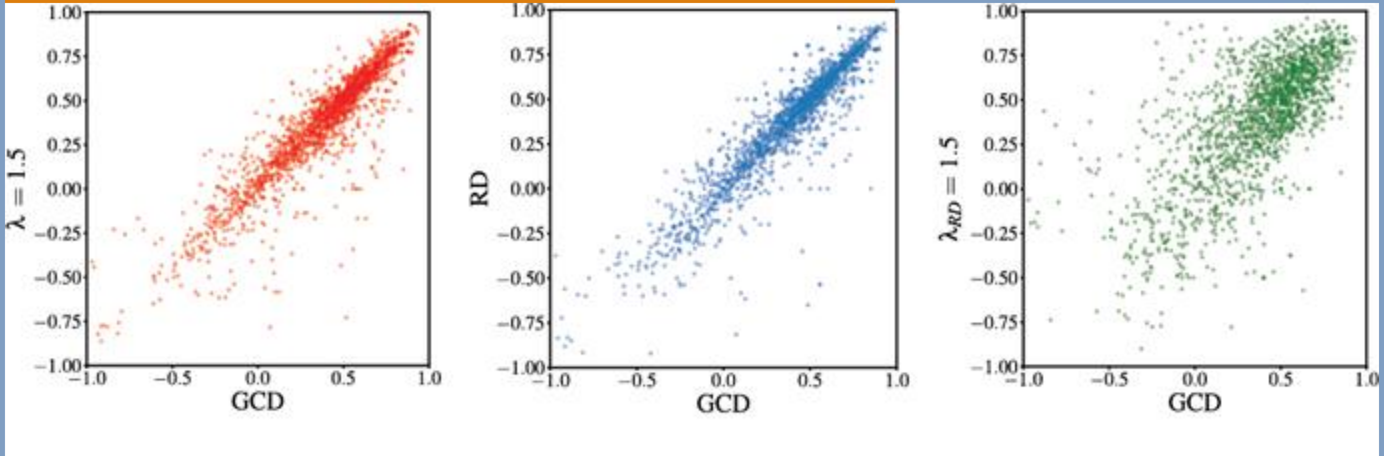

Changing granularities

Selecting quotient over difference for remaining latency

Quotient is a good metric when identifying smaller scale patterns

Changing granularities

Selecting quotient over difference for remaining latency

Quotient is a good metric when identifying smaller scale patterns

Selecting smaller time-windows

Changing time used for selecting the minRTT provides different view

Alternative graph metrics

Selecting Ricci curvature over different bottleneck metrics

Cut metric and betweenness centrality are not robust metrics.

Lemma:

The maximal change in the L^1 of the Ricci curvature distribution that the addition or deletion of an edge can induce is 6.

Future steps

Sharing the manifold visualization tool and keep on improving its outputs

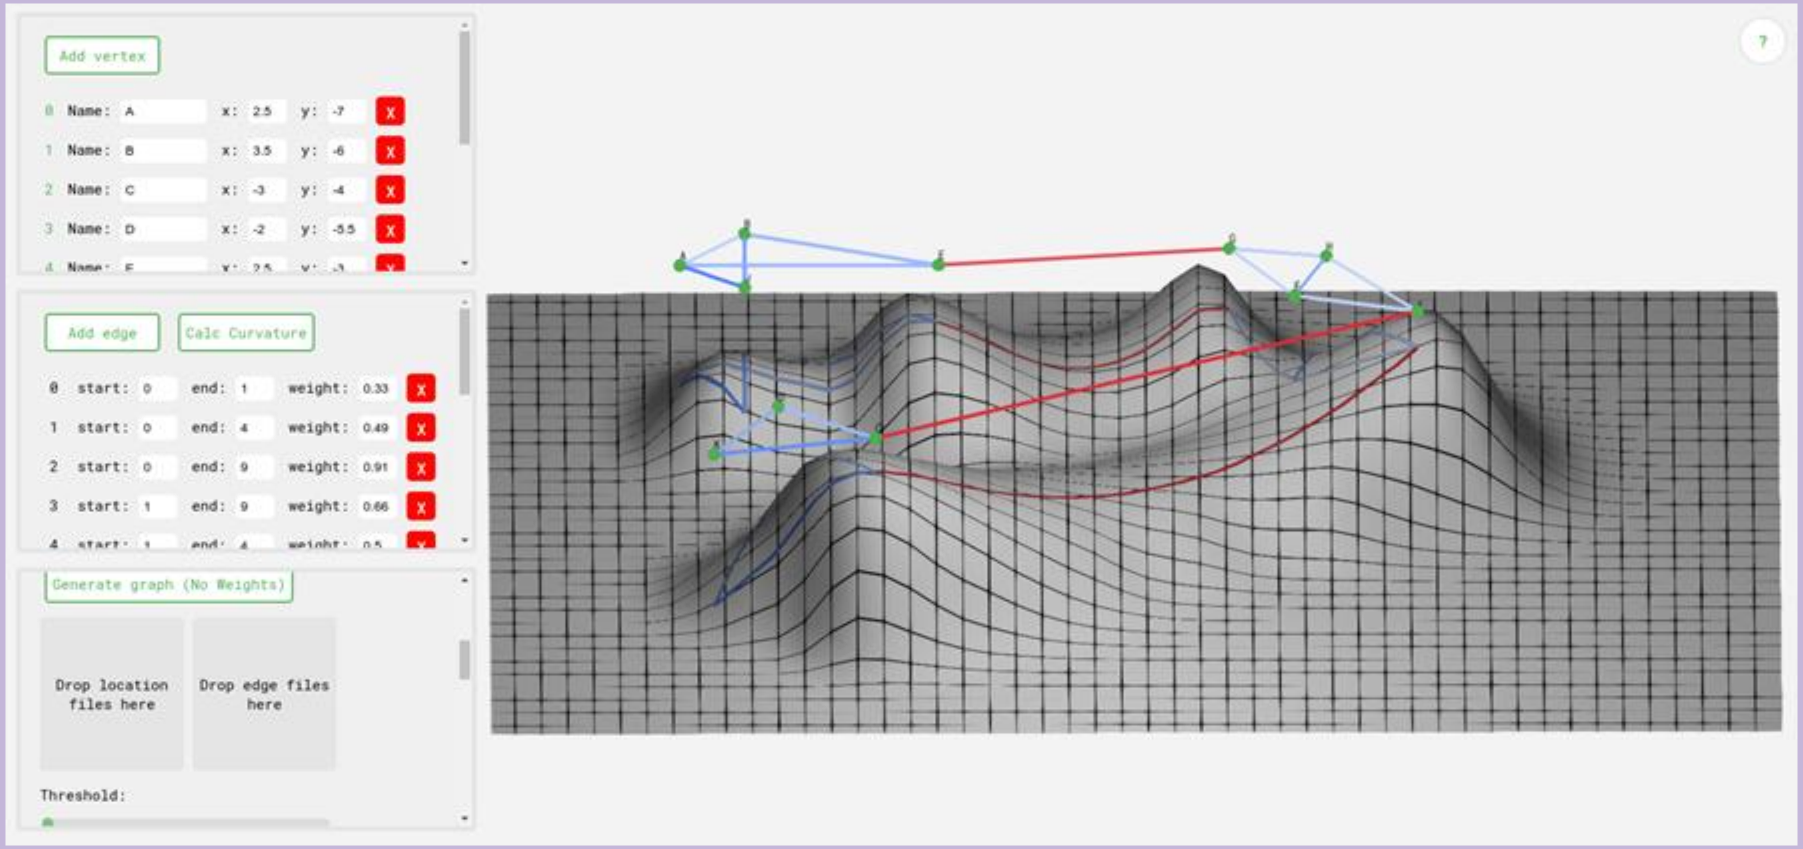

COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Future steps

Sharing the manifold visualization tool and keep on improving its outputs

Applying the technique for Public Internet

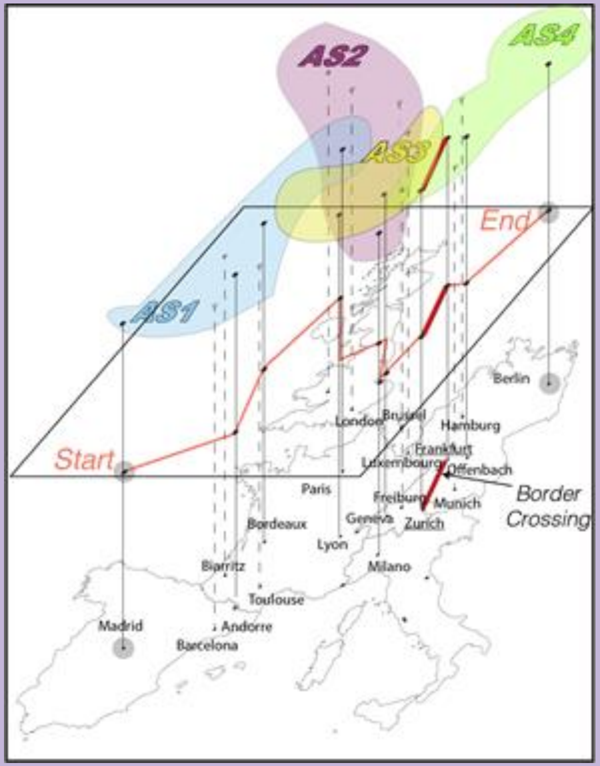

Routing is way harder for public Internet

Conclusions

In this presentation, we have:

1. introduced a new technique to uncover topology that was based on Riemannian geometry.

Conclusions

In this presentation, we have:

1. introduced a new technique to uncover topology that was based on Riemannian geometry.

2. built visualization that facilitates the utilization of our technique.

Conclusions

In this presentation, we have:

1. introduced a new technique to uncover topology that was based on Riemannian geometry.

2. built visualization that facilitates the utilization of our technique.

3. compared and contrasted the network of a couple of cloud providers.

Conclusions

In this presentation, we have:

1. introduced a new technique to uncover topology that was based on Riemannian geometry.

2. built visualization that facilitates the utilization of our technique.

3. compared and contrasted the network of a couple of cloud providers.

Questions?