Network Curvature: Structure of the Internet

Rik Sarkar

Last class: Gossip Protocols

- The push-sum protocol
  - In every round
  - Every node takes a fraction of its value and sends to a random neighbor
  - It adds all received values to its current value
- The pairwise averaging protocol
  - In every round, a node talks to one other random neighbor
  - Both nodes set their values to the average of the two

Gossip averaging protocols

- On a complete graph
  - Both protocols converge to the average fast
  - O(log n) rounds
- On small world graphs/small world distributions
  - Convergence not known
- On a geometric graph (nodes connected to nearby nodes):
  - Convergence slow
- We will omit proofs.

Internet

- An interconnection network of “network of routers”
- Thousands of networks together form the Internet
- The “center” consists of big routers in highly connected networks, many connections between adjacent networks
- Outer layers have smaller routers and sparser connections

Internet

- Has a layered structure with higher connectivity at the core
    - A routed packet tends to use high connectivity regions to get shorter/faster routes
- Known to have power law distribution of degrees

Metrics

- A metric is a set of (shortest) distances (d) between points
- E.g. distances on a plane satisfy a Euclidean metric
- Distances on a sphere satisfy a spherical metric
- Distances between nodes on a grid graph satisfy a grid metric
- Nodes on a tree satisfy a tree metric

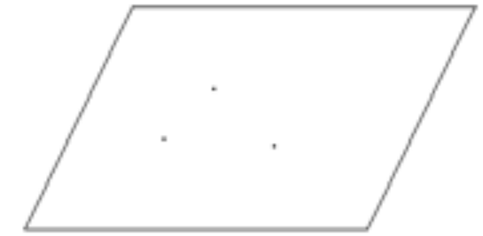

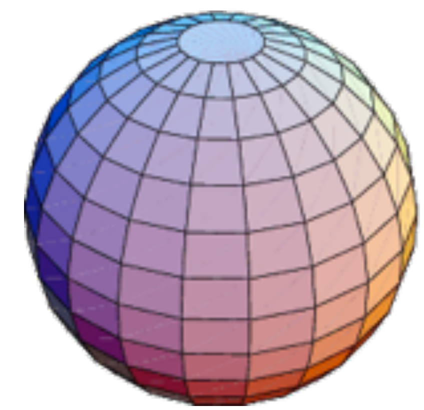

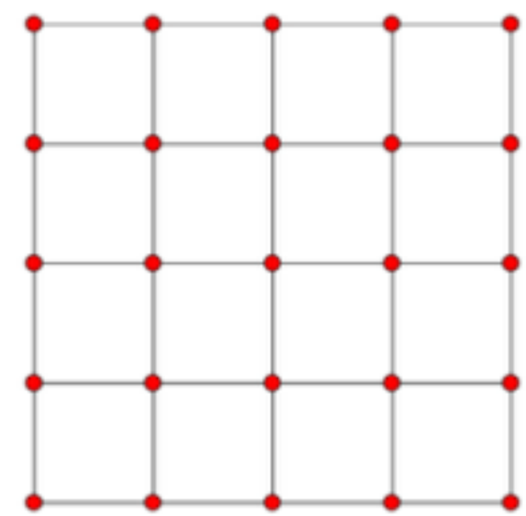

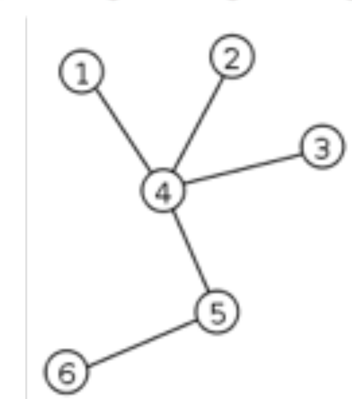

Comparing metrics

- We can say two metric spaces A and B are similar if one can be embedded in the other with small distortion

- That is, There is a function f A: -> B such that

    - $A(x,y) \sim B(f(x),f(y))$

A test for tree metrics

- A metric is a tree metric if and only if it satisfies this 4 Point Condition:
  - Any 4 nodes (points in the metric space) can be ordered as w,x,y,z such that:
  - $d(w,x) + d(y,z) \leq d(w,y) + d(x,z) \leq d(w,z) + d(x,y)$ and
  - $d(w,y) + d(x,z) = d(w,z) + d(x,y)$

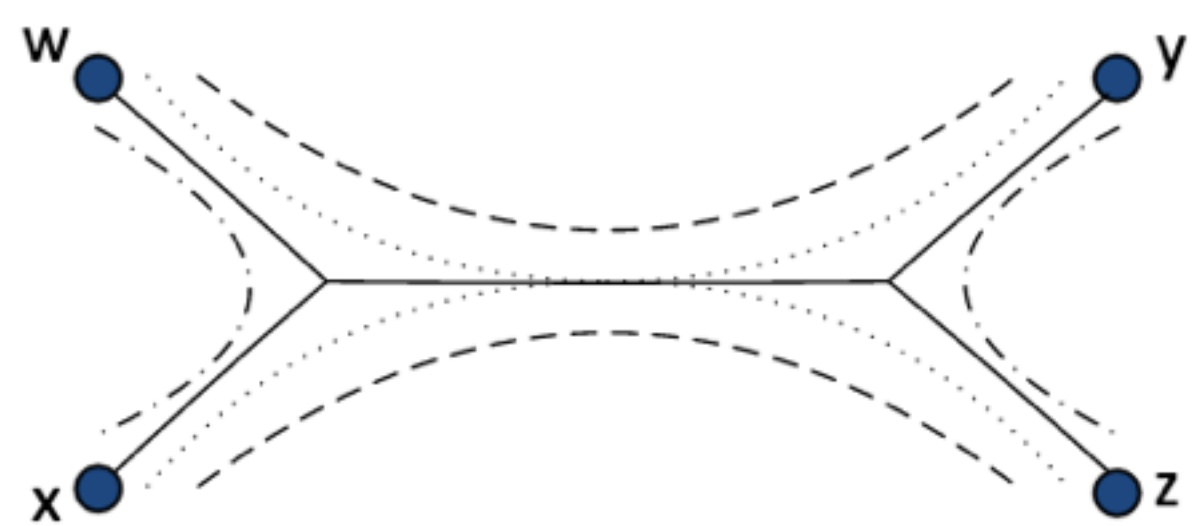

Trees tend to have high loads in “center”

- Since many routes will have to go through the center

Almost tree metrics

- Real networks are not exactly trees
- Let's measure how far a network is from a tree
- 4PC-ε for a set of 4 nodes is the smallest ε that satisfies:
  - $d(w,x) + d(y,z) \leq d(w,y) + d(x,z) \leq d(w,z) + d(x,y)$ and
  - $d(w,z) + d(x,y) \leq d(w,y) + d(x,z) + 2\varepsilon \cdot \min\{d(w,x), d(y,z)\}$

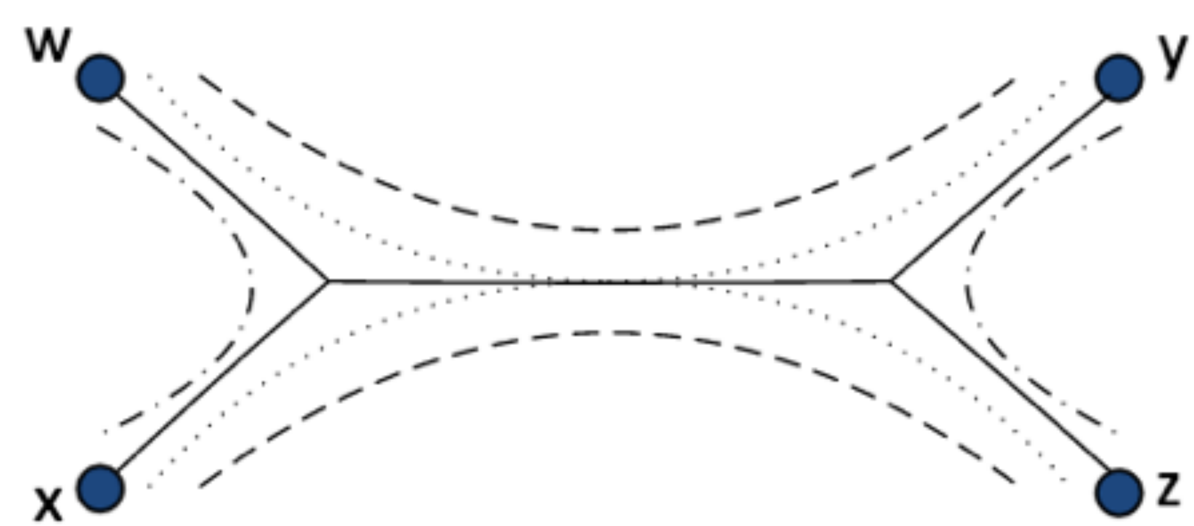

Almost tree metrics

- A tree has $\varepsilon = 0$
- A metric space with smaller ε implies that it is more similar to a tree
  - Theorem: A metric space with small ε can be embedded into a tree with correspondingly small distortion
  - Ref: I Abraham et al. Reconstructing approximate tree metrics, PODC 07.

Treeness of Internet

- PlanetLab: A distributed collection of servers around the world
- Experiment based on latency (communication delay) as an estimate of distance
- Shows the distance metric between servers is similar to a tree, and far from a sphere

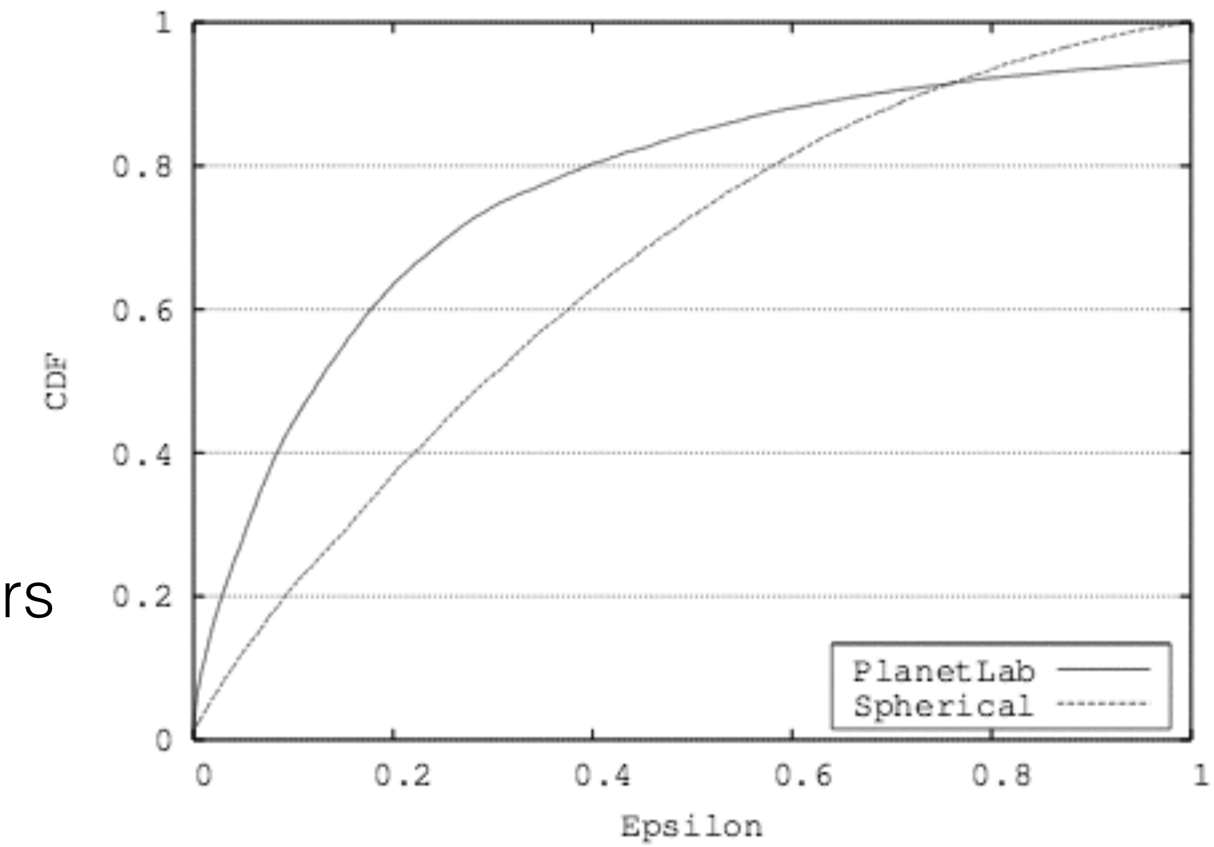

- Ref: I Abraham et al. Reconstructing approximate tree metrics, PODC 07.
- V. Ramasubramanian etal. On treeness of internet latency and bandwidth, Sigmetrics 09.

Course and Projects

- More office hours:
  - Thursday 19th, 10:00 - 11:30
- Remaining Classes:
  - Last usual class (some remaining material) Friday 20th
  - No class on Tuesday 24th
  - General review/discussion class on Friday 27th.

Treeness of metrics

- δ-hyperbolic metrics
- $d(w,x) + d(y,z) \leq d(w,y) + d(x,z) \leq d(w,z) + d(x,y)$ and
- $d(w,z) + d(x,y) = d(w,y) + d(x,z) + \delta$
- Uses an absolute value δ
  - Instead of a multiplicative factor

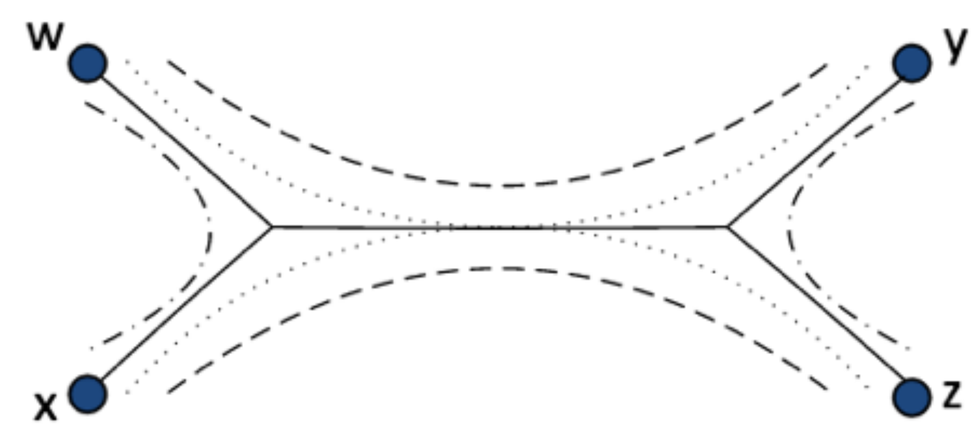

δ-hyperbolic metrics: Thin triangles

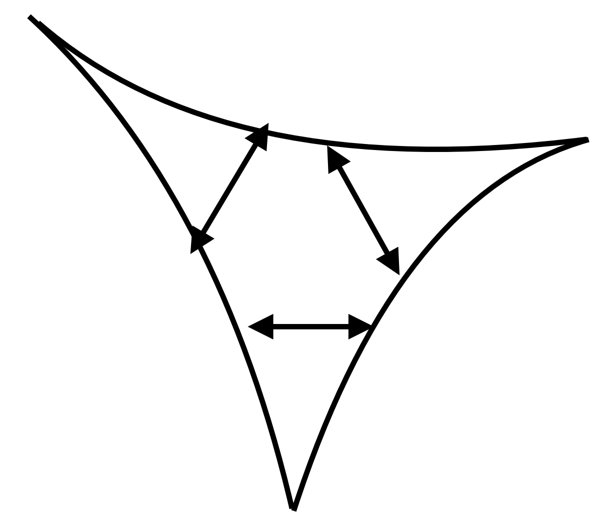

- Alternative definition
- Any point on a triangle must be within distance δ of one of the *other* sides
- The middle of the triangles are squeezed together
- trees have δ = 0: most hyperbolic

Curvatures of spaces

- Spherical : +ve curvature
  - Triangle centers are “Fat”
- Flat (Euclidean): 0 Curvature
- Hyperbolic: -ve curvature


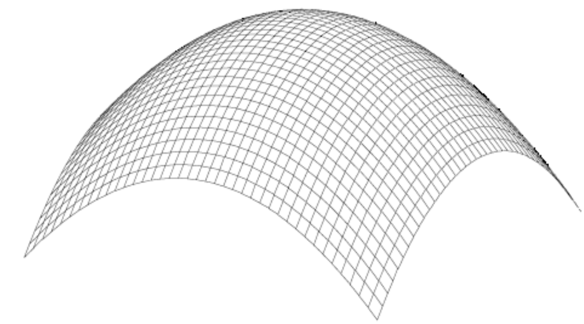
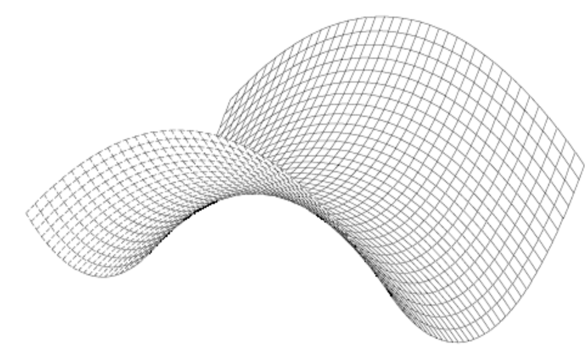

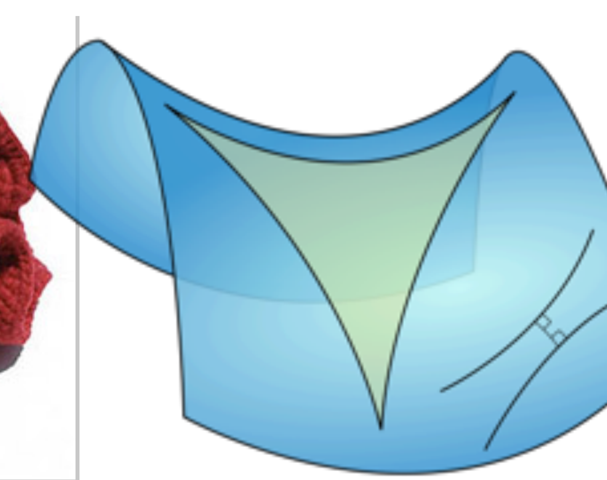

Angle properties

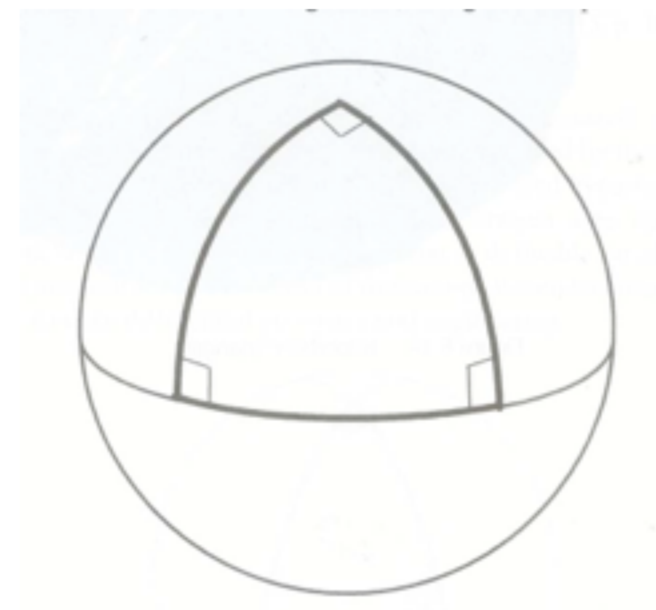

Figure 6.13 Triple-right triangle on a sphere

- Sum of angles of a triangle
- In Spherical space $\geq \pi$
- In flat space $= \pi$
- In hyperbolic space $\leq \pi$

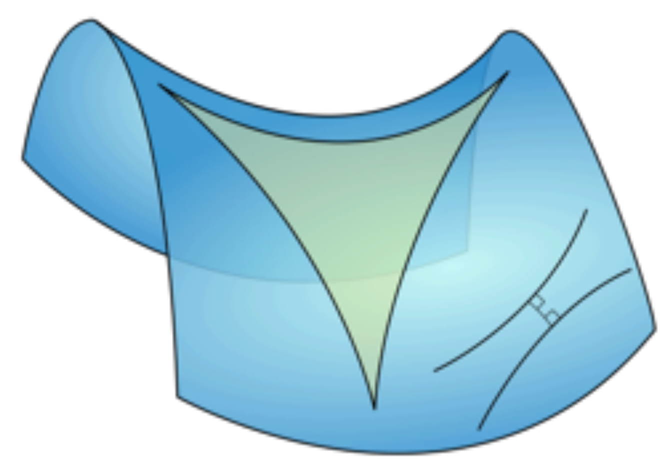

Growth properties

- Growth of a circle circumference with radius
  - Flat space: grows linearly: $2\pi r$
  - Spherical: grows sub-linearly: $\leq 2\pi r$
    - In fact, as less than cr for any const c
    - Shrinks after covering half the sphere
  - Hyperbolic: Grows super-linearly: $\geq 2\pi r$
    - In fact, exponentially as $\sinh(r) \sim r^e$
    - Does not fit on the plane or even 3D

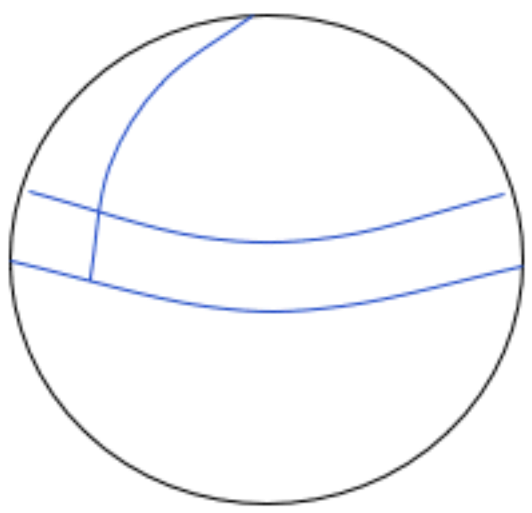

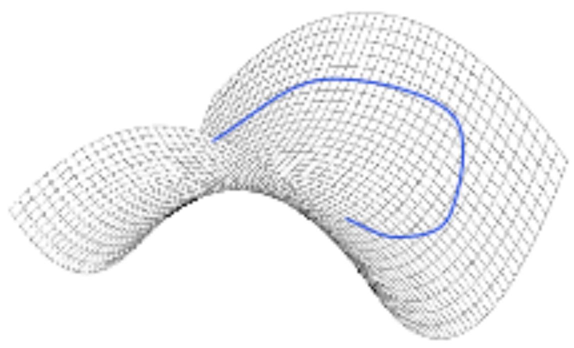

- Any hyperbolic space is δ-hyperbolic for some finite delta
  - Not the case for Euclidean and spherical spaces
  - For more on δ-hyperbolic spaces, See: Gromov hypoerbolic spaces

- Any tree can be embedded in a hyperbolic space with a low distortion

- R. Sarkar. “Low distortion delaunay embedding of trees in hyperbolic plane.” GD 2011.

Internet

- Internet has good embedding in hyperbolic spaces

- Ref. Shavitt and Tankel 2008, Narayan and Saniee 2011

Modeling metrics

- We need a set of points
- A notion of “straight line distances”

- On euclidean plane, straight lines are straight lines
- On a sphere, straight lines are “great circles”

A model for hyperbolic metrics

- The Poincare model
  - Take a disk
  - Straight lines are circular arcs that meet the boundary at right angles
  - Note that there can be many more parallel lines than in flat space

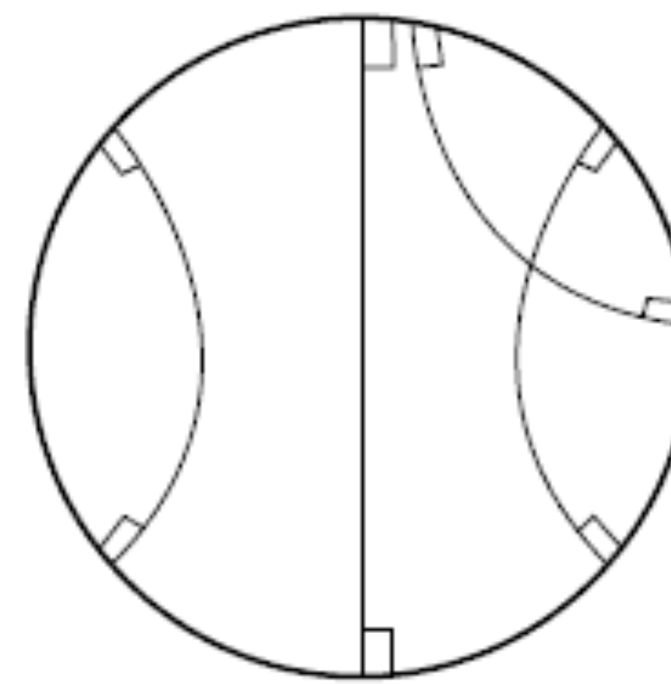

The internet structure

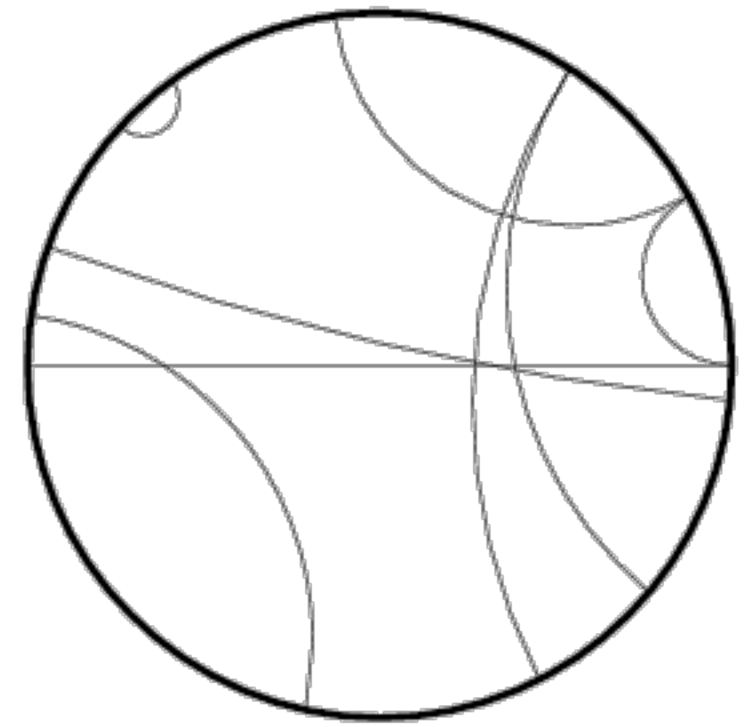

- Observe:
  - Shortest paths tend to come close to the center and then move away
  - Explains “internet core”

Visualizing hyperbolic metrics

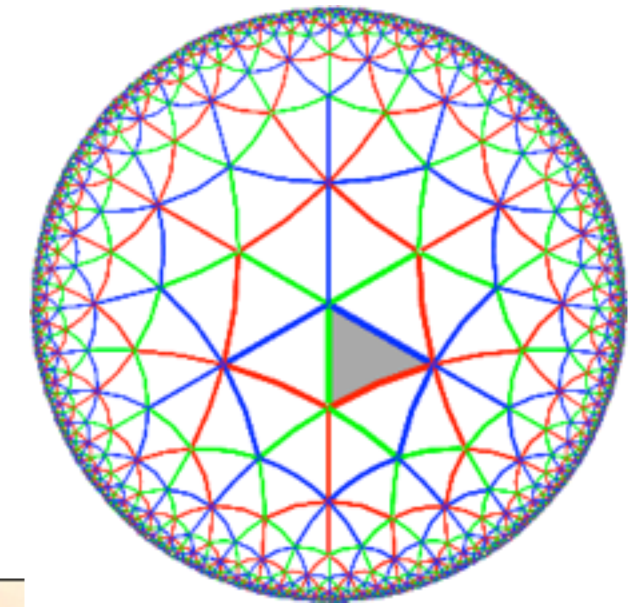

- All elements are same size
- Things close to the boundary “look” small
    - Perspective
    - When drawing an infinite surface, some things have to be compressed
- See Escher’s paintings

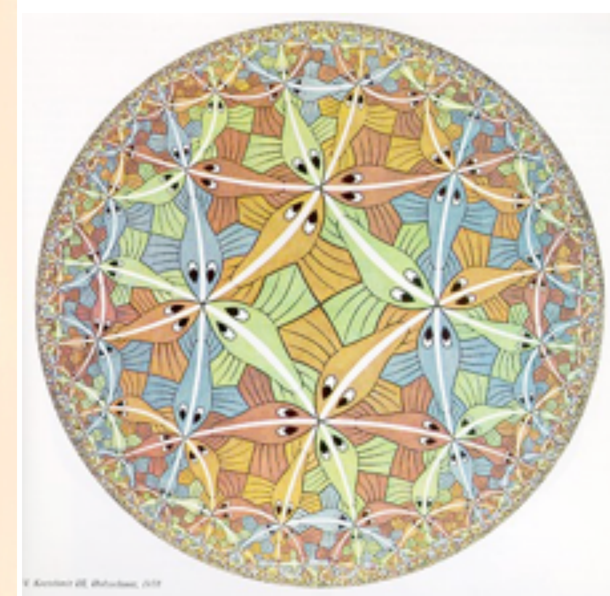

A model for hyperbolic metrics