Geohyperbolic Routing and Addressing Schemes

Ivan Voitalov, Rodrigo Aldecoa,
Lan Wang, Dmitri Krioukov

NDNcomm, March 2017

dk lab

netsi
Northeastern University
Network Science Institute

networking research lab
COMPUTER SCIENCE · UNIVERSITY OF MEMPHIS

Routing Scalability — Why Do We Care?

- Best forwarding strategy: zero routing overhead and smallest FIB size.
- BGP: linear growth of FIBs, unbounded routing overhead (persistent routing oscillations).
- Data-centric paradigms require immense addressing spaces.
- New approach to routing is needed for future networks.

Greedy Geometric Routing

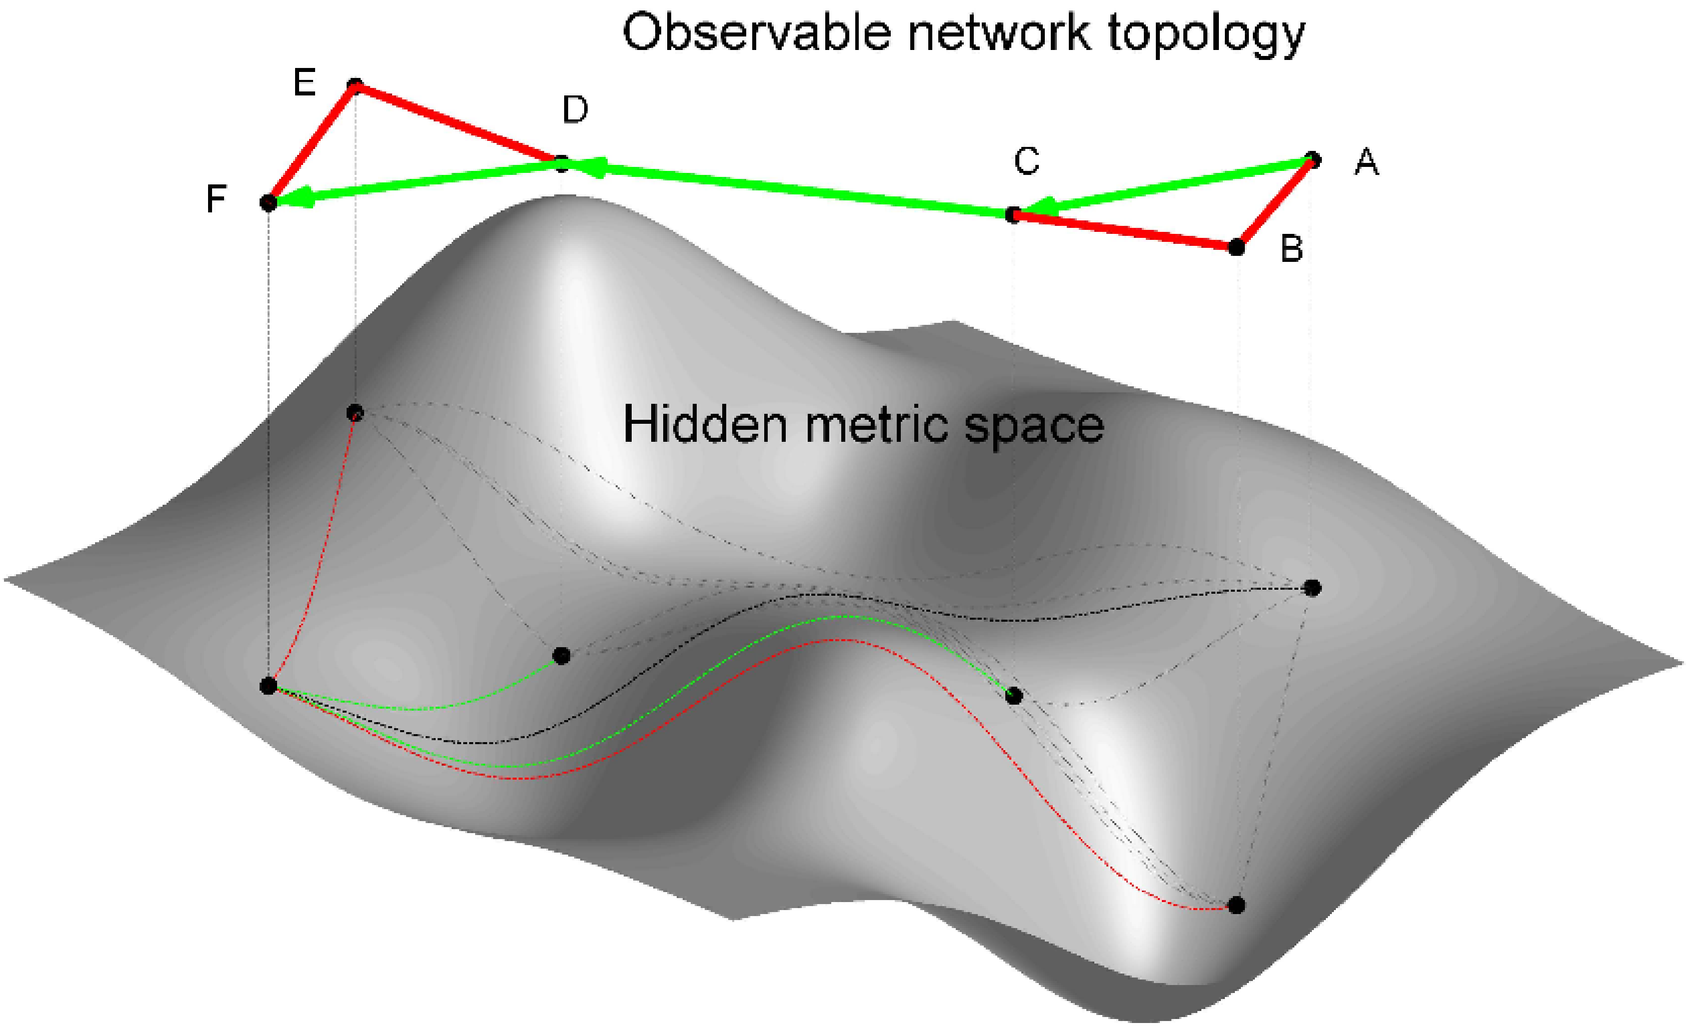

From Boguna et al. “Navigability of Complex Networks” (2009)

Greedy *Hyperbolic* Routing

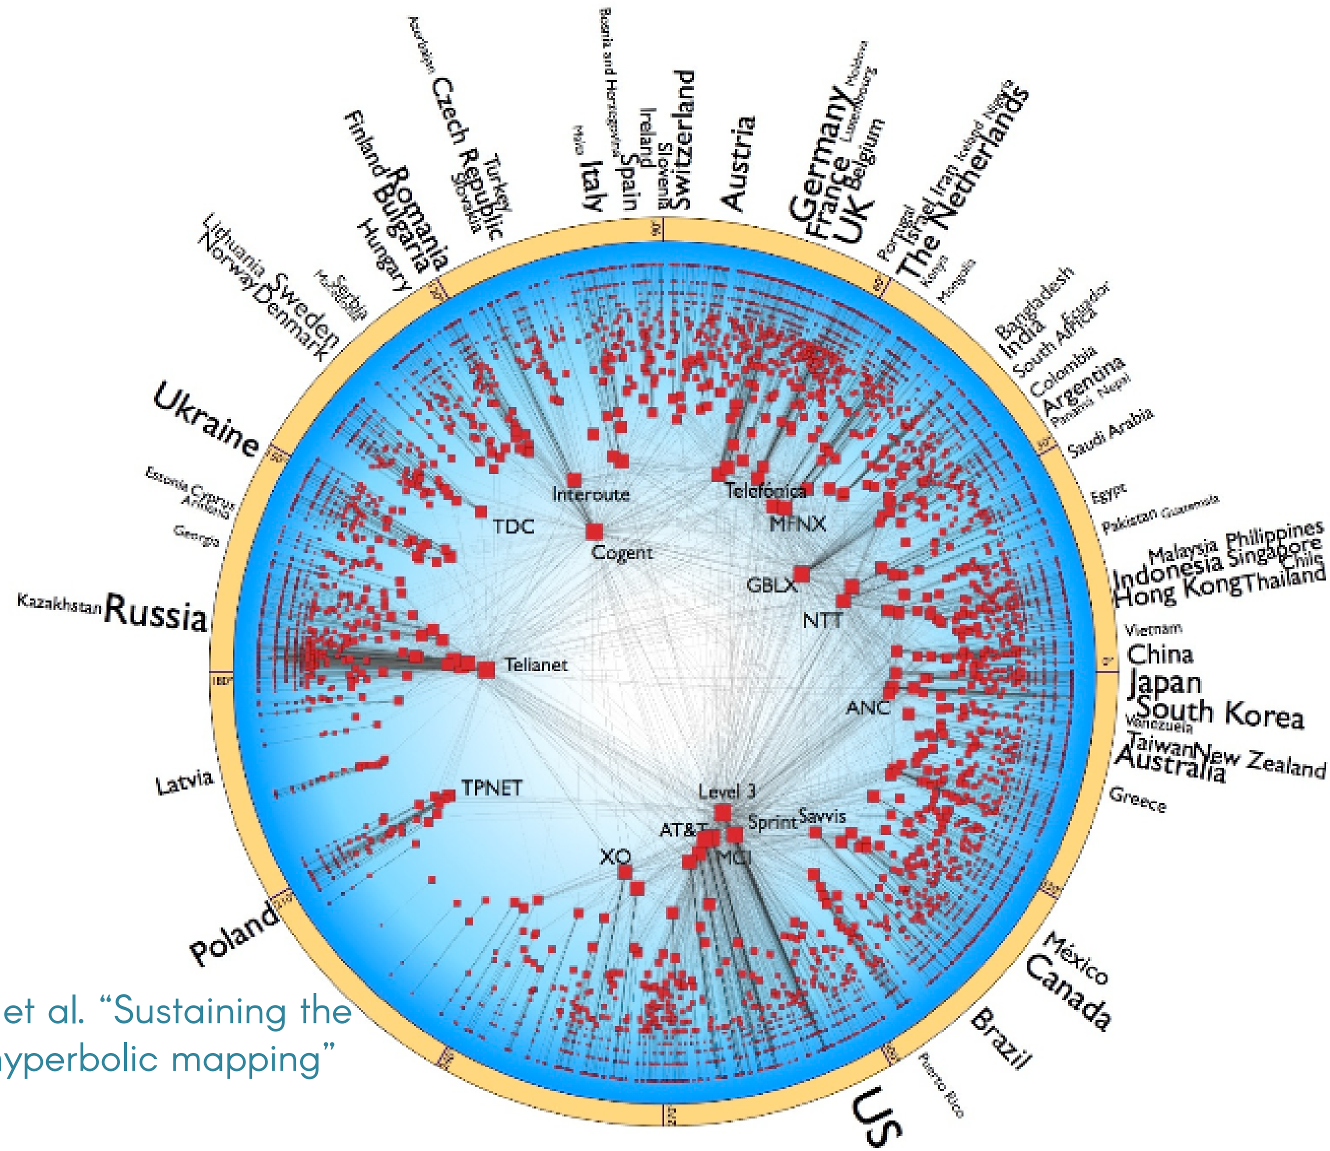

From Boguna et al. "Sustaining the Internet with hyperbolic mapping" (2010)

Greedy *Hyperbolic* Routing

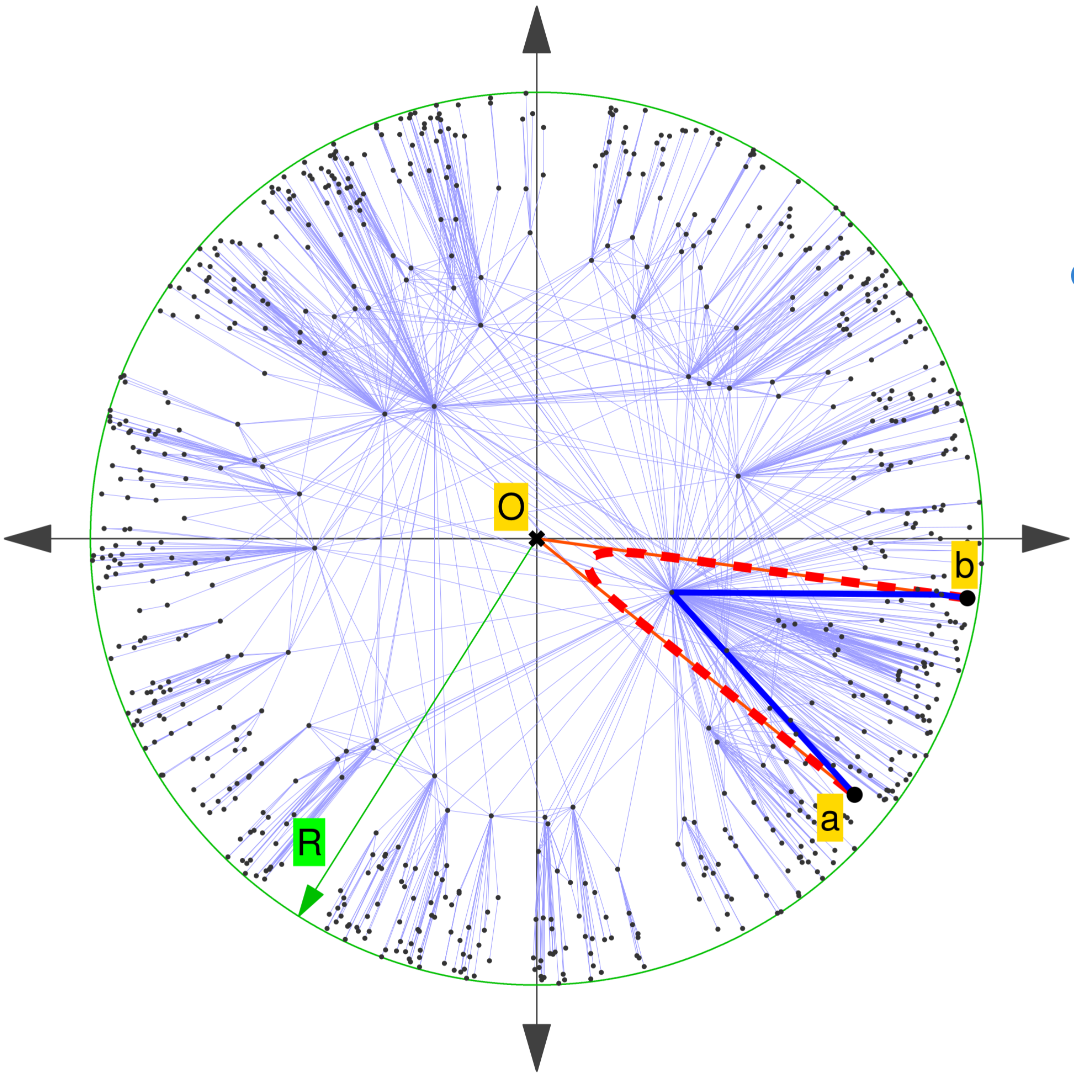

Geodesic lines (shortest distances) in hyperbolic space are congruent with shortest paths in network embedded in this hyperbolic space

From Boguna et al. "Sustaining the Internet with hyperbolic mapping" (2010)

Greedy *Hyperbolic* Routing

- It's not easy to make such a map: need to know substantial part of network topology; coordinate computations are not straightforward.
- Such mapping is not aware of time delays that affect user experience.

What to do?

- Include delay information in network design: geodistance is related to round trip time (RTT).
- Map nodes to the underlying space “on-the-go”, as the network grows.
- Network should evolve in a specific manner.

What to do?

- General network construction scheme: a node arrives in the network and gets an address — a set of coordinates in the underlying space.
- A node connects to *m* geometrically closest neighbors, as defined by the distance function of the underlying space.
- Greedy routing uses the same distance function to forward packets.
- That's it!

Geographic Scheme (GEO)

- Start with the simplest model: use latitude and longitude of a node to assign an address and connect new nodes to *m* geographically closest neighbors.
- Resulting underlying space — two-dimensional sphere.
- Let's consider an example with $m = 1$.

Geographic Scheme (GEO)

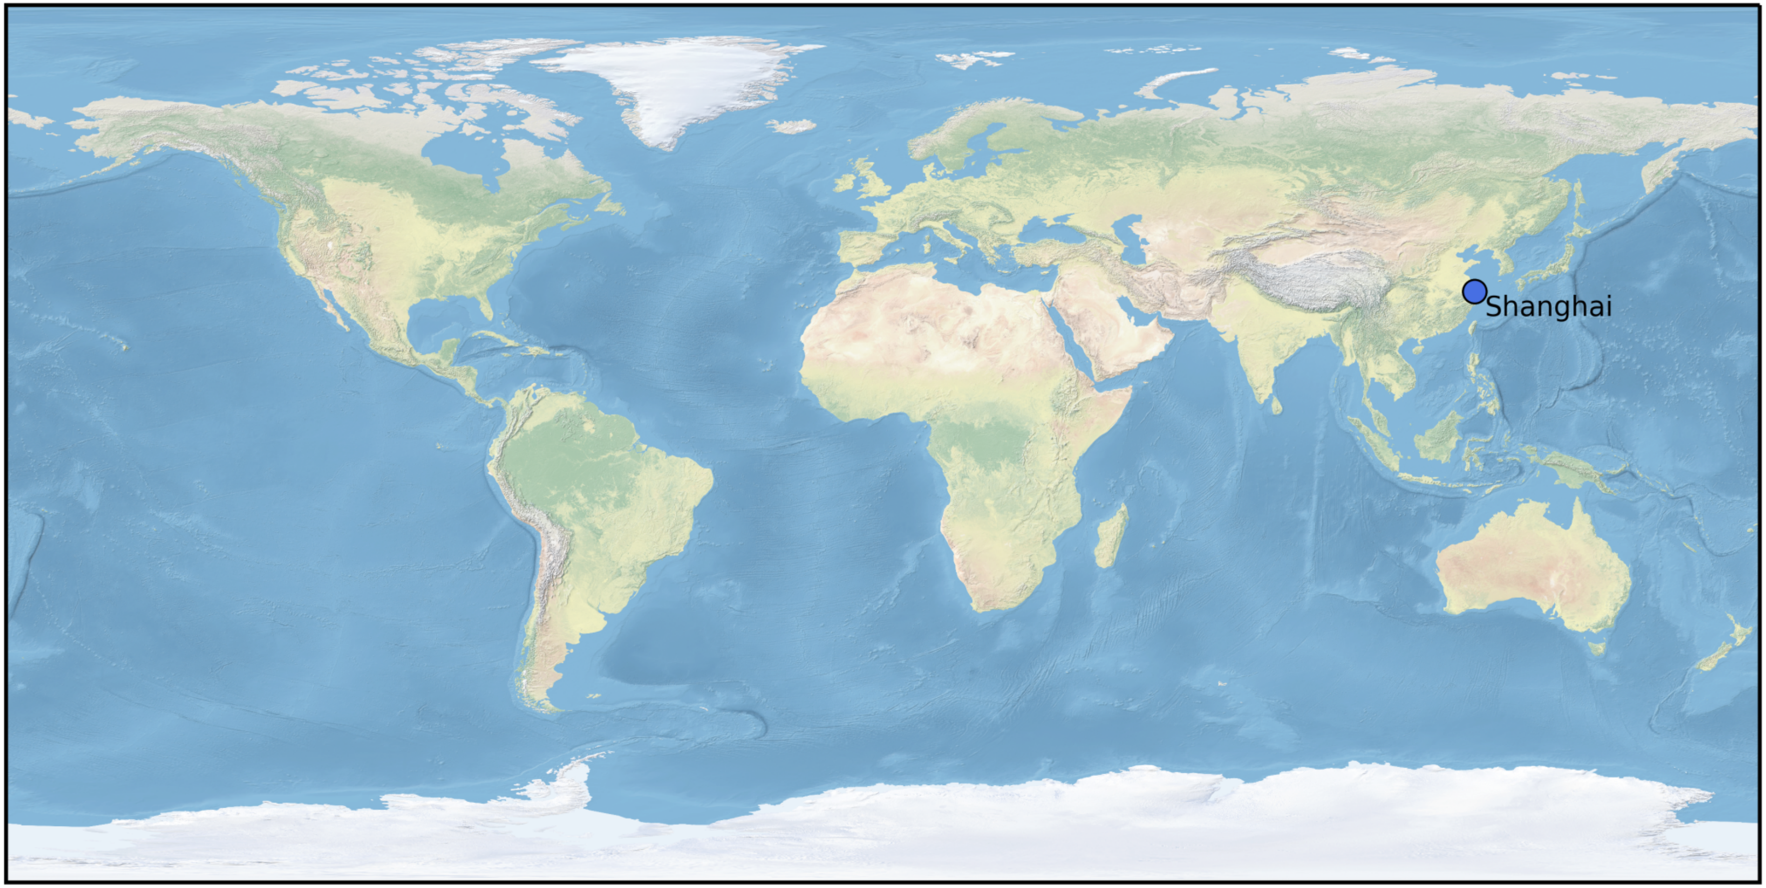

Geographic Scheme (GEO)

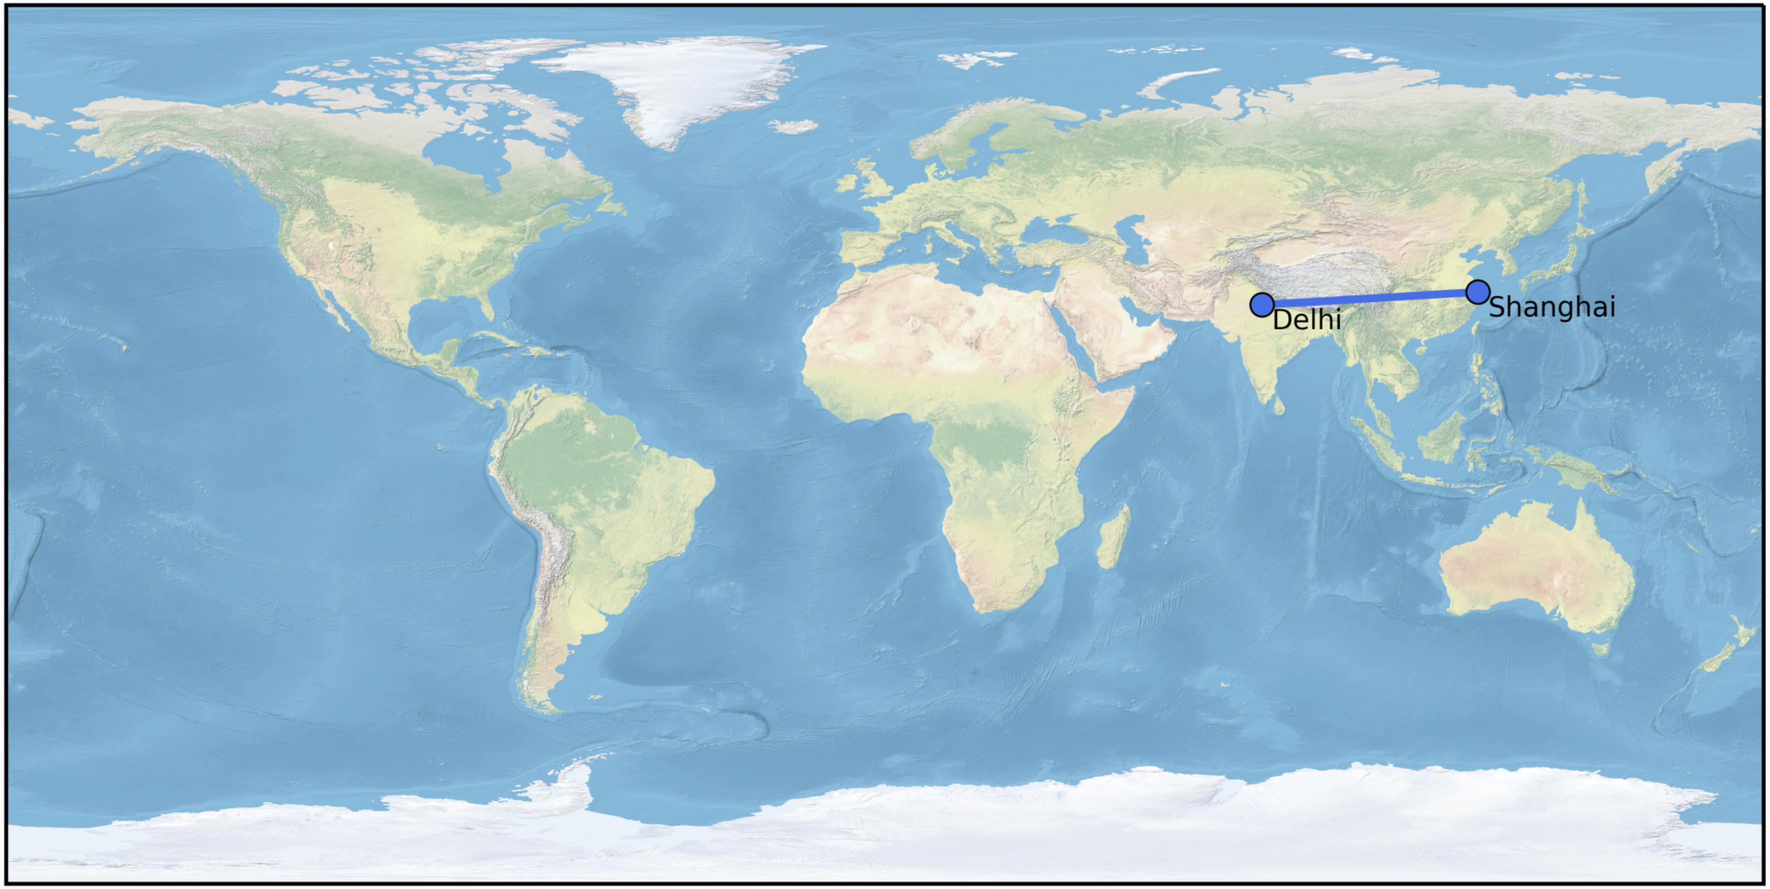

Geographic Scheme (GEO)

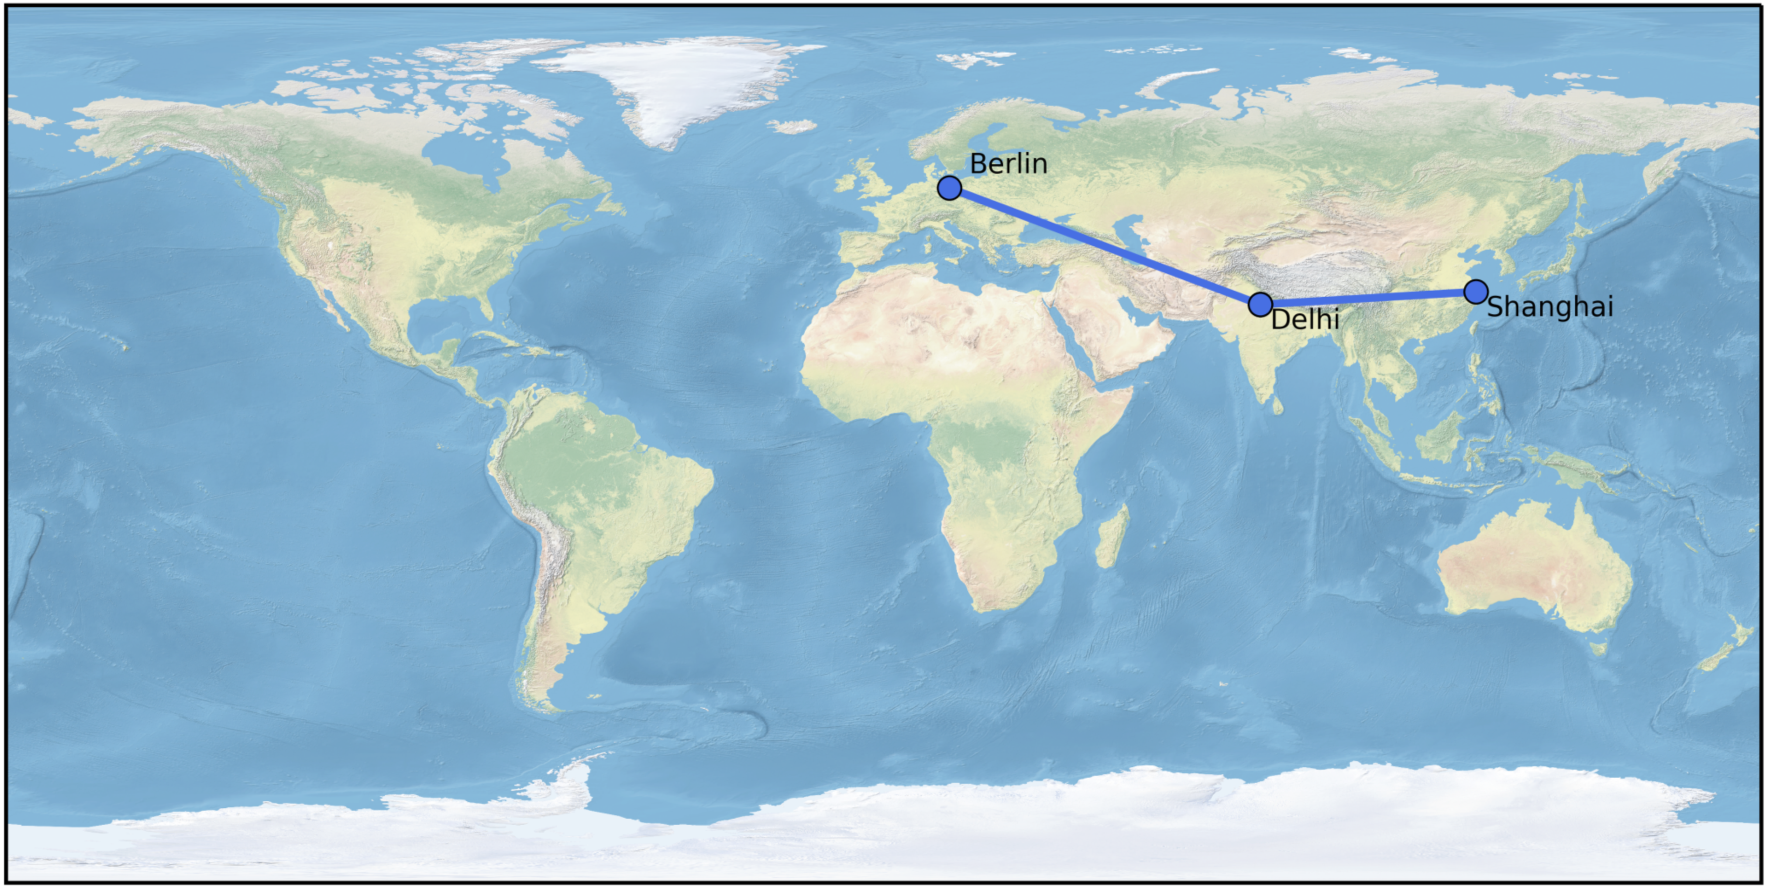

Geographic Scheme (GEO)

- This simple scheme does not guarantee routing scalability: success ratio degrades quickly with network size, demanding for auxiliary forwarding algorithms.
- The resulting networks do not have robust topology: they resemble random geometric graphs on sphere, i.e. graphs with narrow degree distribution.

Geohyperbolic Scheme (GH)

- Modification to the last scheme: add third, **centrality**, coordinate to a node's address that captures how "central" a node is in the network, i.e. how likely new nodes are to establish their connections with this node.
- Assign node's *radial coordinate* as a ~log(centrality).
- If centrality scores have heavy-tailed distribution, this assignment of coordinates naturally maps nodes to 3-dimensional *hyperbolic space.*
- Use hyperbolic distance instead of geographic distance to establish network links and perform GGR.

Geohyperbolic Scheme (GH)

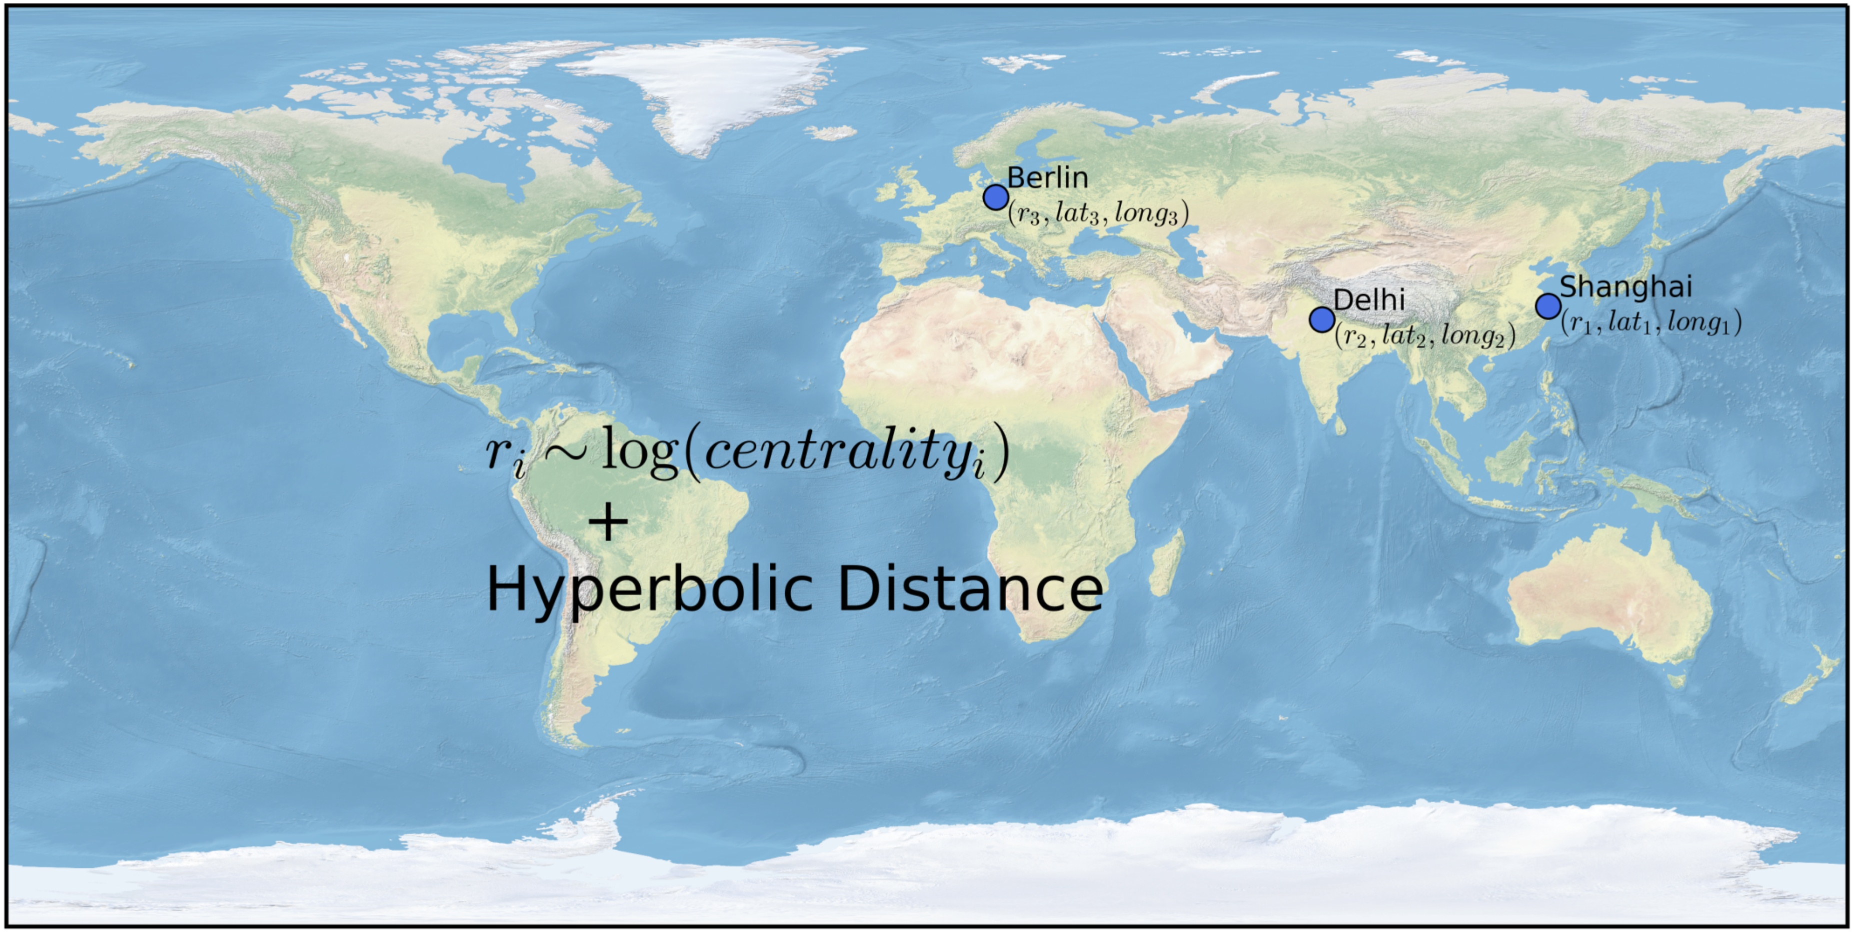

Geohyperbolic Scheme (GH)

- Centrality scores can be chosen according to certain needs.
- Any geographic location is properly mapped to 3-dimensional hyperbolic space.
- Resulting networks are robust to random failures.
- Success ratio of GGR is almost 1 for any network size and under severe connectivity failures.
- However, suboptimal delay-wise performance is observed.

Geohyperbolic Scheme (GH)

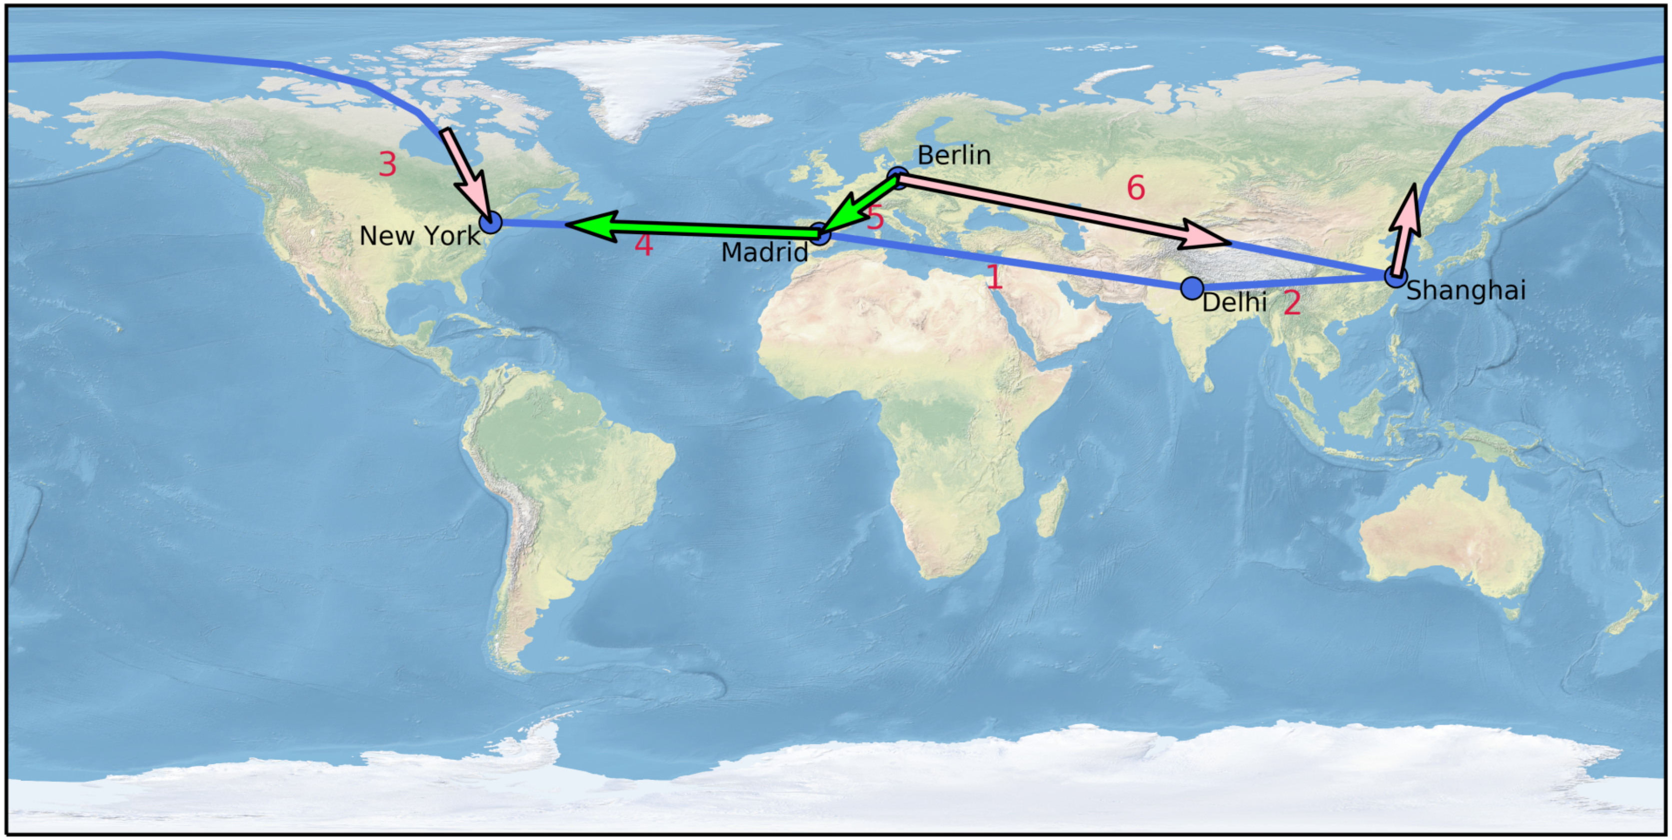

- Bad delay performance example: packet forwarding from Berlin to New York via "super central" node in Shanghai.

Regionalized Geohyperbolic Scheme (RGH)

- Small tweak of a previous scheme: place multiple "local hubs" within large geographic regions to "attract" greedily forwarded packets from peripheral local nodes.
- Local hubs have the same radial coordinates, i.e., greedily forwarded packets are attracted to *geographically closest* local hubs, which reduces long-delay paths.

Regionalized Geohyperbolic Scheme (RGH)

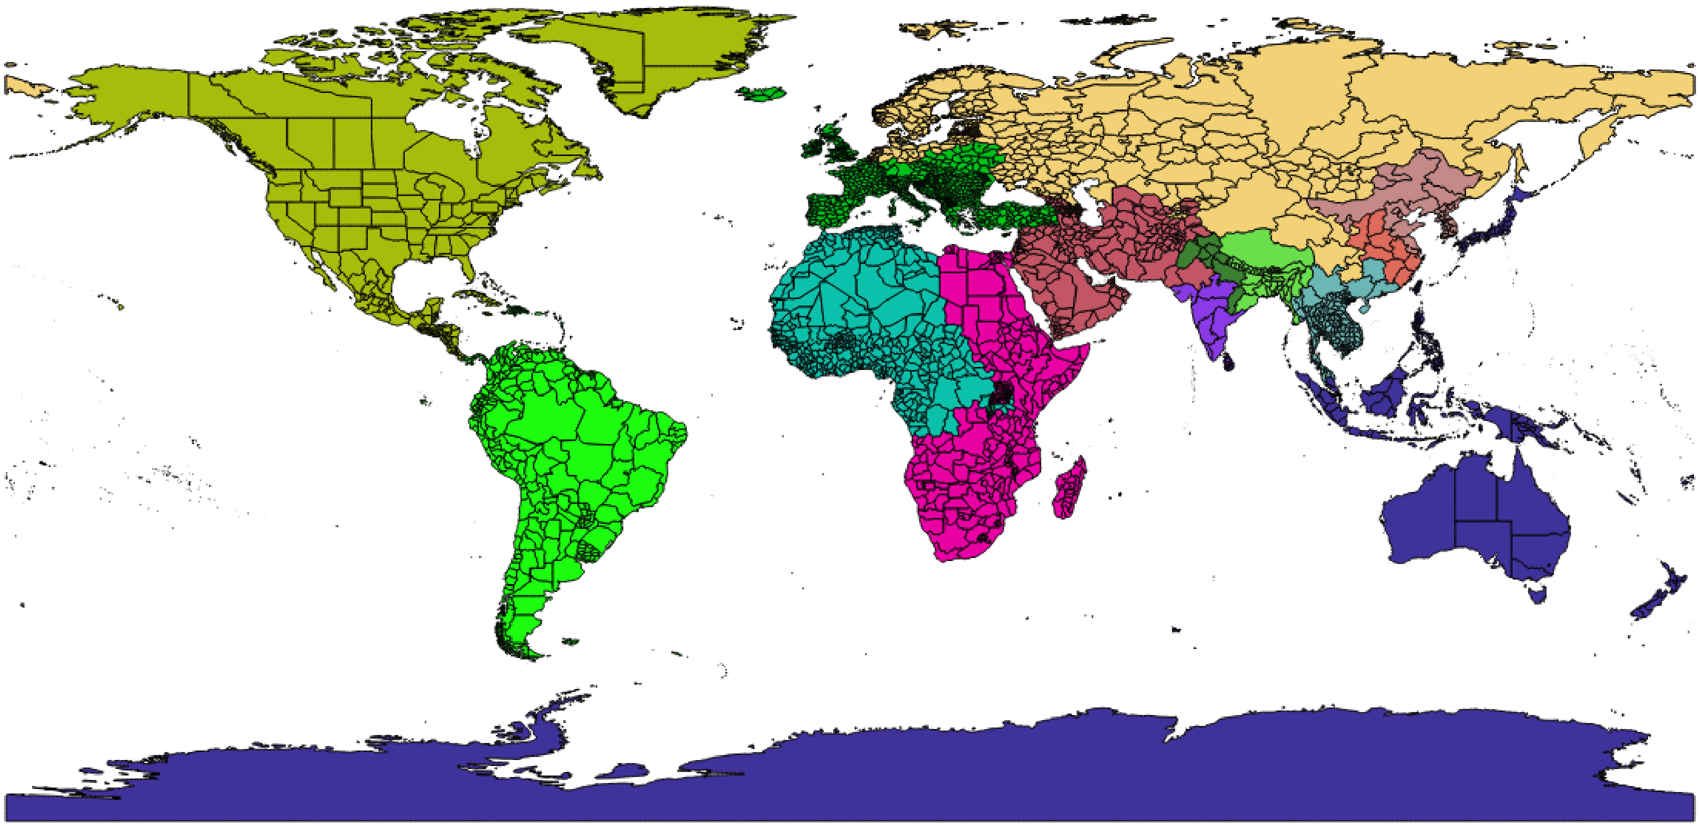

Administrative Level 1 Units (e.g., states or provinces) are merged into large regions equipped with local hubs.

Validation

- Tested GEO, GH and RGH schemes both in large synthetic networks and NDN testbed simulations.
- Nodes appearing in cities, towns, etc.
- Centrality score of a node is set to population of corresponding populated place:

$$r \sim \log(\text{population rank})$$

- Randomized order of nodes' arrivals in RGH scheme mimics real-world situation: central nodes are more likely to appear in densely populated places.
- Tested connectivity disruptions: 20% of links were randomly removed and greedy routing is tested again.

Validation

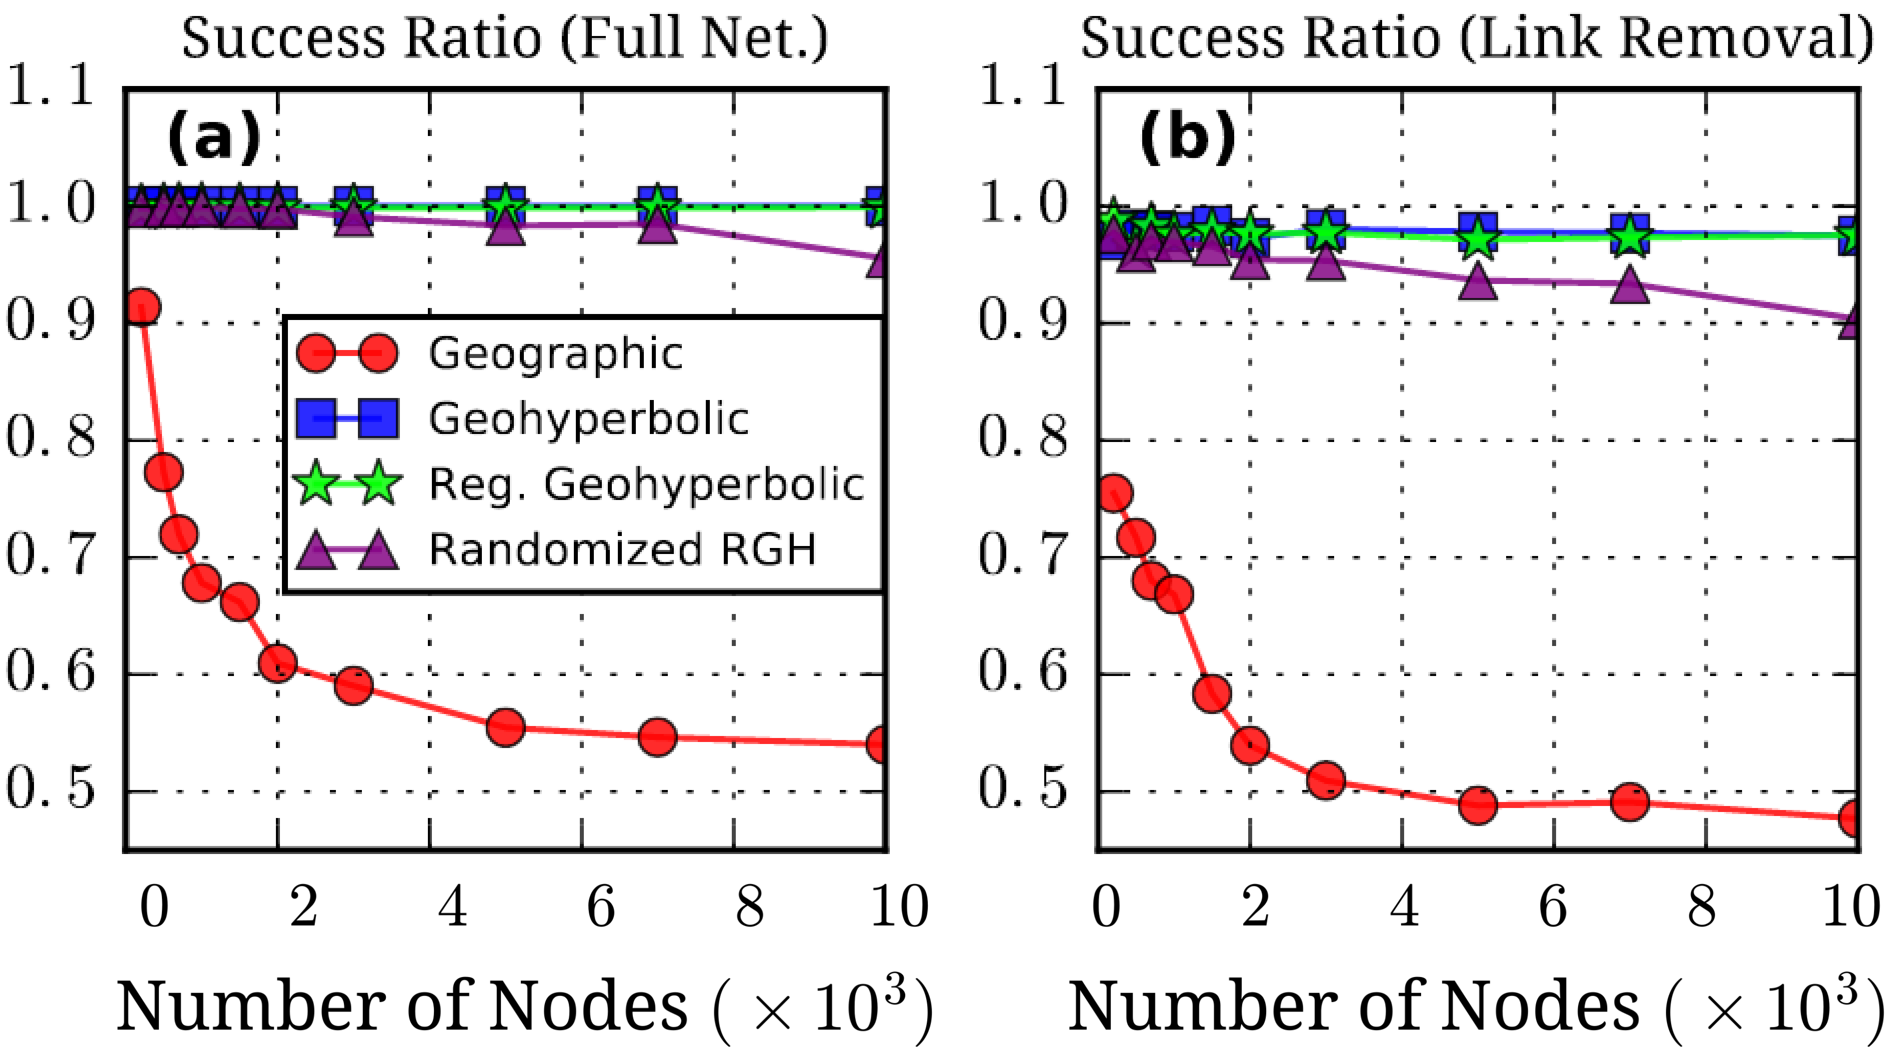
Success Ratio (Full Net.)
Success Ratio (Link Removal)
(a)
(b)
1.1
1.0
0.9
0.8
0.7
0.6
0.5
0
2
4
6
8
10
Geographic
Geohyperbolic
Reg. Geohyperbolic
Randomized RGH
Number of Nodes (× 10^3)
Number of Nodes (× 10^3)

Delay Performance Metrics

Underlay delay stretch (UDS):

greedy path delay / direct (underlay) delay

Overlay delay stretch (ODS):

Greedy path delay / Dijkstra path overlay delay

Validation

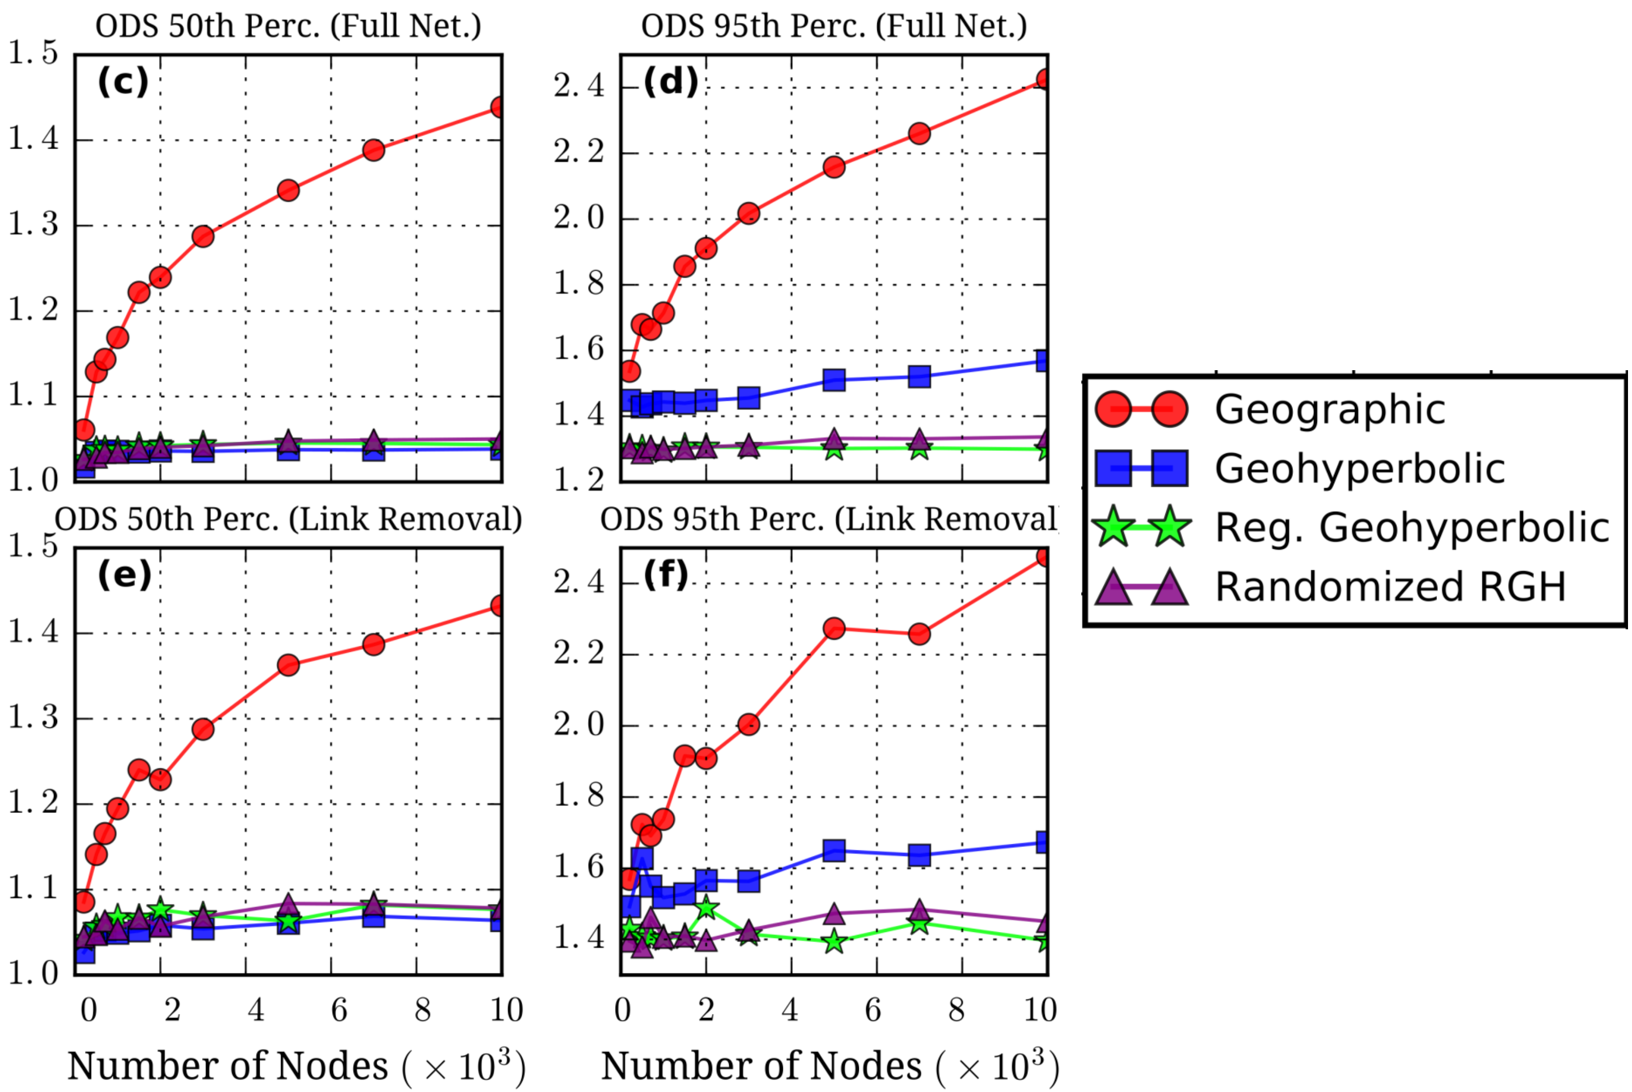

ODS 50th Perc. (Full Net.)
(c)
ODS 95th Perc. (Full Net.)
(d)
ODS 50th Perc. (Link Removal)
(e)
ODS 95th Perc. (Link Removal
(f)
1.0
1.1
1.2
1.3
1.4
1.5
1.2
1.4
1.6
1.8
2.0
2.2
2.4
0
2
4
6
8
10
Number of Nodes ($\times 10^3$)
Geographic
Geohyperbolic
Reg. Geohyperbolic
Randomized RGH

Validation

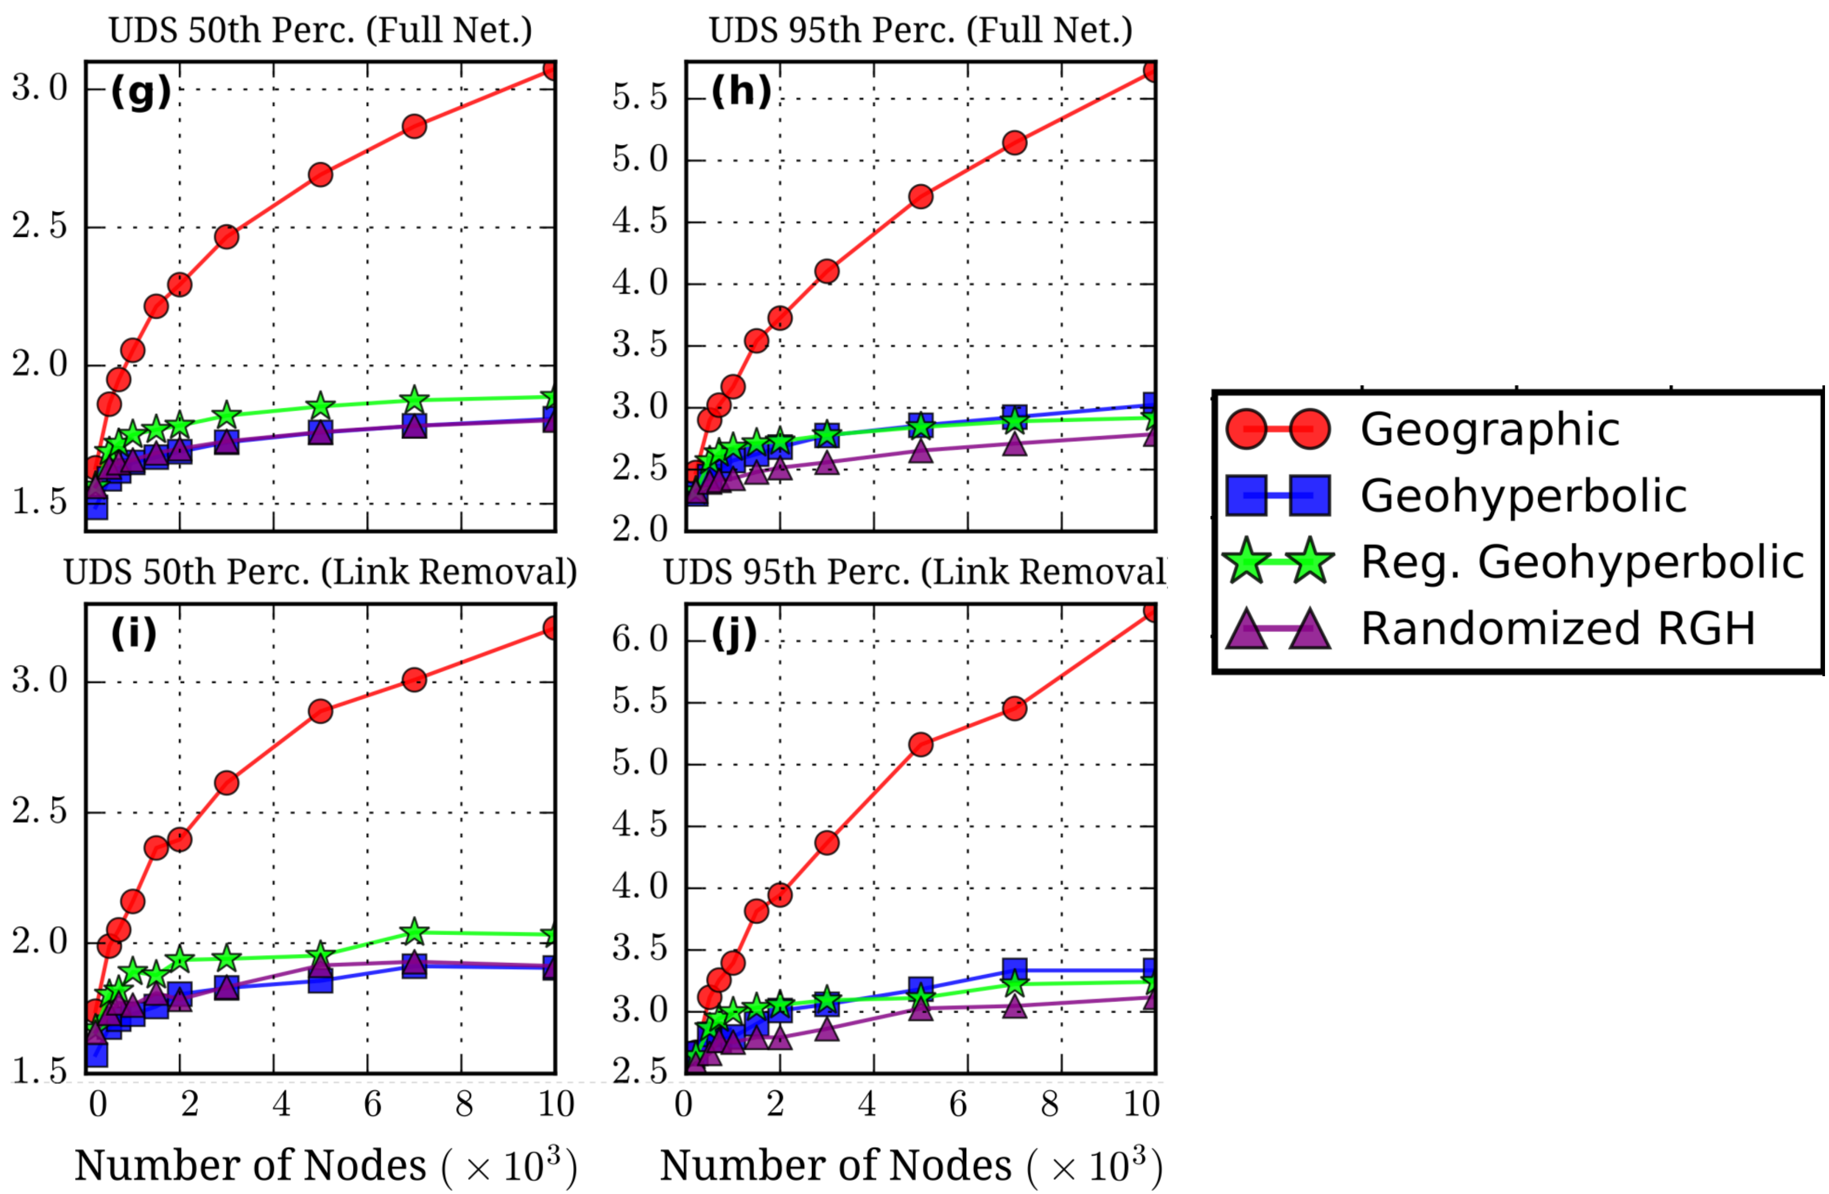
UDS 50th Perc. (Full Net.)
(g)
UDS 95th Perc. (Full Net.)
(h)
UDS 50th Perc. (Link Removal)
(i)
UDS 95th Perc. (Link Removal)
(j)
Number of Nodes (× 10^3)
Number of Nodes (× 10^3)
Geographic
Geohyperbolic
Reg. Geohyperbolic
Randomized RGH

Validation

- NDN testbed: used 29 functional nodes and links according to three schemes.
- Best SR performance: GH and RGH.
- Under 20% links damage SR is still high: 0.98 and 0.97.
- Median ODS: 1 in all cases.
- 95-th percentile ODS: 1.54 for RGH, 2.65 for GH (regionalization helps!)

Conclusion

- GEO scheme does not offer scalable solution.
- SR is almost independent of network size for GH and RGH.
- SR is high even under severe connectivity failures.
- Delay stretch is bounded.
- Only small fraction of paths should be optimized via auxiliary forwarding algorithms!

Conclusion

- Proposed network design offers light-weight and scalable routing solution.
- Only geolocation of nodes are used: no virtual coordinates now!
- Flexible centrality scheme allows to implement such networks in different circumstances.
- Future work: mobile nodes?

Thank you!