




Sequoia: Virtual-Tree Models for Internet Path Metrics

Rama and Dahlia, MSR SVC

Goals



- ▶ Study “treeness” of Internet path metrics such as bandwidth and latency
- ▶ Construct tree models that represent network metrics
- ▶ Provide network-aware functionalities:
 - ▶ bandwidth and latency estimation
 - ▶ selection of “best” servers
 - ▶ topological clustering of peers




Applications

- ▶ *“what’s the server with the largest bandwidth that a client can download content from?”*
 - ▶ *Content distribution*
- ▶ *“what’s the closest peer that a host in a P2P network can connect to?”*
 - ▶ *P2P Networking*
- ▶ *“what’s the best game server to coordinate the online game between a set of players?”*
 - ▶ *Online gaming*

Approach: Metric Embedding into Trees




- ▶ End hosts (A, B, C, R) are leaf nodes



Approach: Metric Embedding into Trees




- ▶ Leaf nodes (A, B, C, R) are end hosts
- ▶ Inner nodes (s, t) are “virtual”



Approach: Metric Embedding into Trees



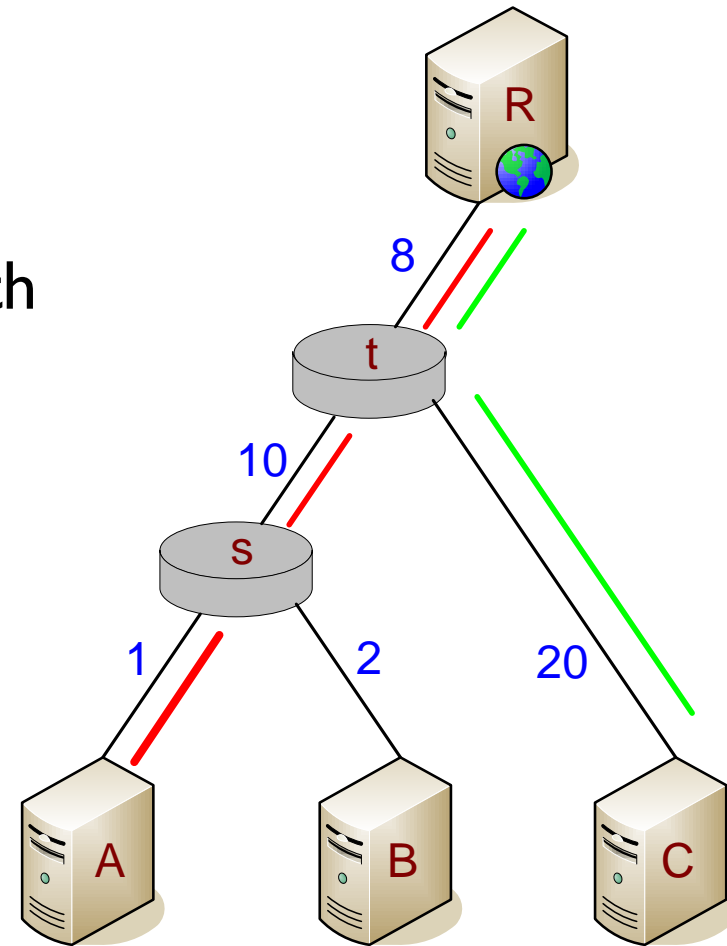
- ▶ Leaf nodes (A, B, C, R) are end hosts
- ▶ Inner nodes (s, t) are “virtual”
- ▶ Edge weights model network measure



Approach: Sequoia and “Coordinates”

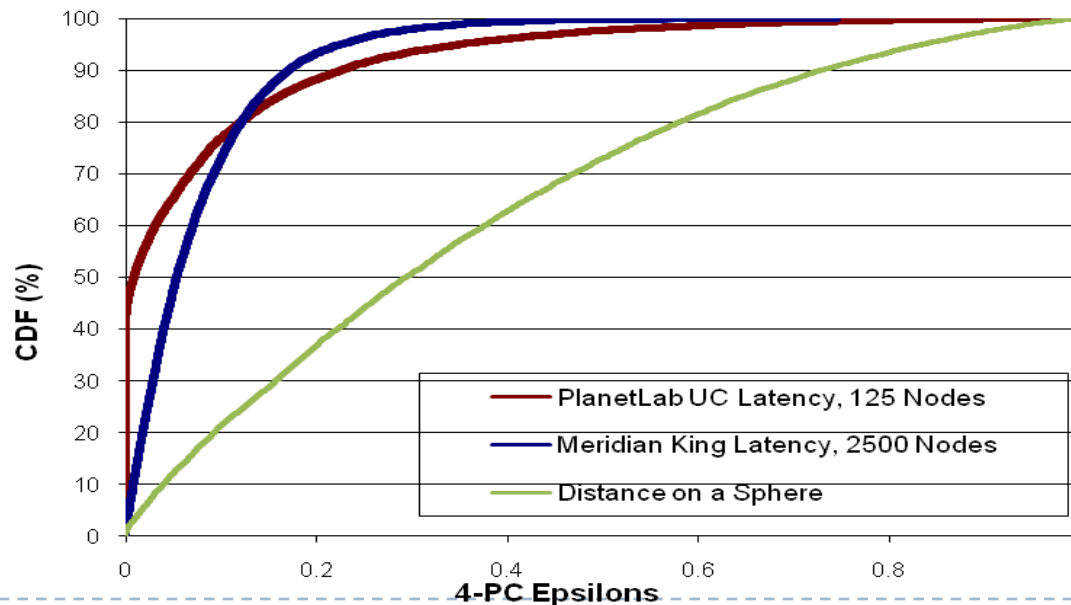


- ▶ Distance Label = path to the root
 - ▶ Example: **A: (s,t,R)** and **C: (t,R)**
- ▶ Estimated Metric = distance on path
 - ▶ Latency: $d(A,C) = d(A,s) + d(s,t) + d(t,C)$
- ▶ Same convenience as network coordinates



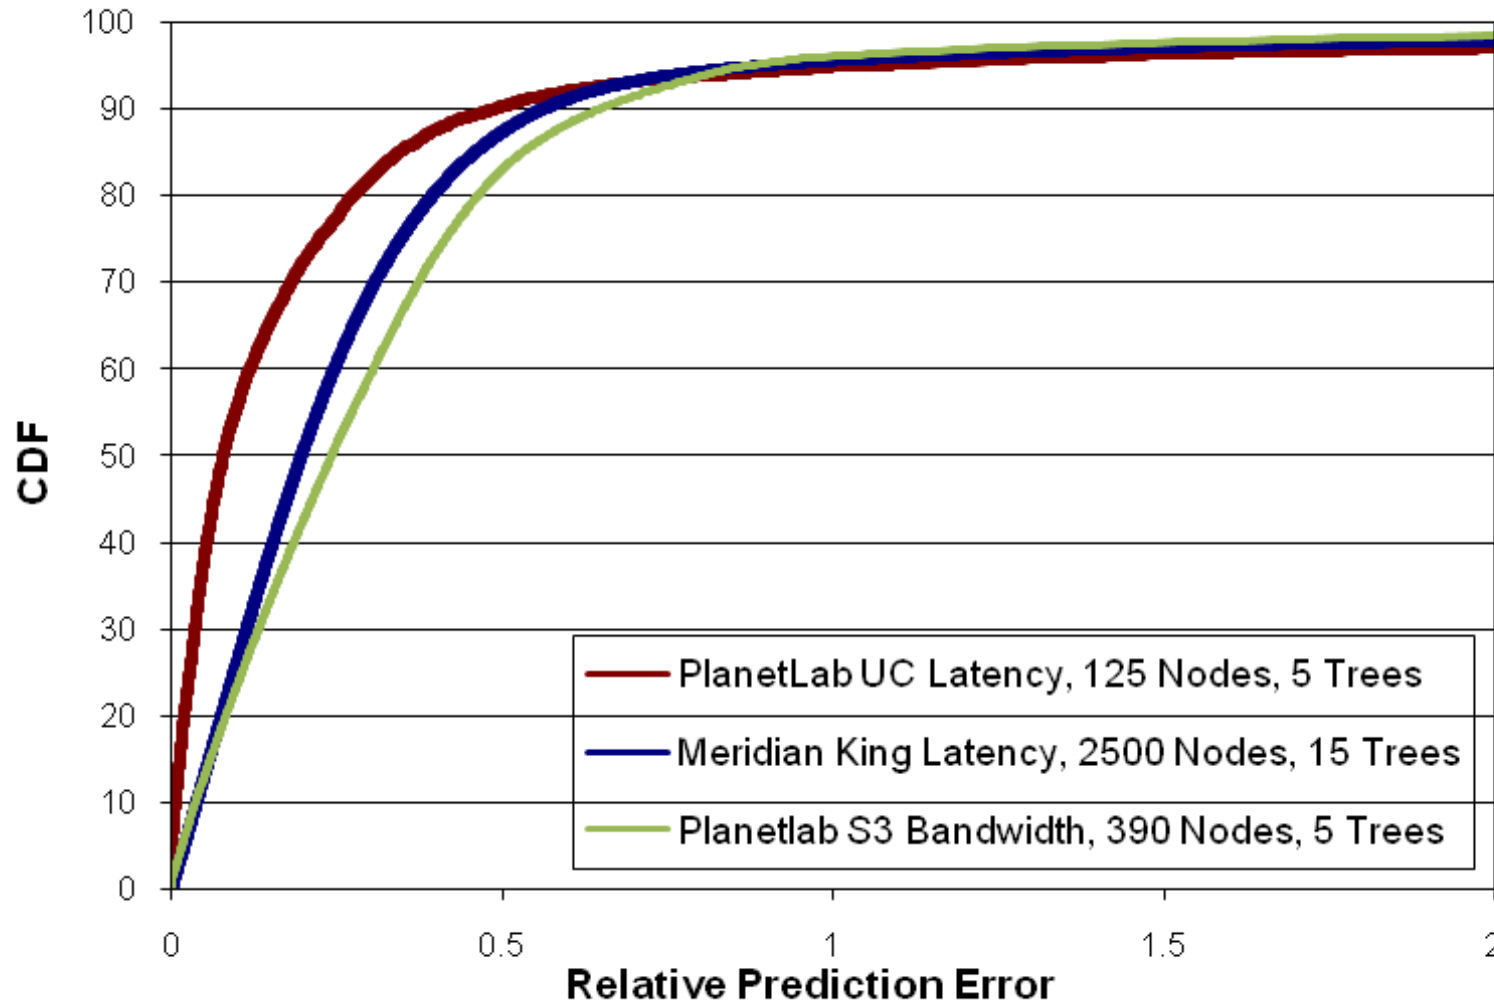
Intuition

- ▶ Internet is inherently hierarchical
 - ⇒ end-to-end network measures are approximate tree metrics
- ▶ bottleneck-link bandwidth along the highest bandwidth path is a tree metric
 - ⇒ bandwidth in a network bottlenecked at the edges is a tree metric



Results:


Bandwidth/Latency Estimation




Results: Server Selection



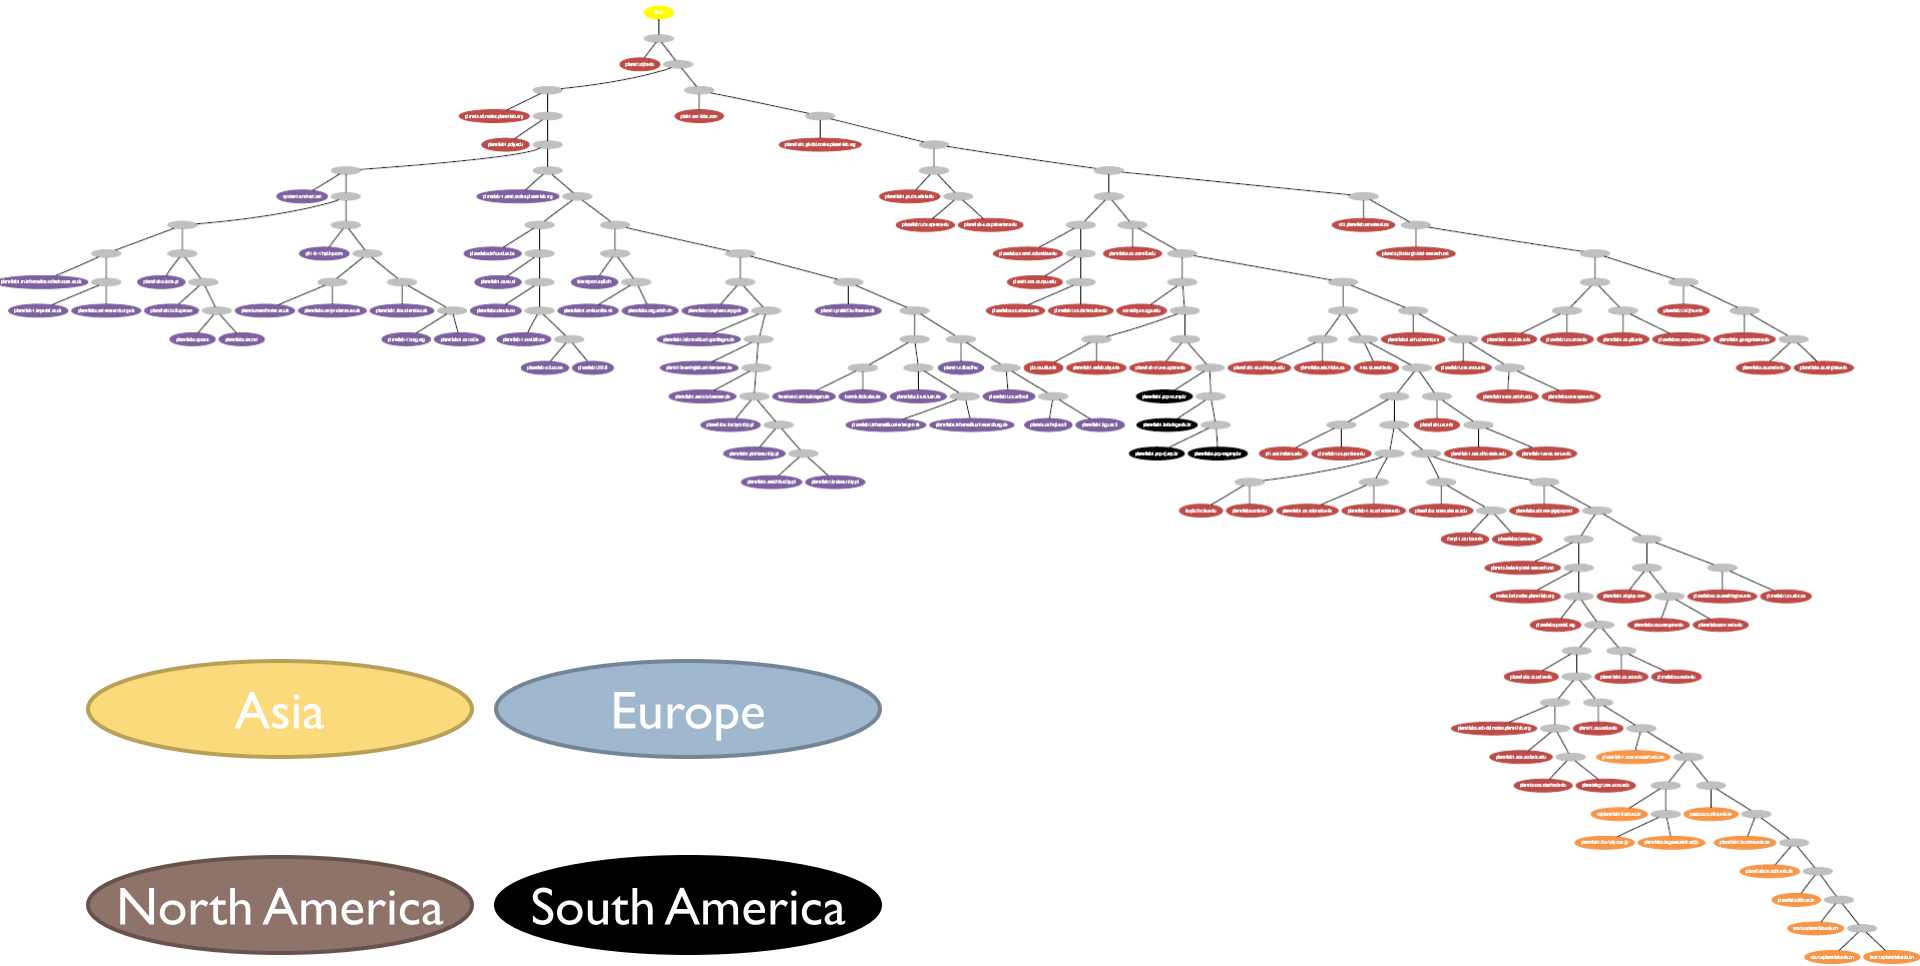
Closest Node



Best-Provisioned Node



Results: Sequoia Tree for PlanetLab



Asia

Europe

North America

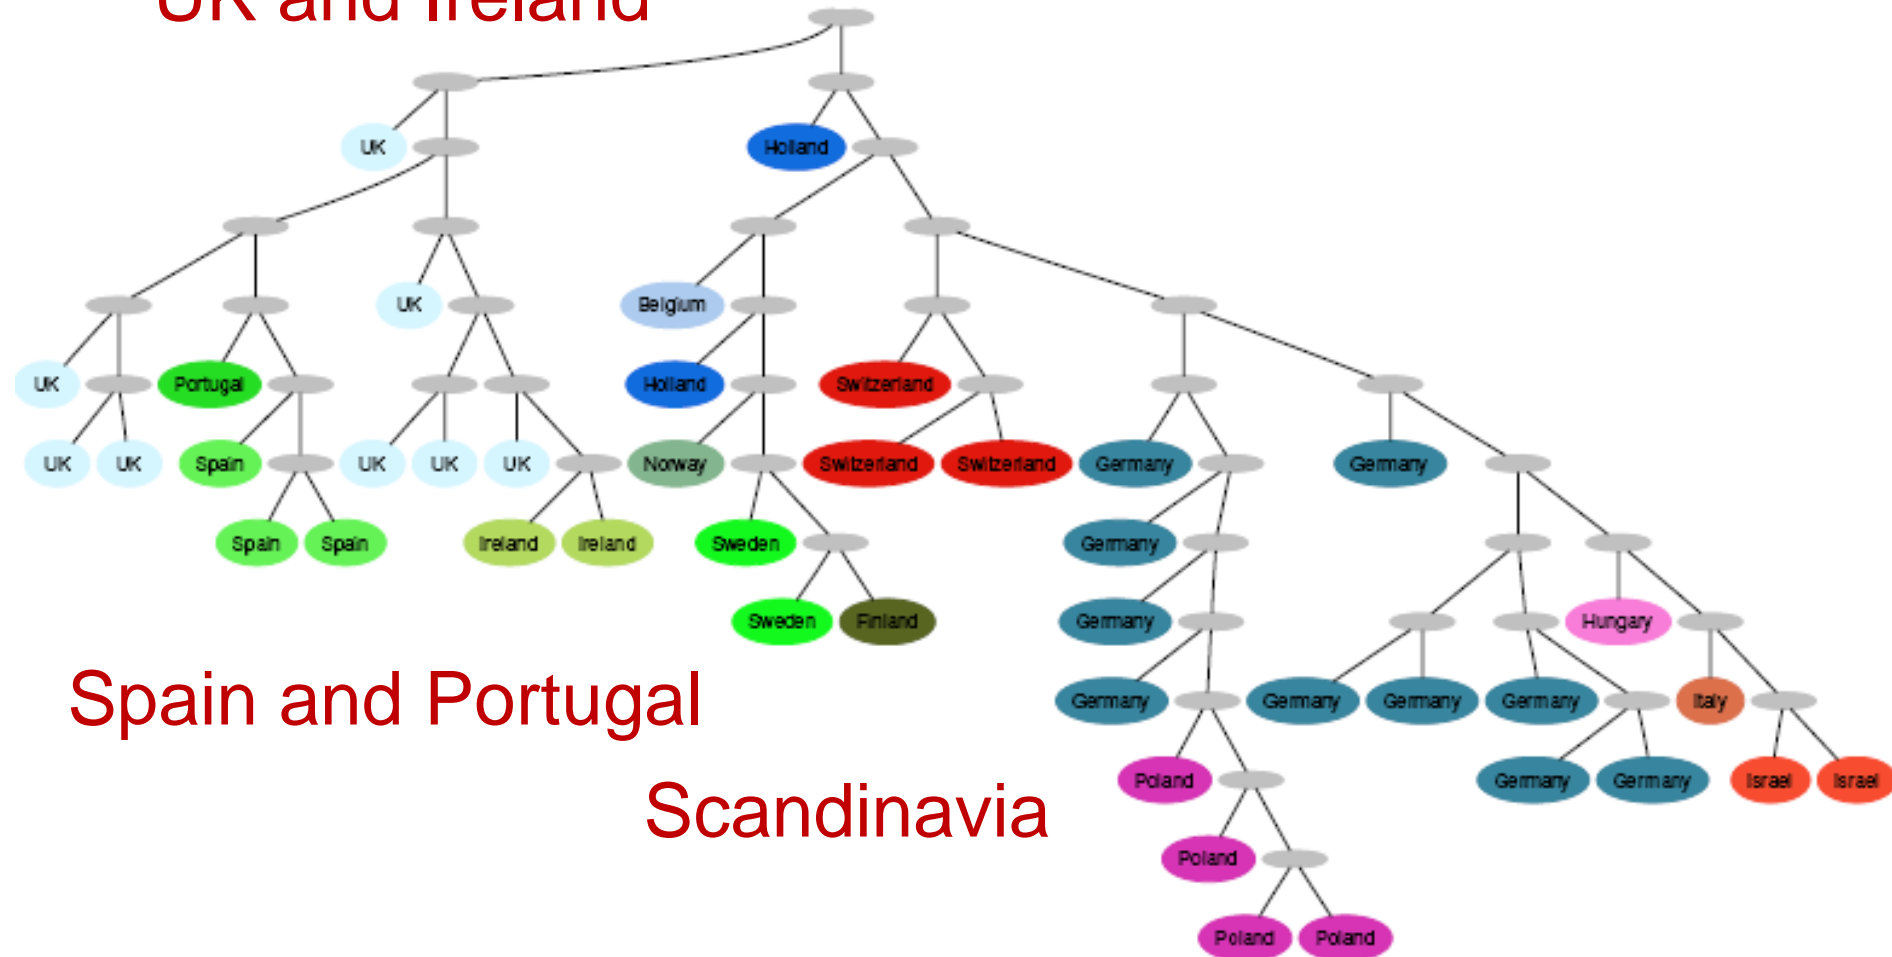
South America



Results: Clustering of European Nodes



UK and Ireland



Summary:

Virtual-Tree Internet Models




▶ Contributions

- ▶ Represents both Bandwidth and Latency
- ▶ Provides Convenient “Coordinates”
- ▶ Shows Topological Correlations

▶ Properties

- ▶ Light-weight end-to-end measurements
- ▶ Tolerates triangle inequality violations
- ▶ Adapts to changes in network measures



Email: rama@microsoft.com or dalia@microsoft.com

<http://research.microsoft.com/research/sv/sequoia>