Predicting Internet Network Distance with Coordinates-Based Approaches

T.S. Eugene Ng and Hui Zhang
Department of Computer Science
Carnegie Mellon University

Problem Statement

- Network distance
  - Round-trip propagation and transmission delay
  - Relatively stable, may be predictable
- Given two Internet hosts, can we accurately estimate the network distance between them without sending any RTT probes between them?

Why Predict Network Distance?

- Want to measure network performance to improve performance of applications
  - Napster, content addressable overlays, overlay multicast
- Huge number of paths to measure
- TCP bandwidth and RTT probes are time-consuming

- Predicted network distance enables fast and scalable first-order performance optimization
  - Eliminate poor choices
  - Refine when needed

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

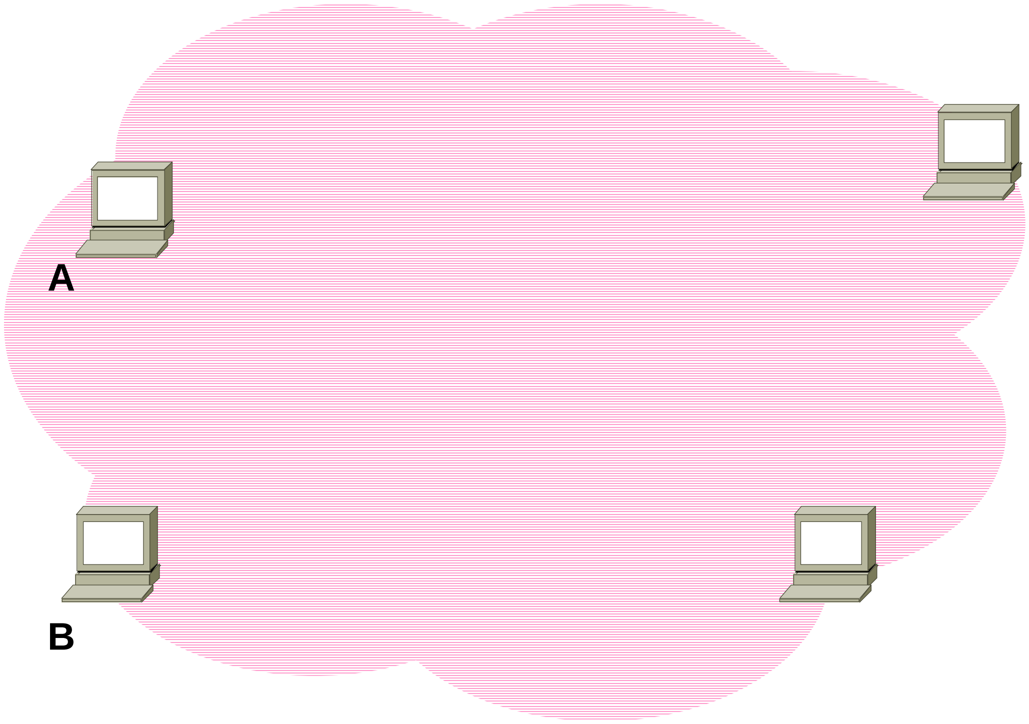

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

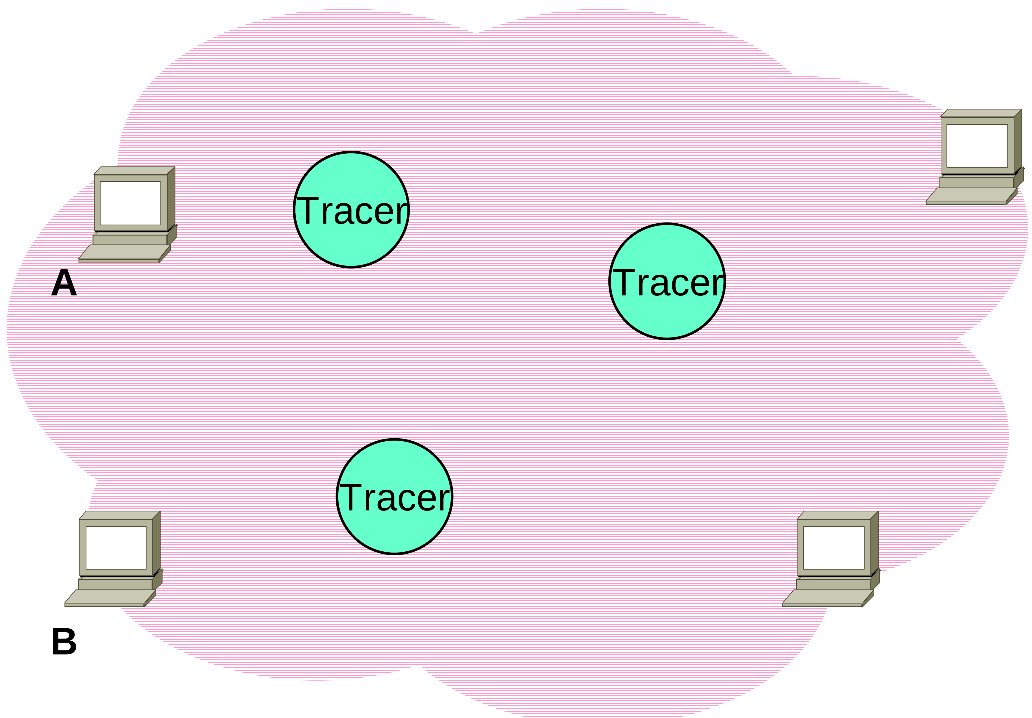

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

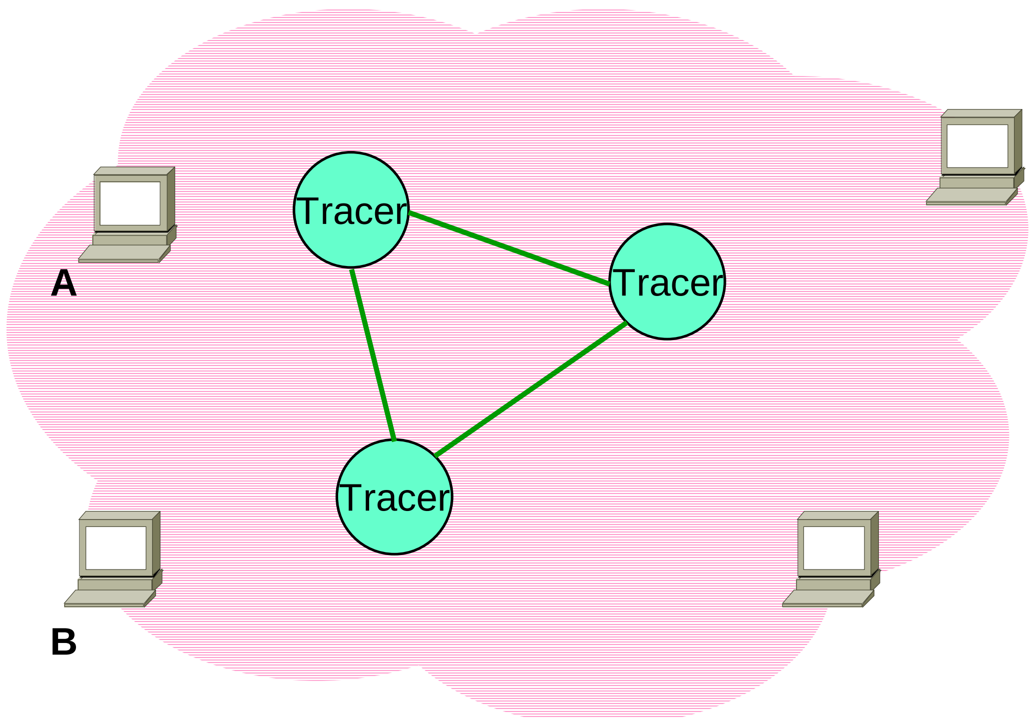

IDMaps [Francis et al. ‘99]

- Servers maintain simplified topological map

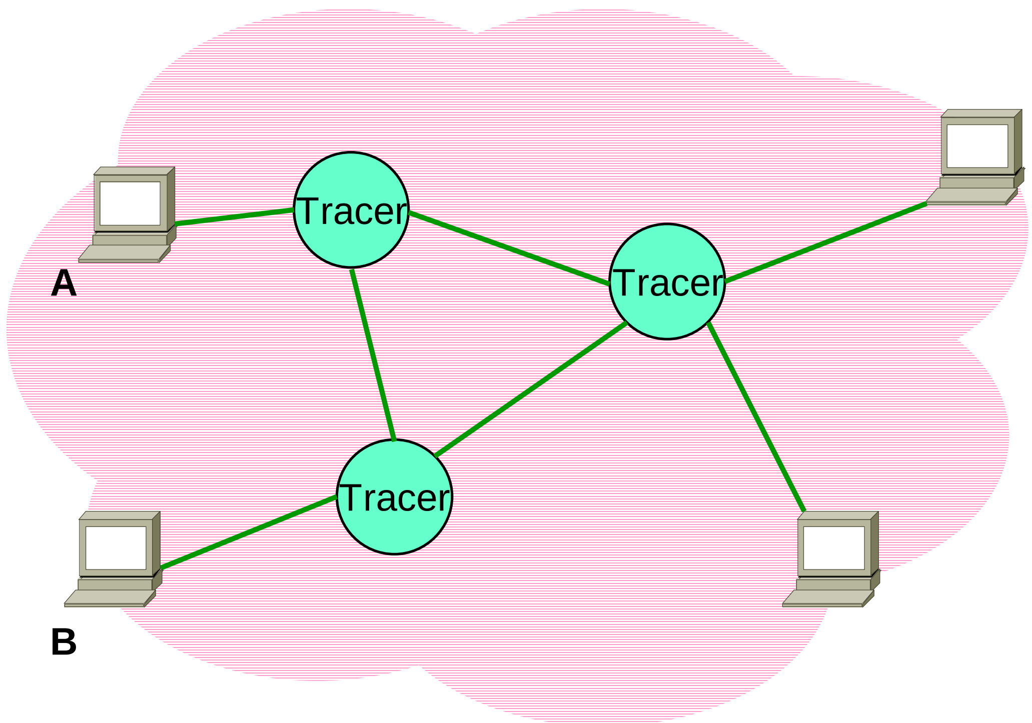

IDMaps [Francis et al. ‘99]

- Servers maintain simplified topological map

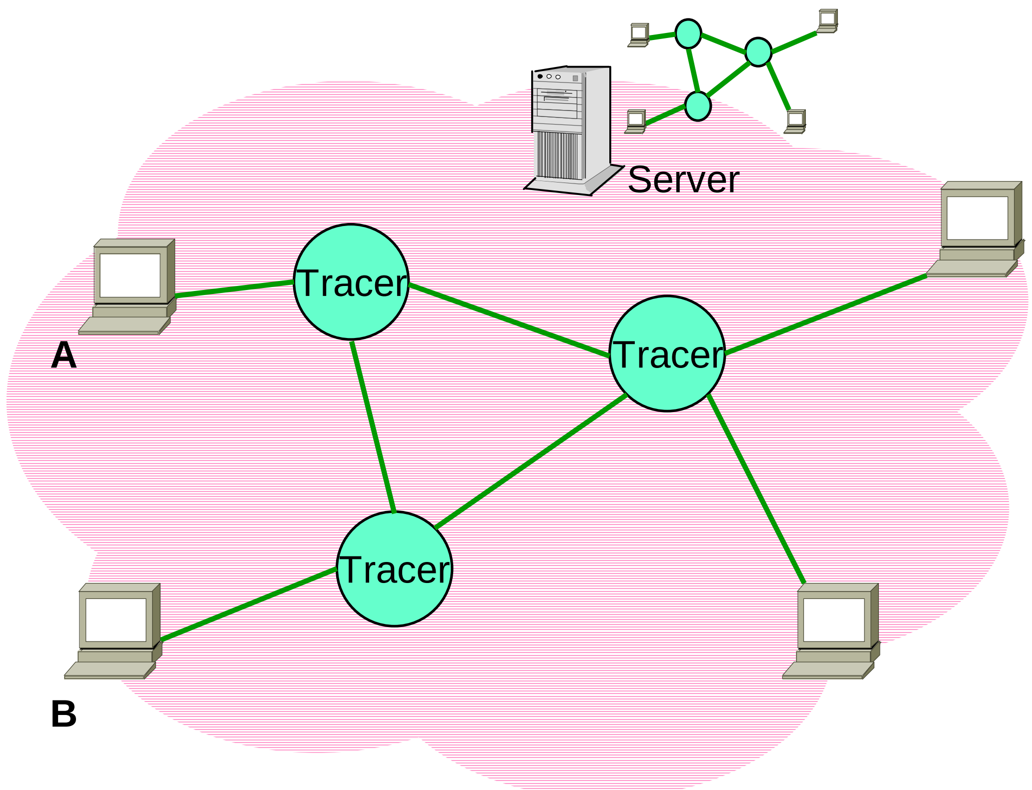

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

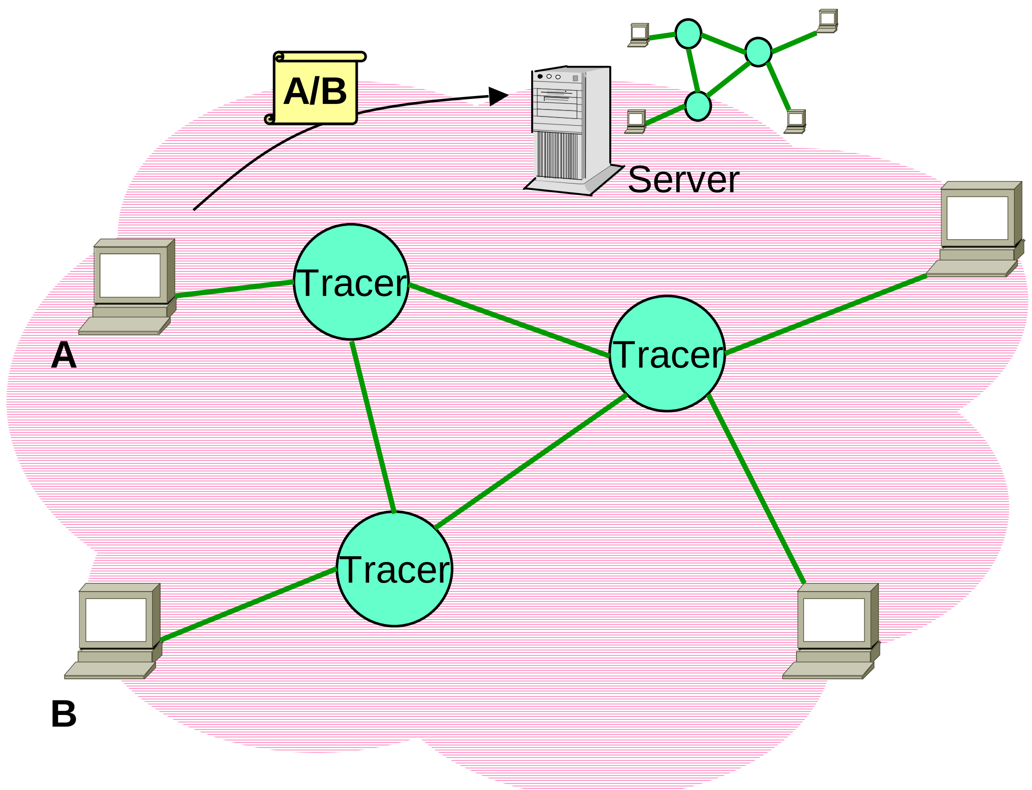

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

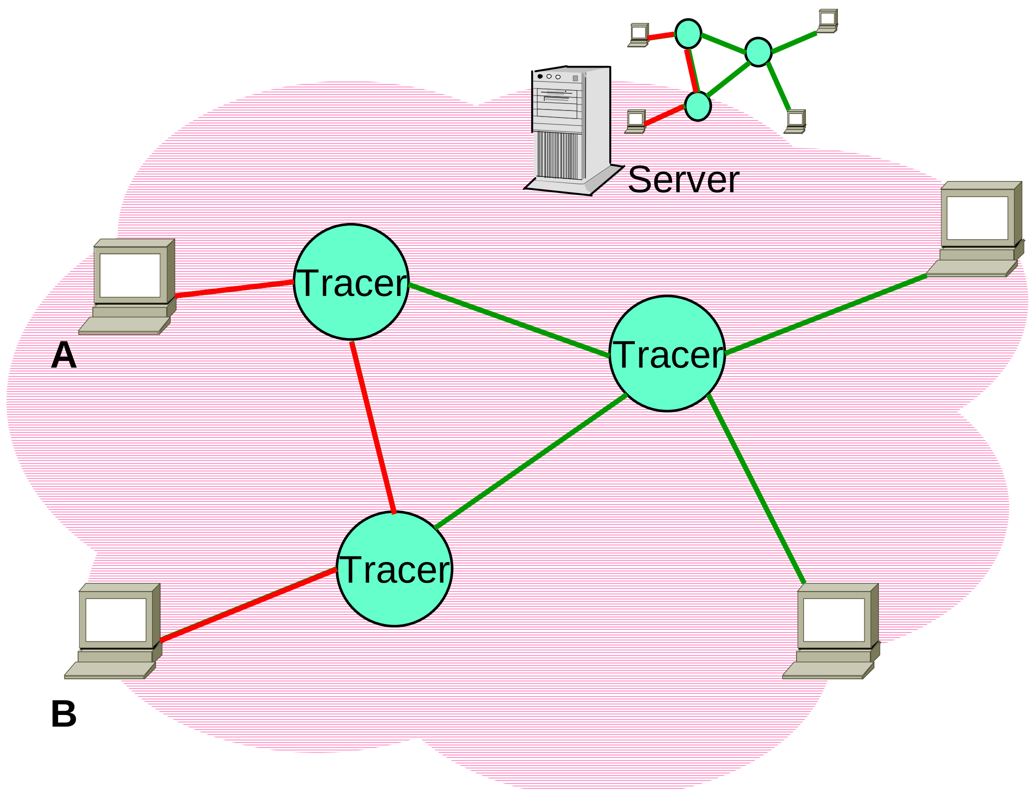

IDMaps [Francis et al. '99]

- Servers maintain simplified topological map

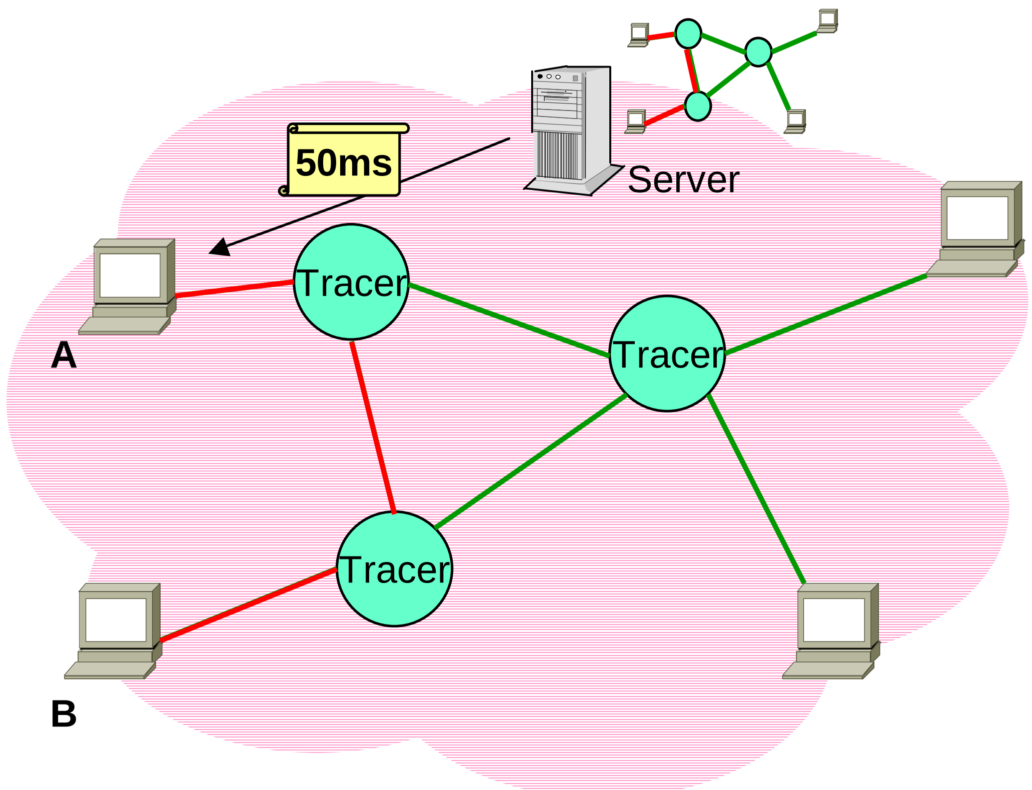

Disadvantage of Client-Server Architecture

- Unavoidable additional delay in communicating with servers
- Shared servers can become performance bottleneck

Can this Problem be Solved with a Peer-to-Peer Architecture?

- Flip the problem: Can end hosts maintain "coordinates" that describe their network locations?
- End hosts exchange coordinates to compute distance
  - High performance, high scalability

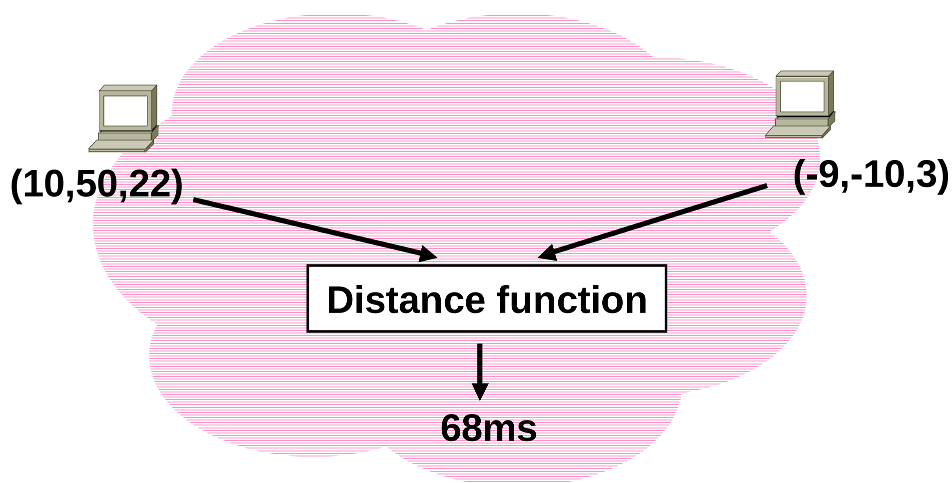

Approach 1: Global Network Positioning (GNP) Coordinates

- Model the Internet as a geometric space (e.g. 3-D Euclidean)
- Characterize the position of any end host with geometric coordinates
- Use geometric distances to predict network distances

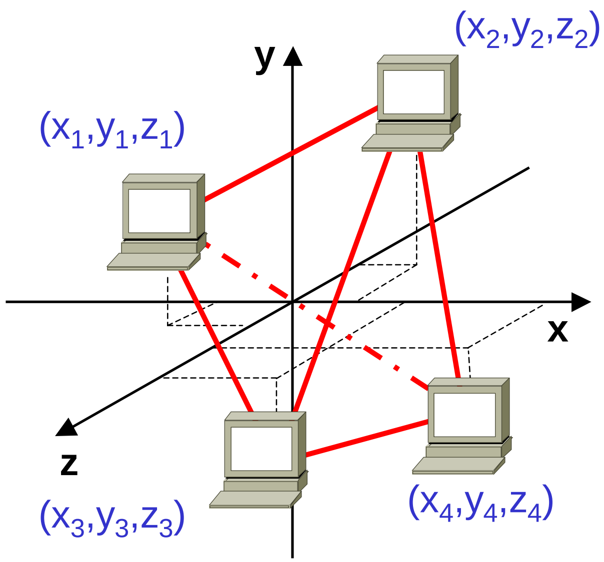

GNP Landmark Operations

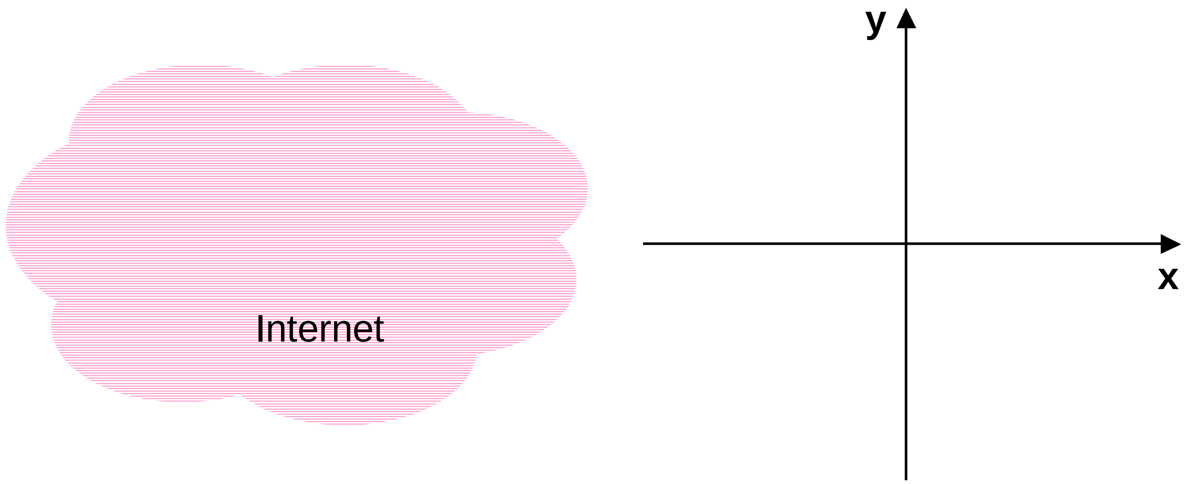

GNP Landmark Operations

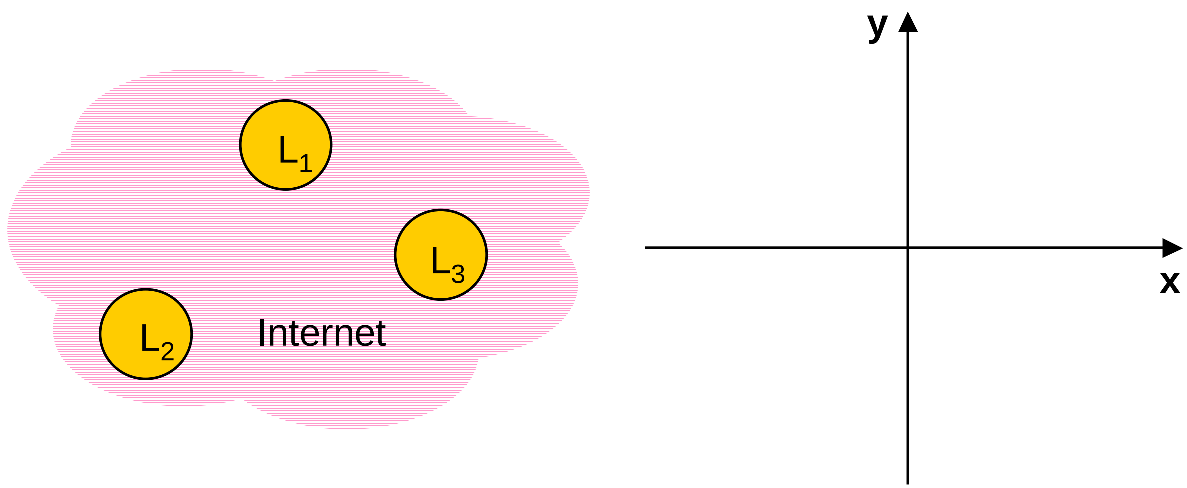

- Small number of distributed hosts called Landmarks measure inter-Landmark distances

GNP Landmark Operations

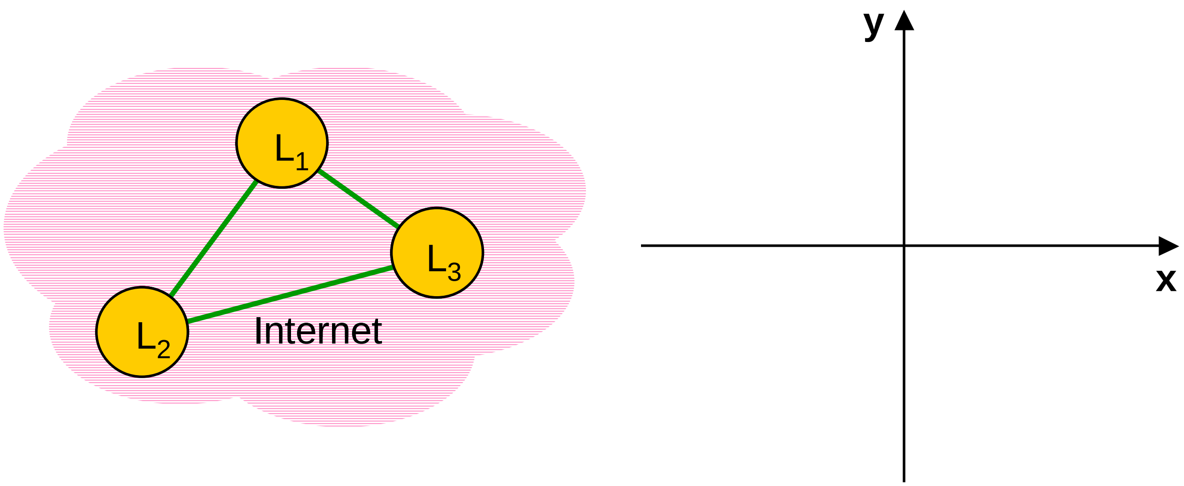

- Small number of distributed hosts called Landmarks measure inter-Landmark distances

GNP Landmark Operations

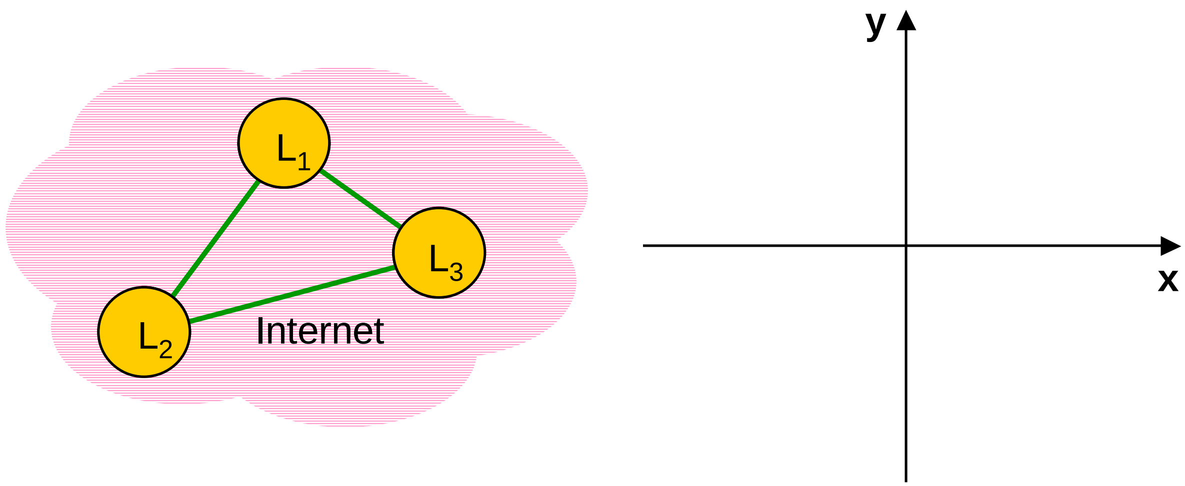

- Small number of distributed hosts called Landmarks measure inter-Landmark distances
- Compute Landmark coordinates by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

GNP Landmark Operations

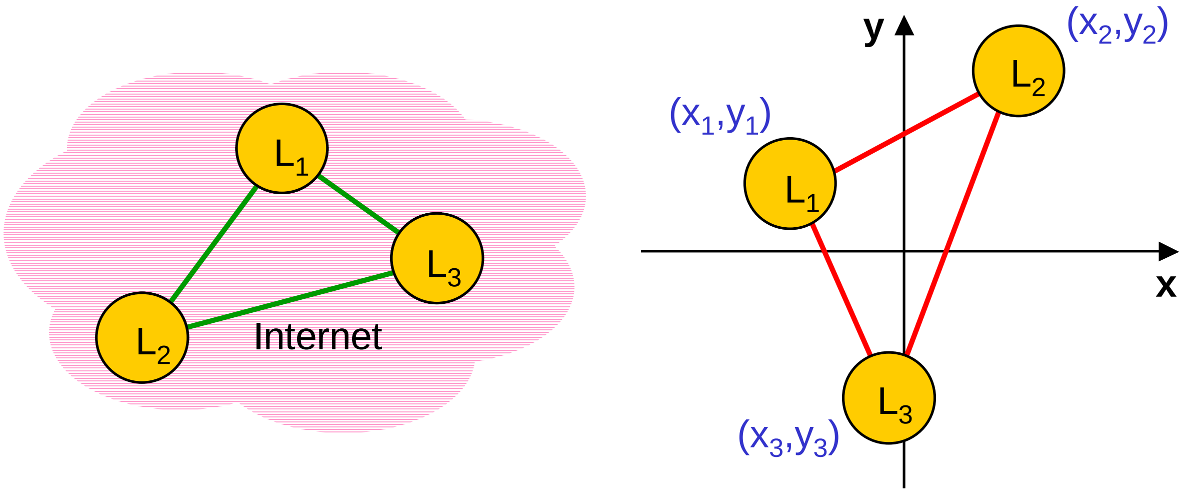

- Small number of distributed hosts called Landmarks measure inter-Landmark distances
- Compute Landmark coordinates by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

GNP Ordinary Host Operations

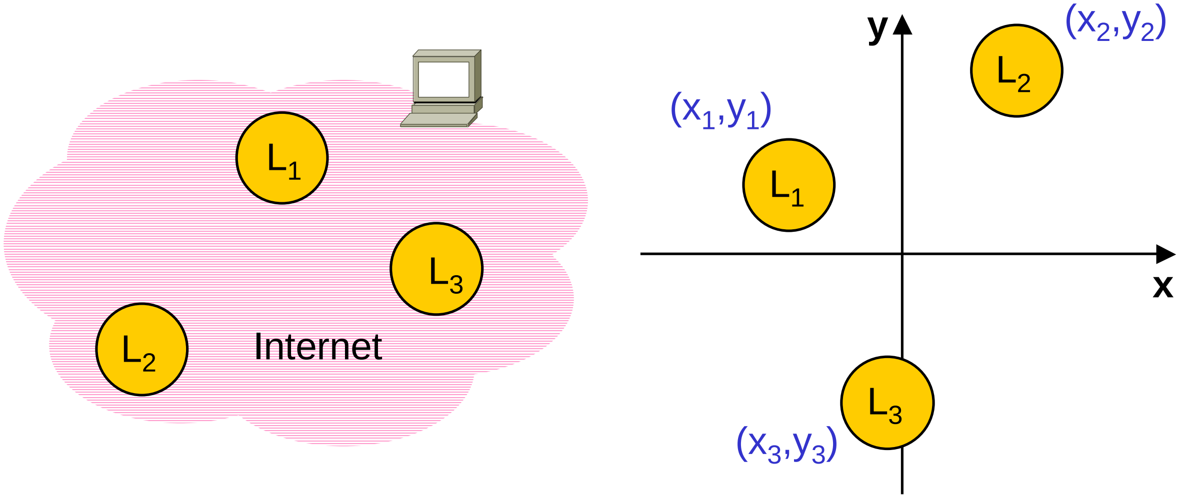

GNP Ordinary Host Operations

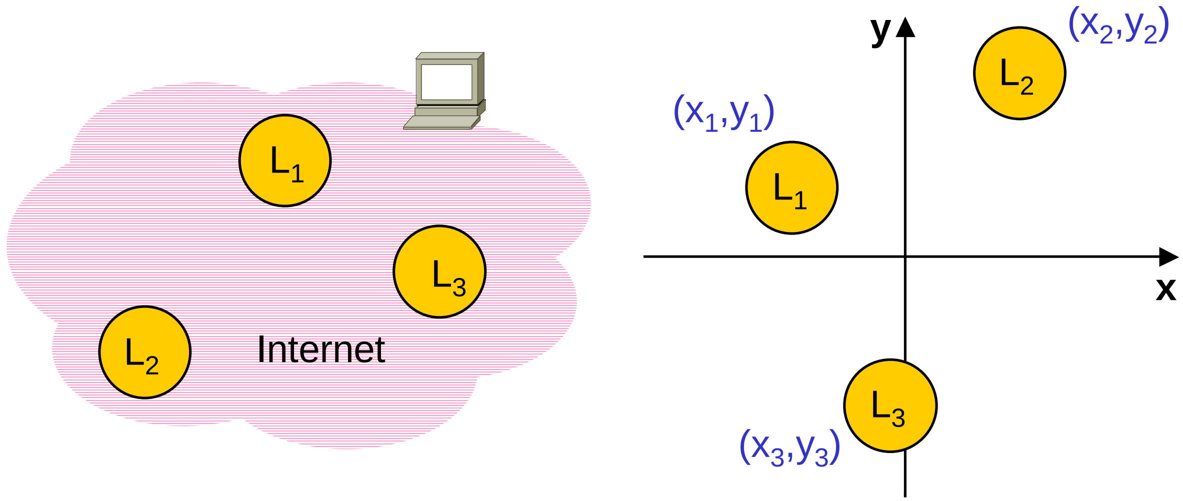

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings

GNP Ordinary Host Operations

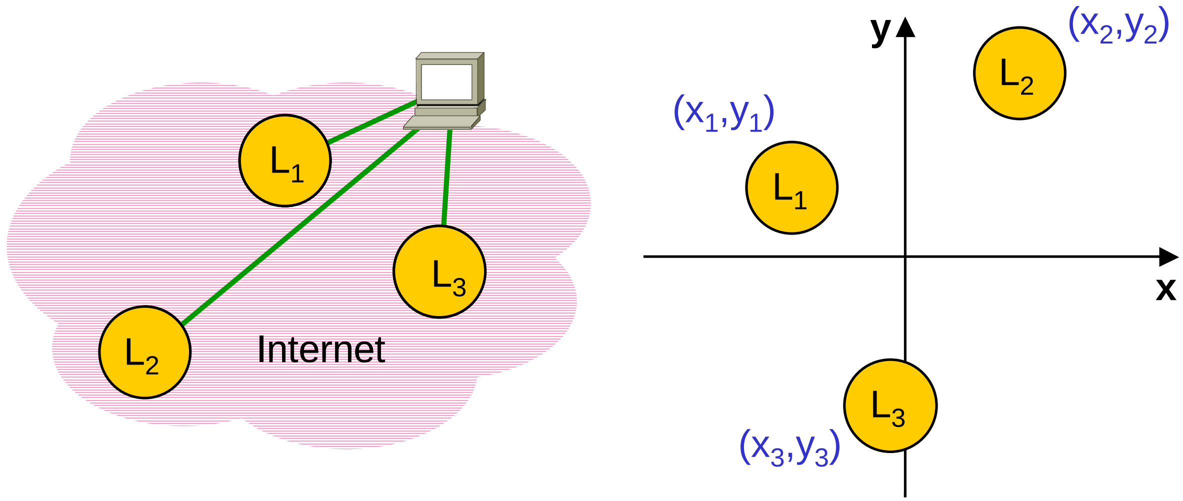

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings

GNP Ordinary Host Operations

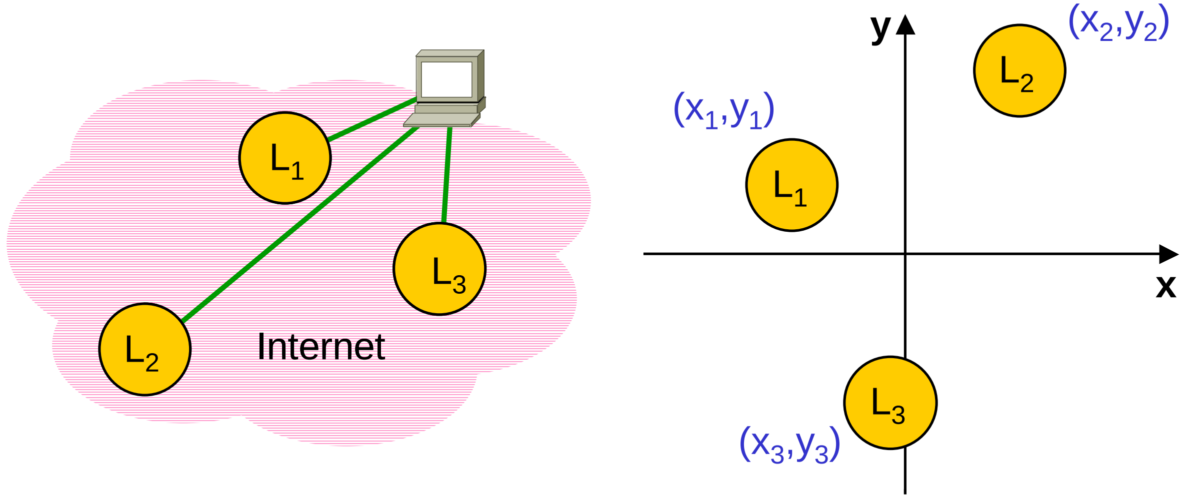

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings
- Ordinary host computes its own coordinates relative to the Landmarks by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

GNP Ordinary Host Operations

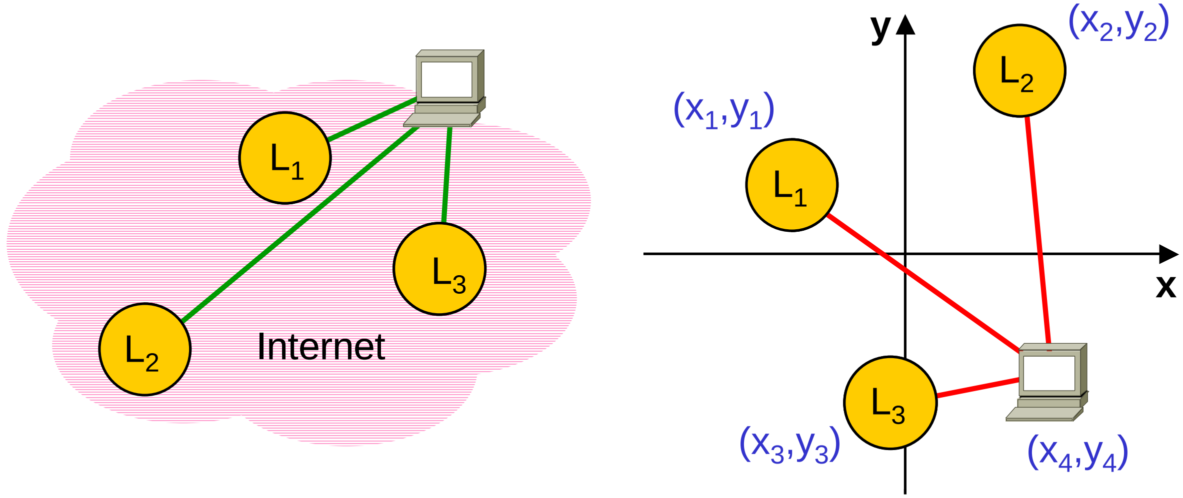

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings
- Ordinary host computes its own coordinates relative to the Landmarks by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

Important Questions

- What geometric model to use?
- How to measure error in minimizations?
- How to select Landmarks?
- How many Landmarks?
- What are the sources of prediction error?
- How to reduce overhead?
- Can we use geographical coordinates?

- Please see our paper
- This talk: focus on performance comparisons

Approach 2: Triangulated Heuristic Coordinates

- Proposed by Hotz in 1994 for *A** heuristic shortest network path search
- Provides upper and lower bounds for network distance
  - Assumes shortest path routing enforced
- This paper is the first study to apply and evaluate triangulated heuristic as a network distance prediction mechanism

Approach 2: Triangulated Heuristic Coordinates

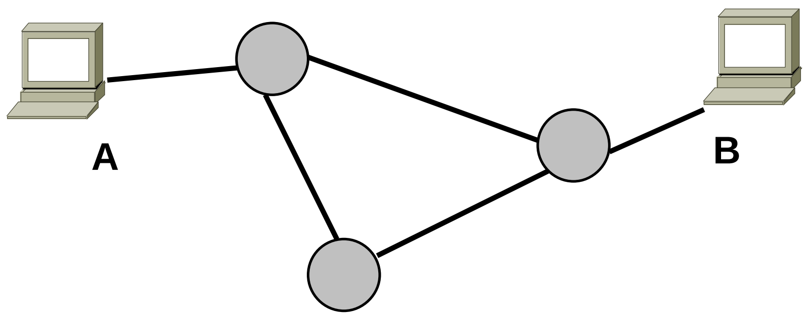

Hop count is used
in this example

Approach 2: Triangulated Heuristic Coordinates

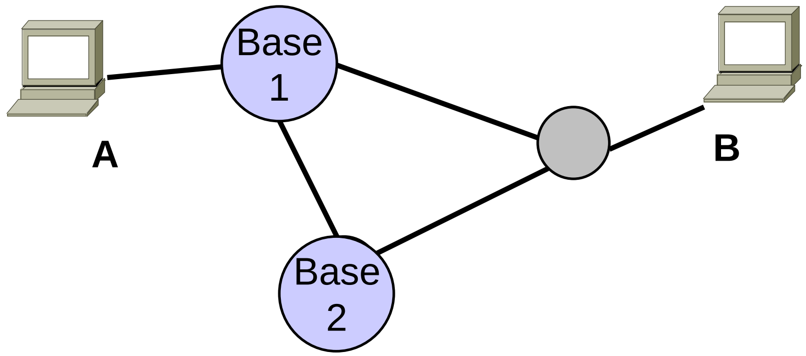

Hop count is used
in this example

Approach 2: Triangulated Heuristic Coordinates

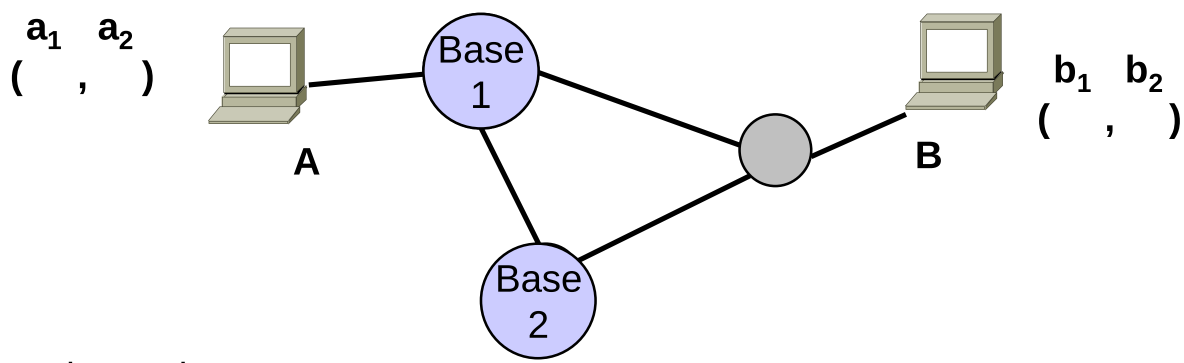

Hop count is used
in this example

Approach 2: Triangulated Heuristic Coordinates

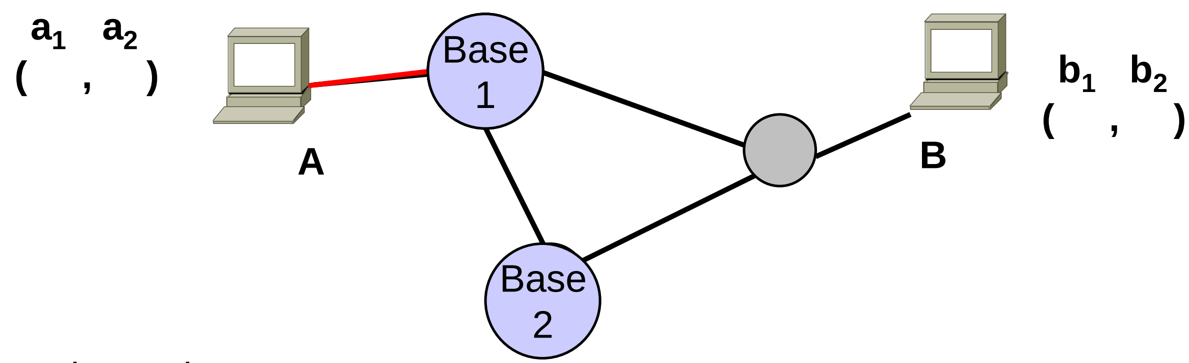

Hop count is used in this example

Approach 2: Triangulated Heuristic Coordinates

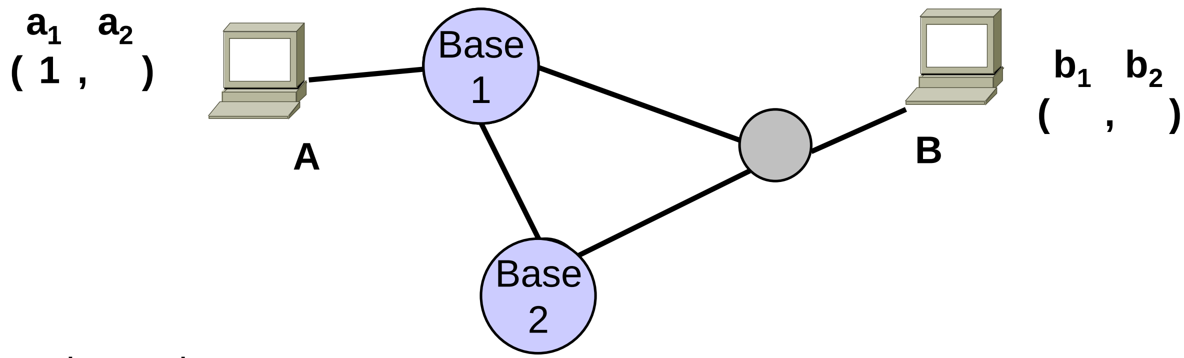

Hop count is used in this example

Approach 2: Triangulated Heuristic Coordinates

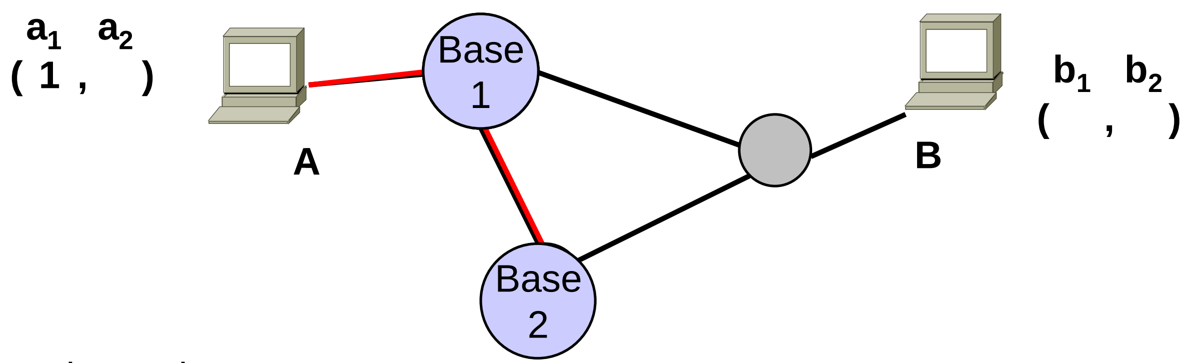

Hop count is used
in this example

Approach 2: Triangulated Heuristic Coordinates

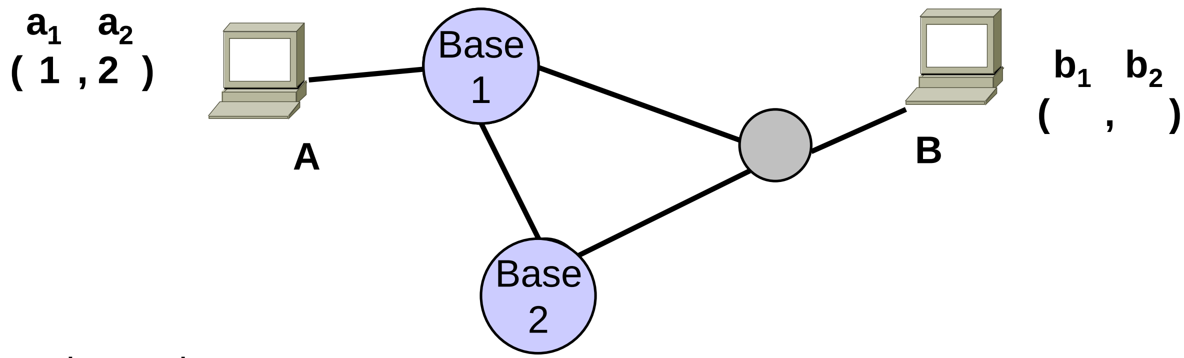

Hop count is used in this example

Approach 2: Triangulated Heuristic Coordinates

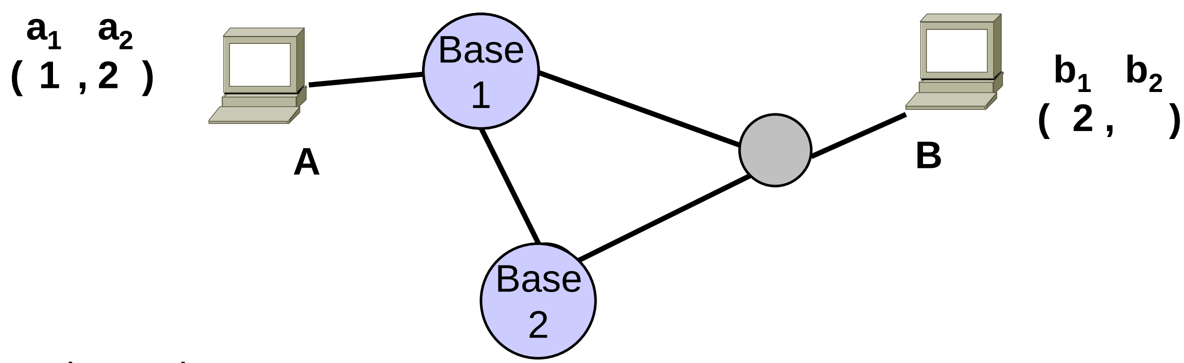

Hop count is used in this example

Approach 2: Triangulated Heuristic Coordinates

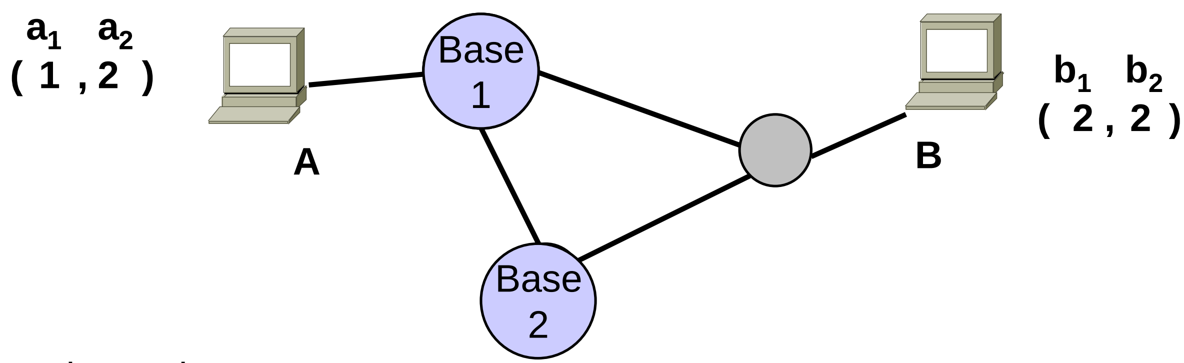

Hop count is used
in this example

Approach 2: Triangulated Heuristic Coordinates

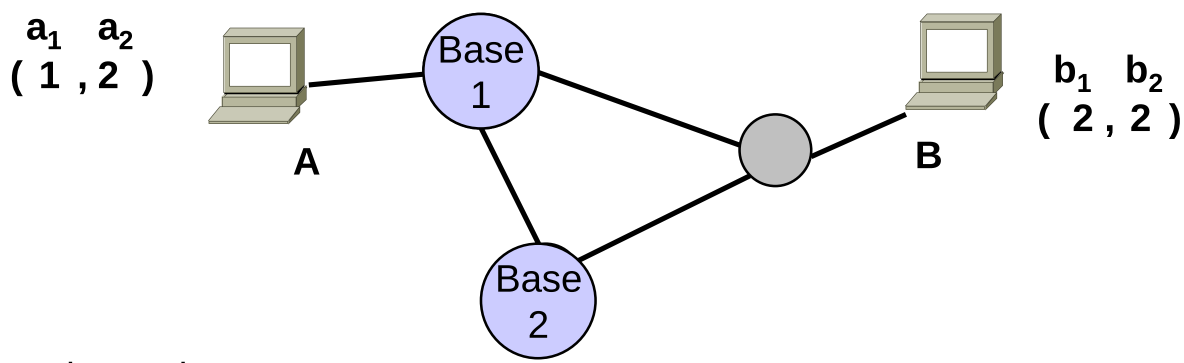

Hop count is used in this example

Upper bound (U) = min (a_i + b_i)

Approach 2: Triangulated Heuristic Coordinates

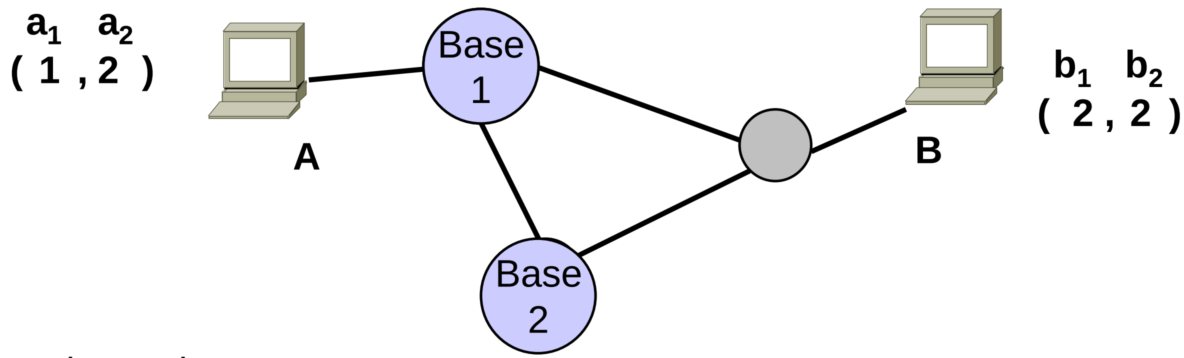

Hop count is used in this example

Upper bound (U) = min ($a_i + b_i$) = 3

Approach 2: Triangulated Heuristic Coordinates

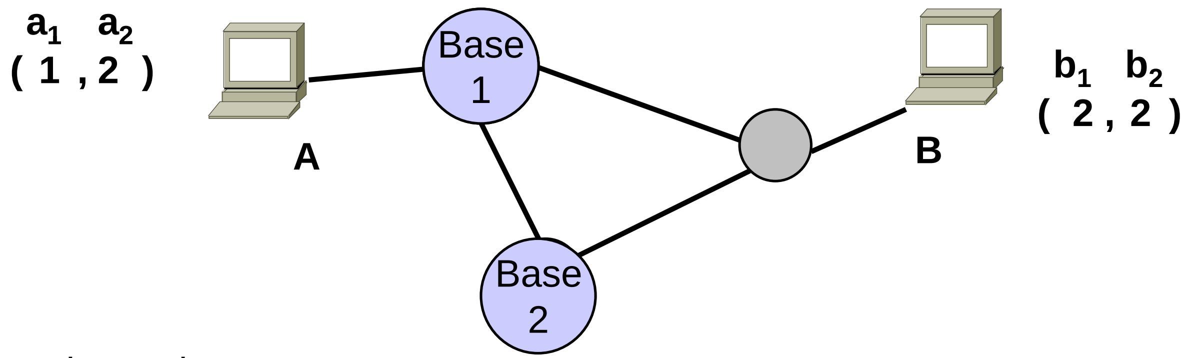

Hop count is used in this example

Upper bound (U) = min $(a_i + b_i)$ = 3 ⇐ Correct!

Evaluation Methodology

- 19 Probes
  – 12 in North America, 5 in East Asia, 2 in Europe
- 869 IP addresses called Targets we do not control
  – Span 44 countries
- Probes measure
  – Inter-Probe distances
  – Probe-to-Target distances
- See paper for more results

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

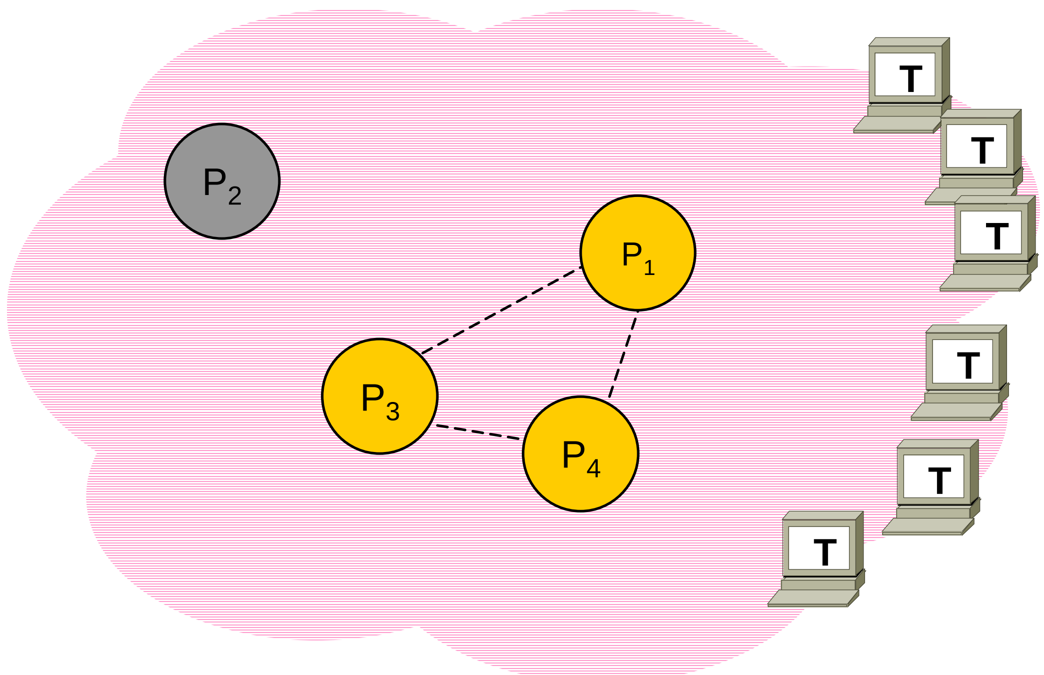

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

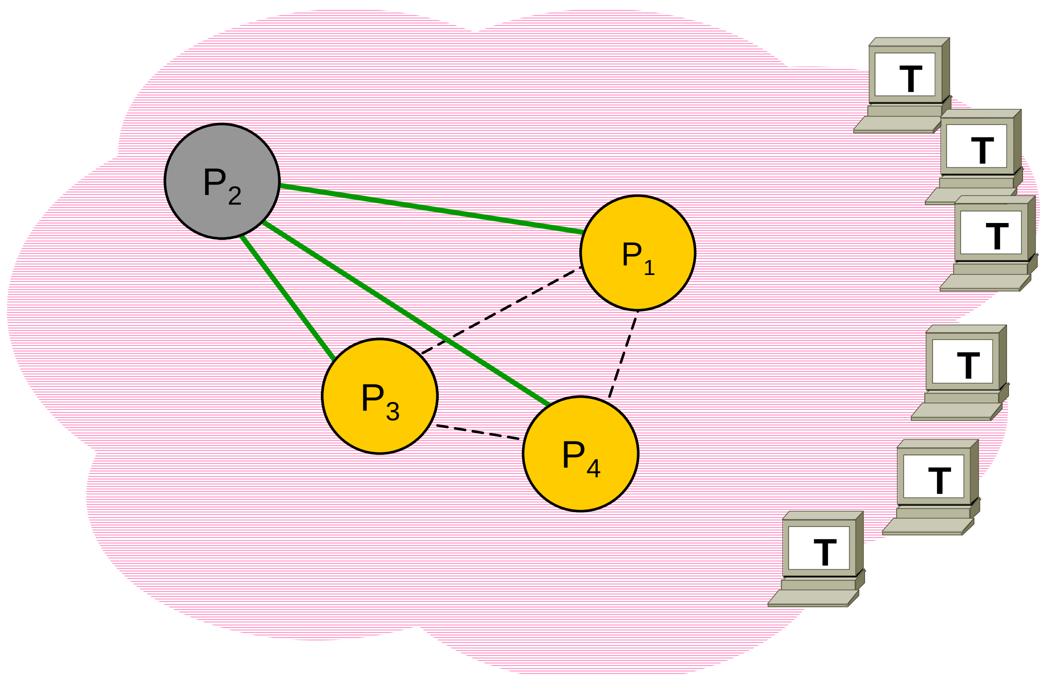

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

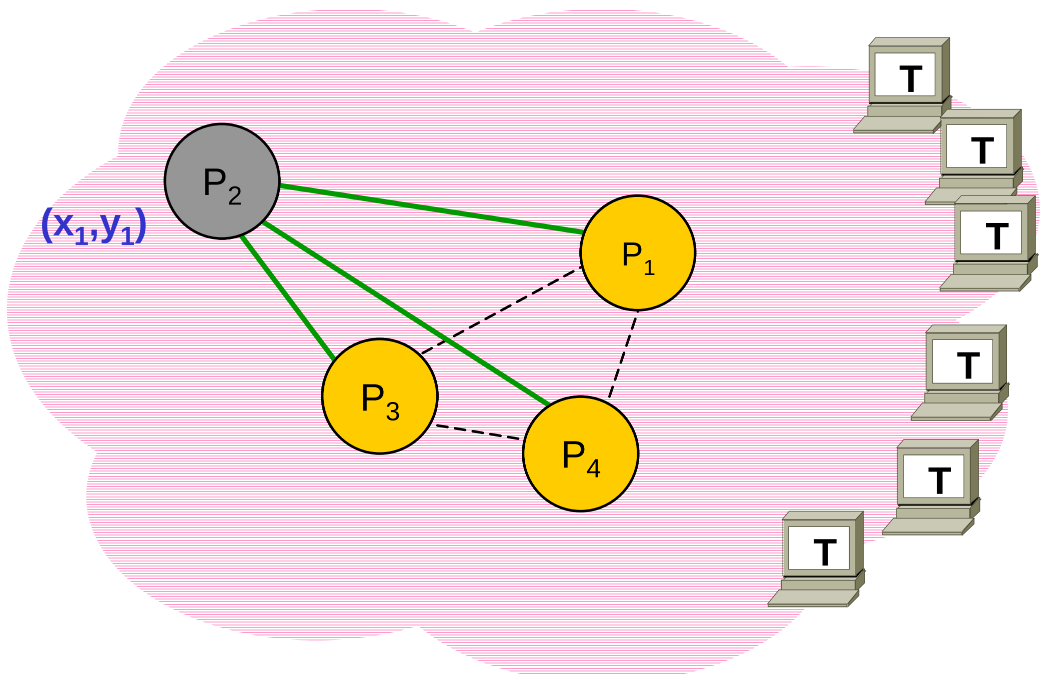

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

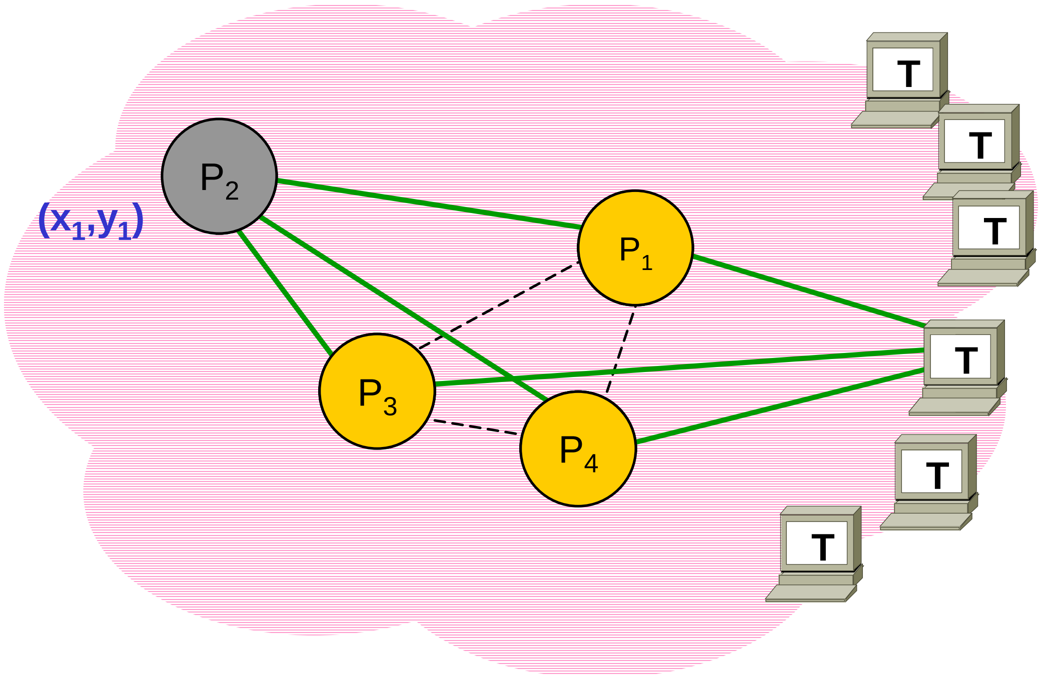

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

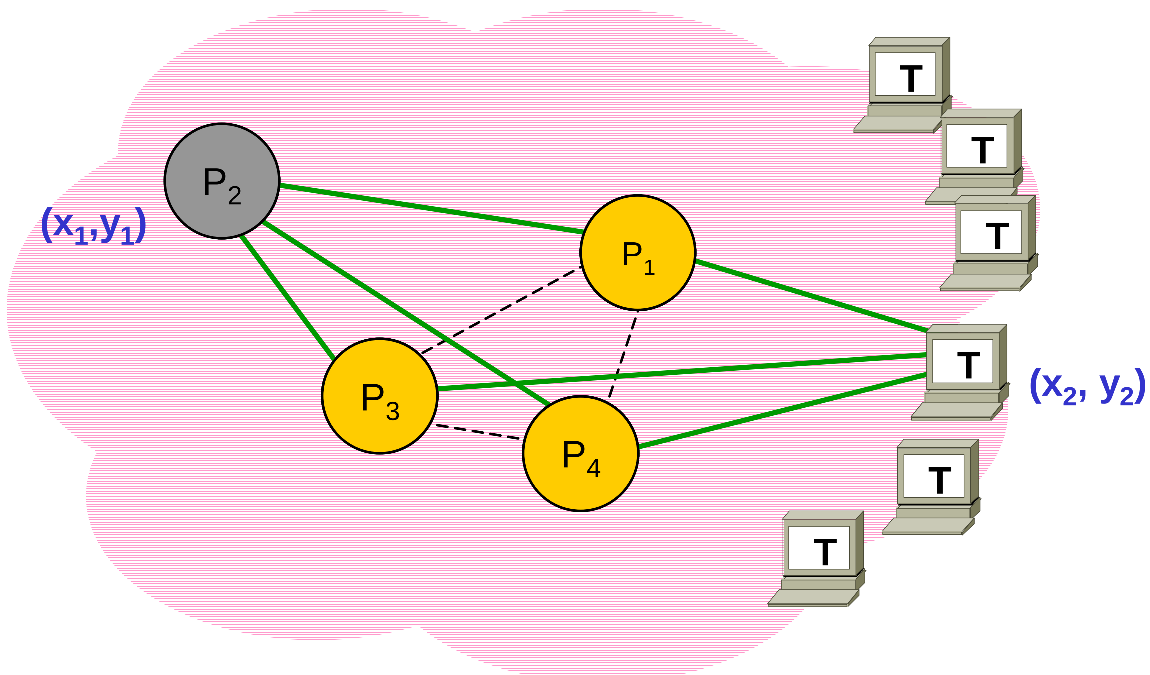

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Tracers/Base nodes/Landmarks
- Use the rest for evaluation

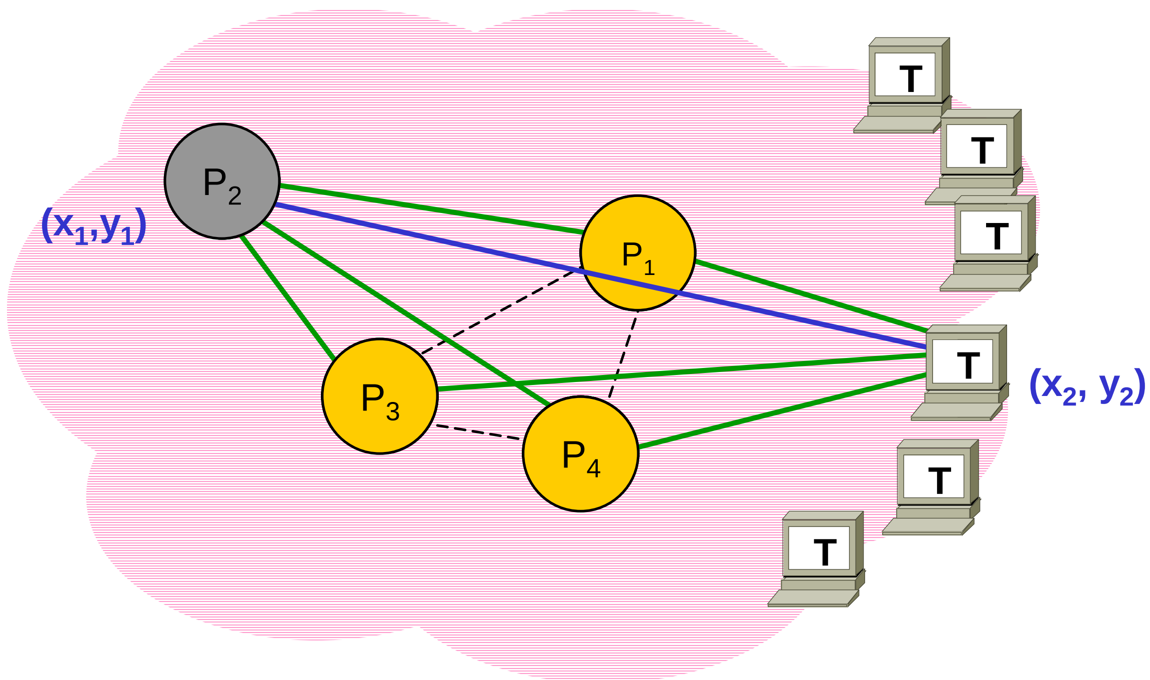

Performance Metrics

- Directional relative error
  - Symmetrically measure over and under predictions

$$\frac{predicted - measured}{\min(measured, predicted)}$$

- Relative error = abs(Directional relative error)

Relative Error Comparison

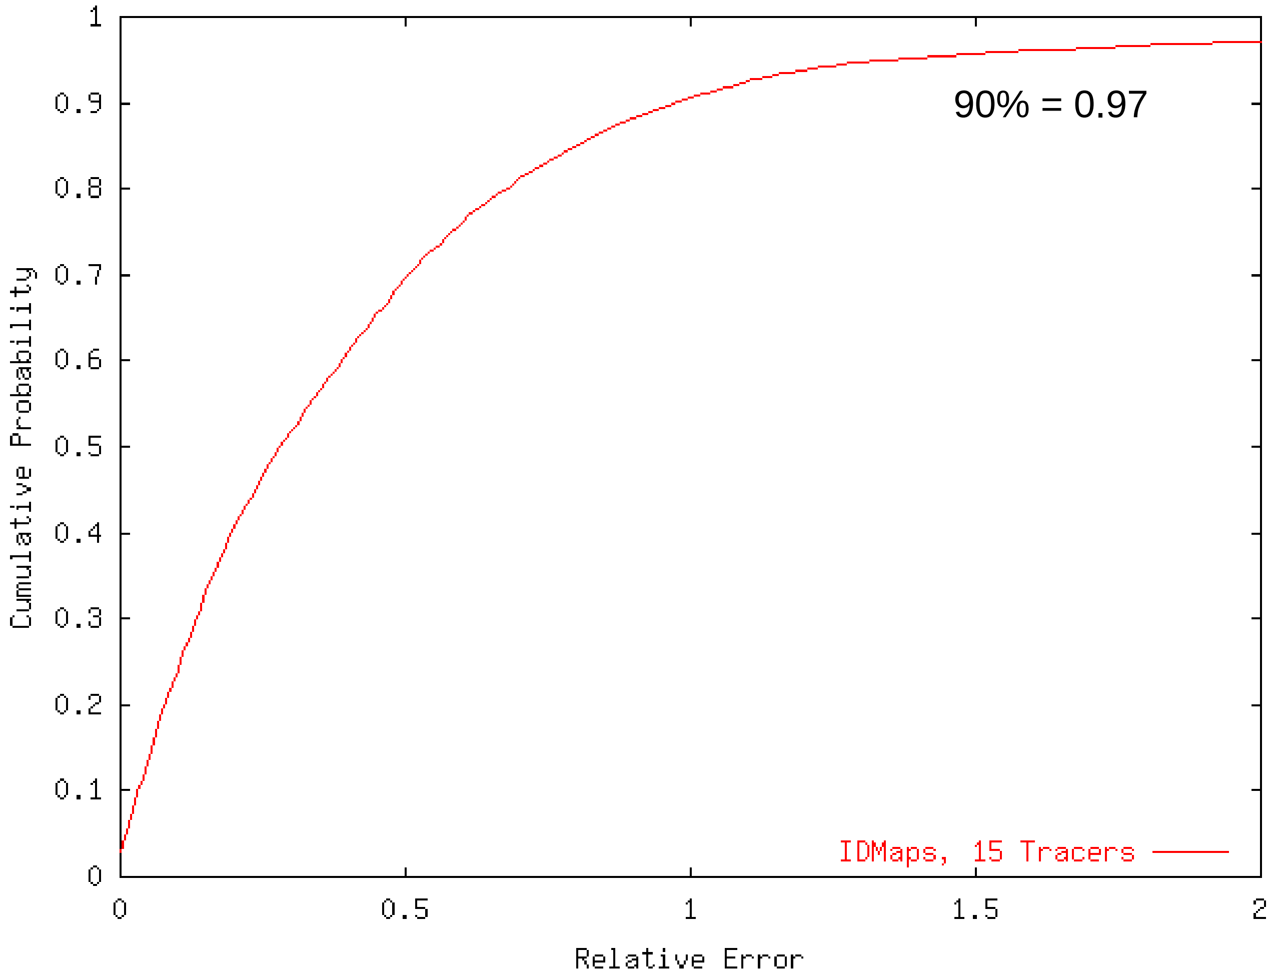

Relative Error Comparison

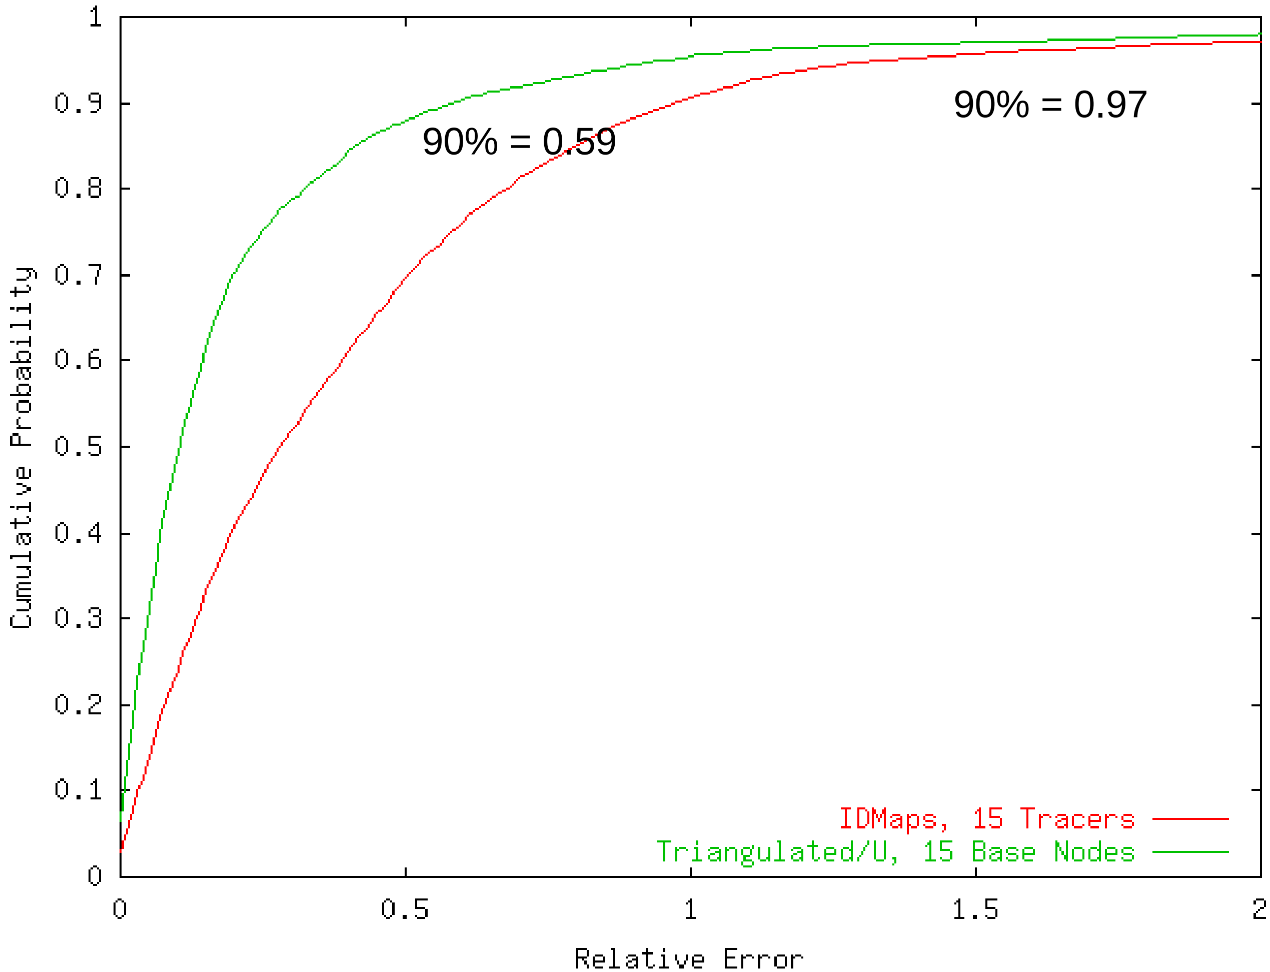

Relative Error Comparison

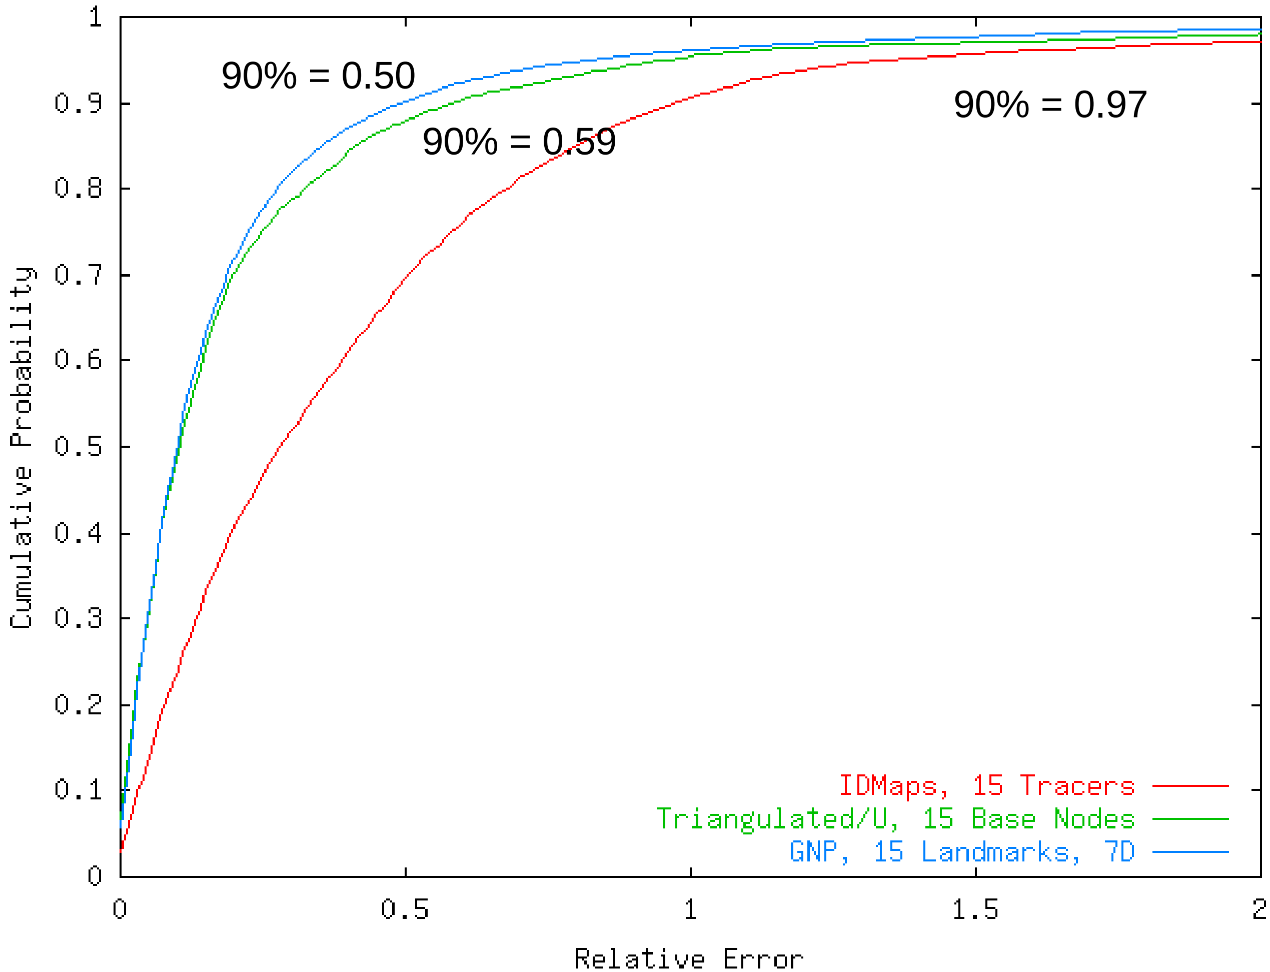

Directional Relative Error Analysis

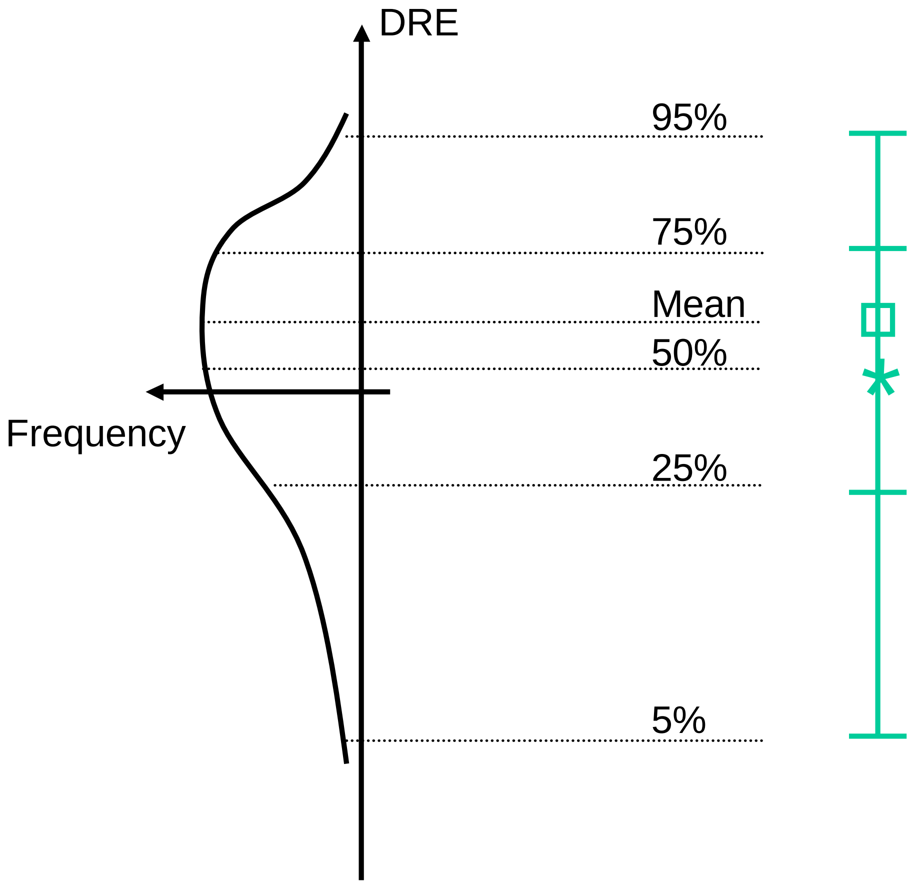

Directional Relative Error Comparison

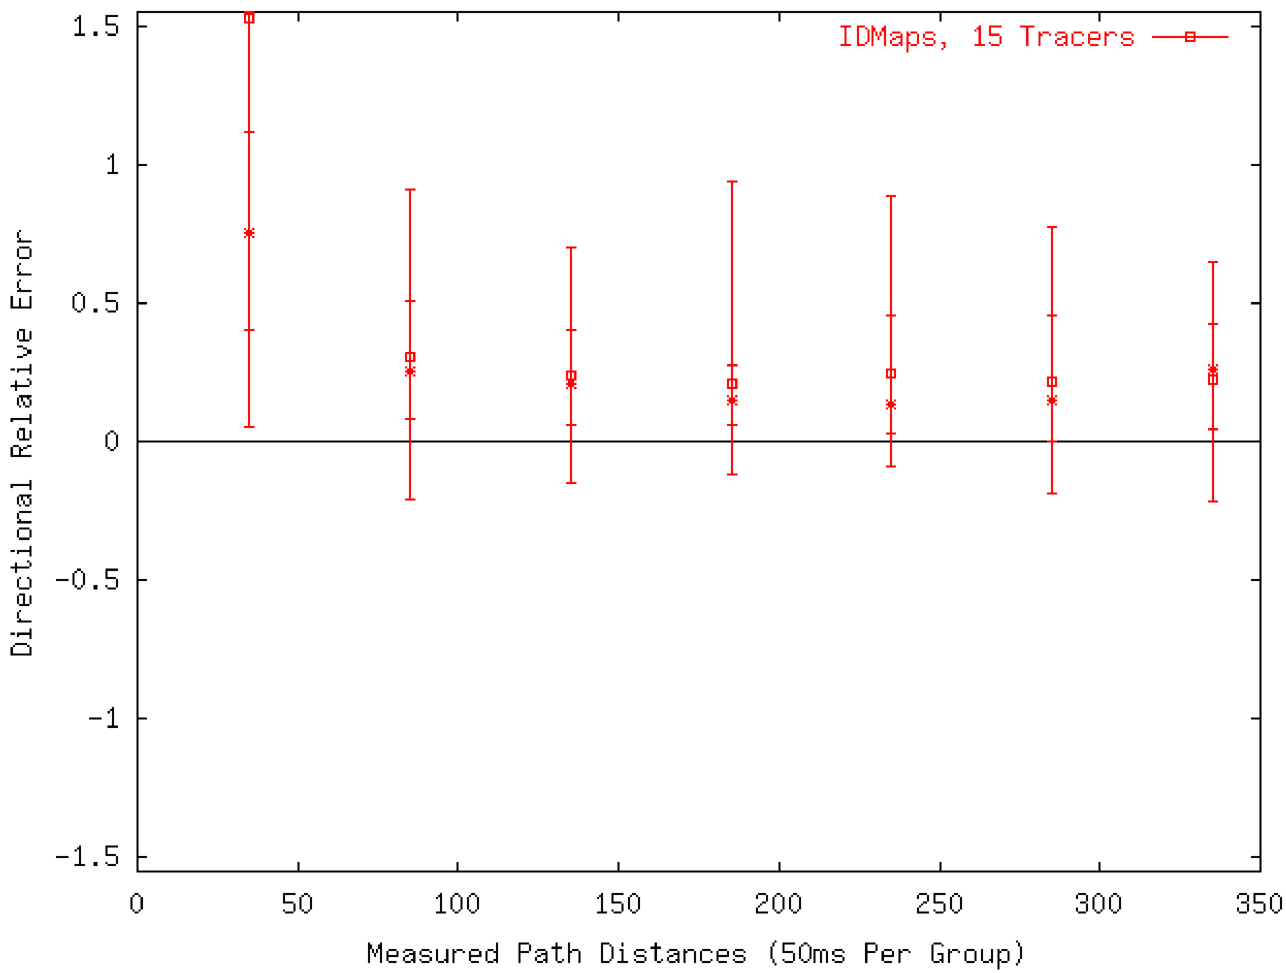

Directional Relative Error Comparison

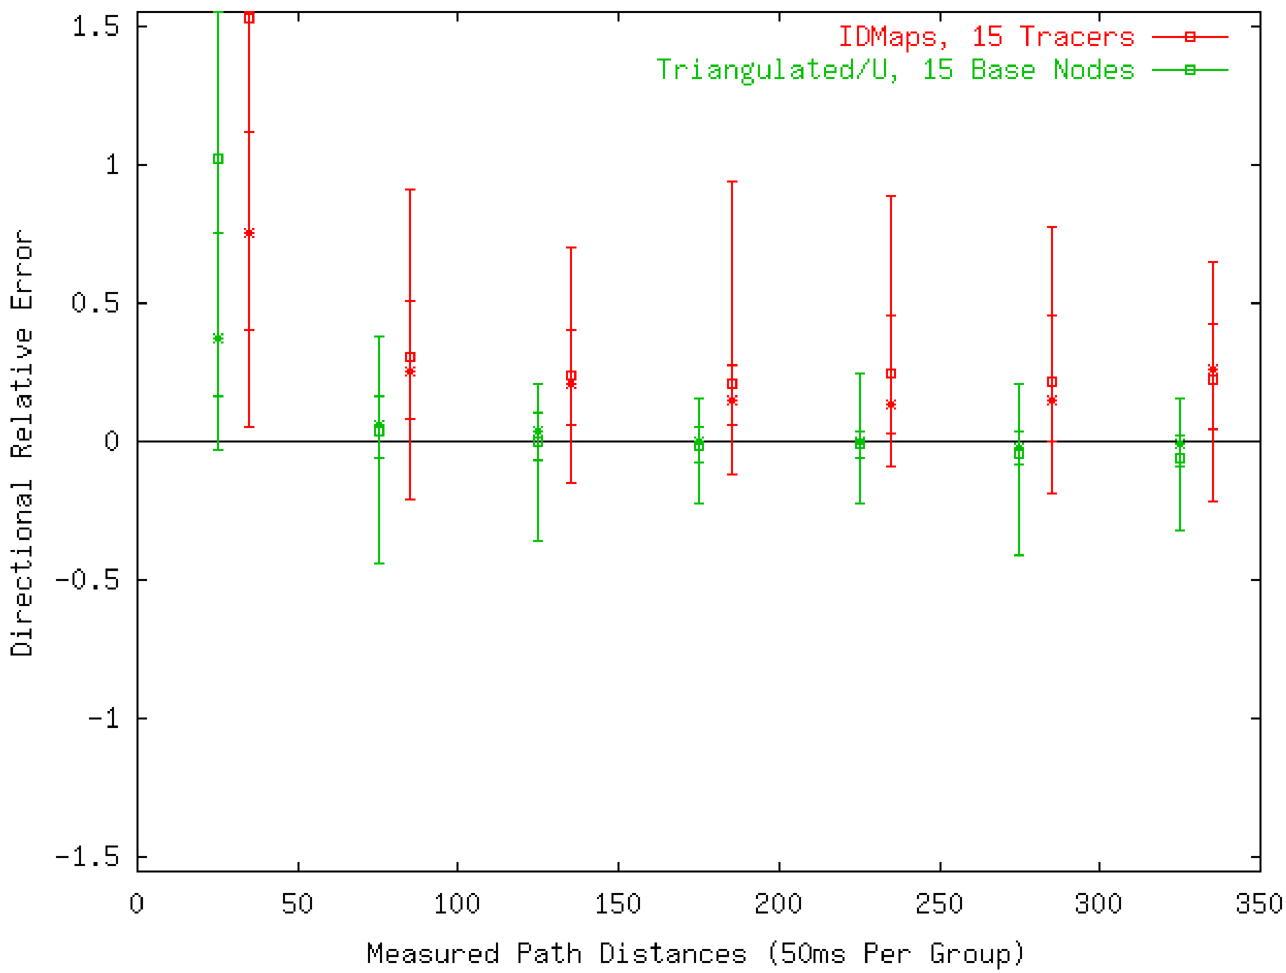

Directional Relative Error Comparison

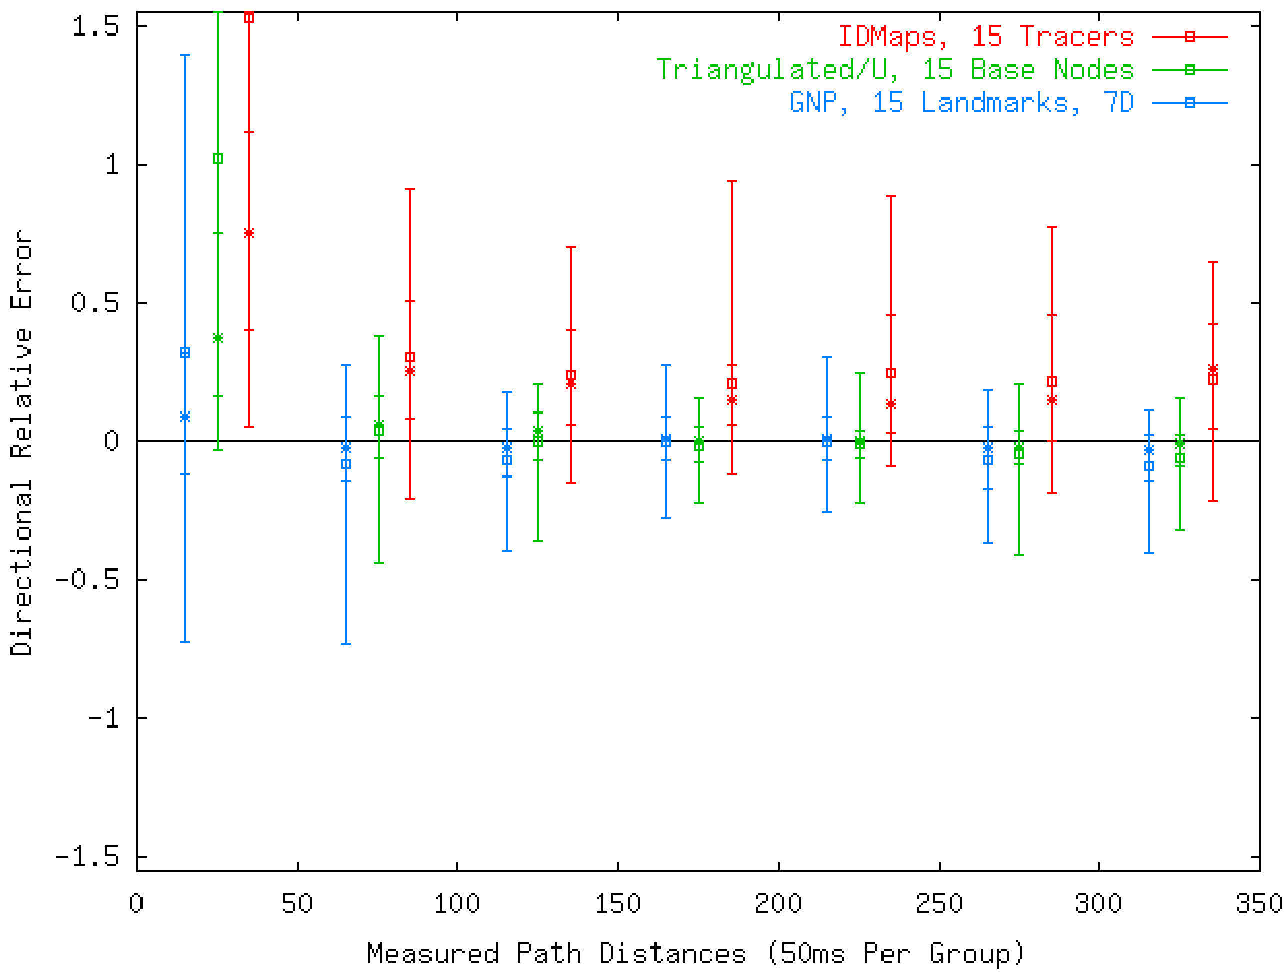

Why the Difference in Over-Predictions?

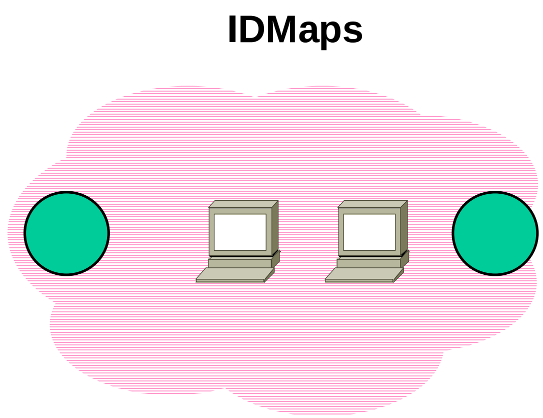

Straight-line distance
used in this example

Why the Difference in Over-Predictions?

IDMaps

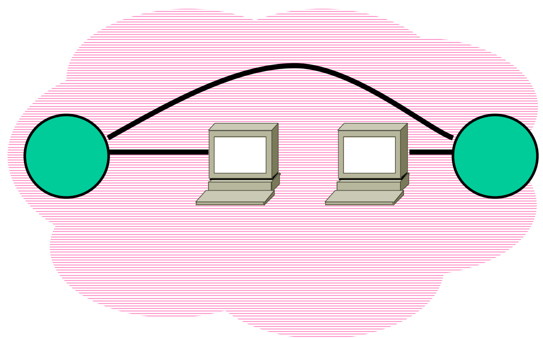

Straight-line distance
used in this example

Why the Difference in Over-Predictions?

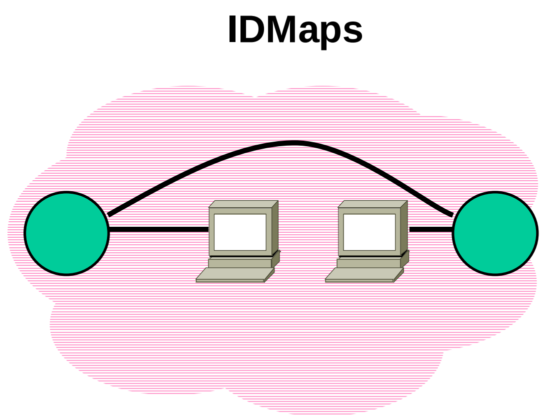

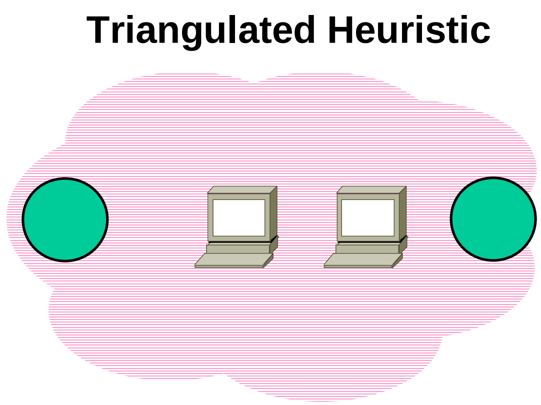

Straight-line distance
used in this example

Why the Difference in Over-Predictions?

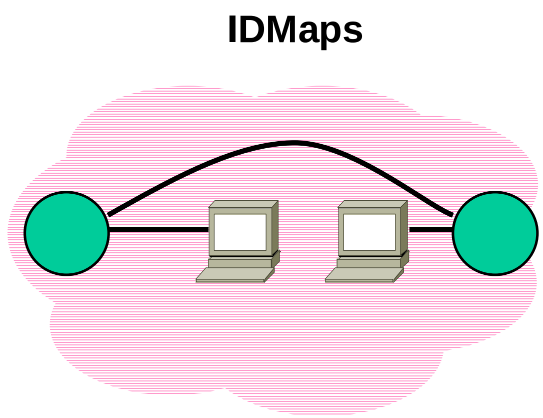

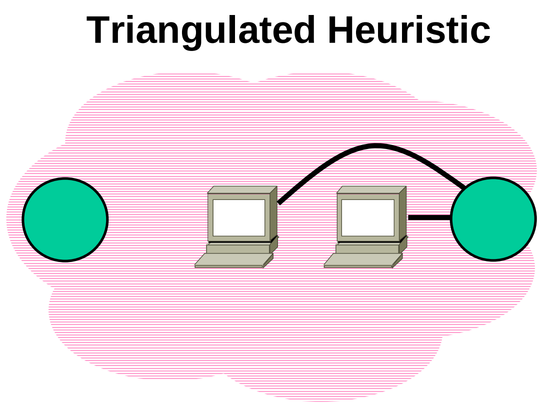

Straight-line distance
used in this example

Why the Difference in Over-Predictions?

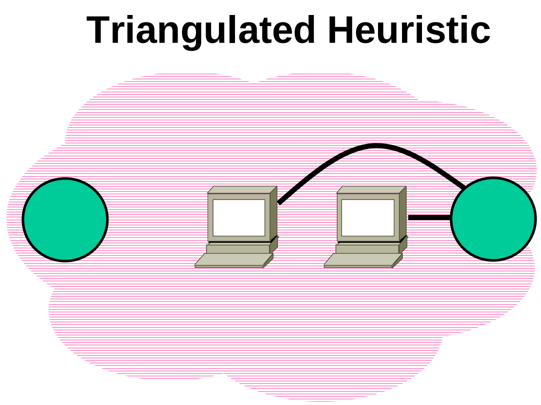

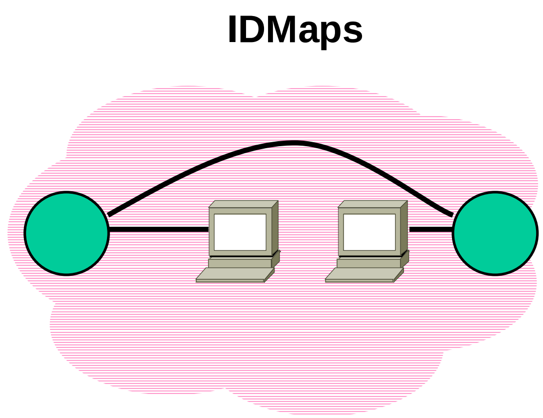

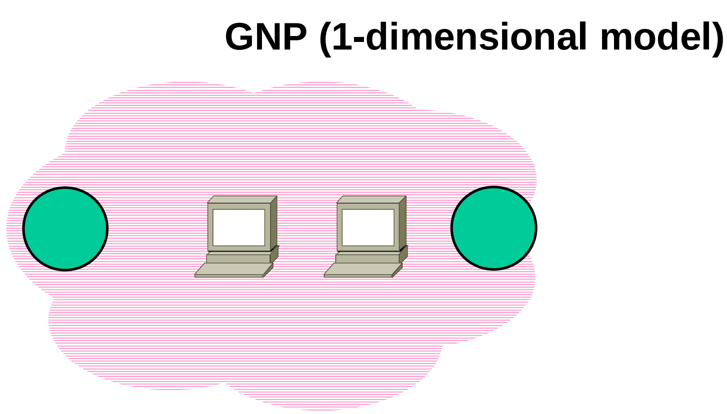

Straight-line distance
used in this example

Why the Difference in Over-Predictions?

Triangulated Heuristic

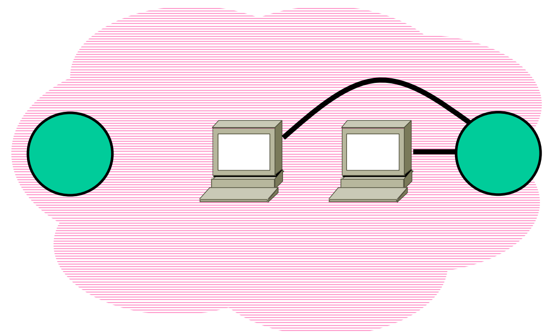

IDMaps

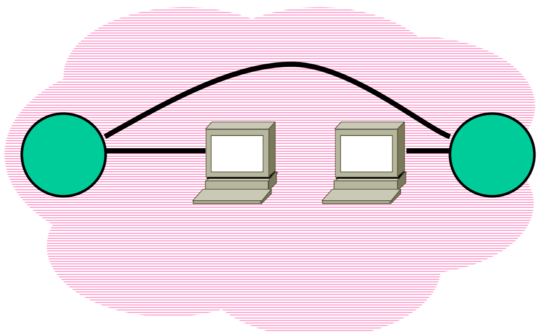

GNP (1-dimensional model)

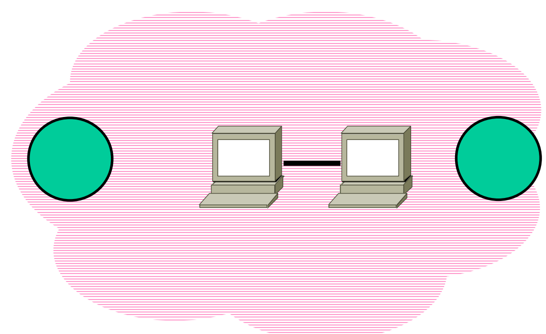

Straight-line distance
used in this example

Sensitivity to Infrastructure Node Placement

- Which nodes are used as Tracers/Bases/Landmarks matter
- High sensitivity means the approach is less robust
- Test sensitivity by picking 20 random combinations of 6 infrastructure nodes and observe performance variance

6 Random Infrastructure Nodes (20 Experiments)

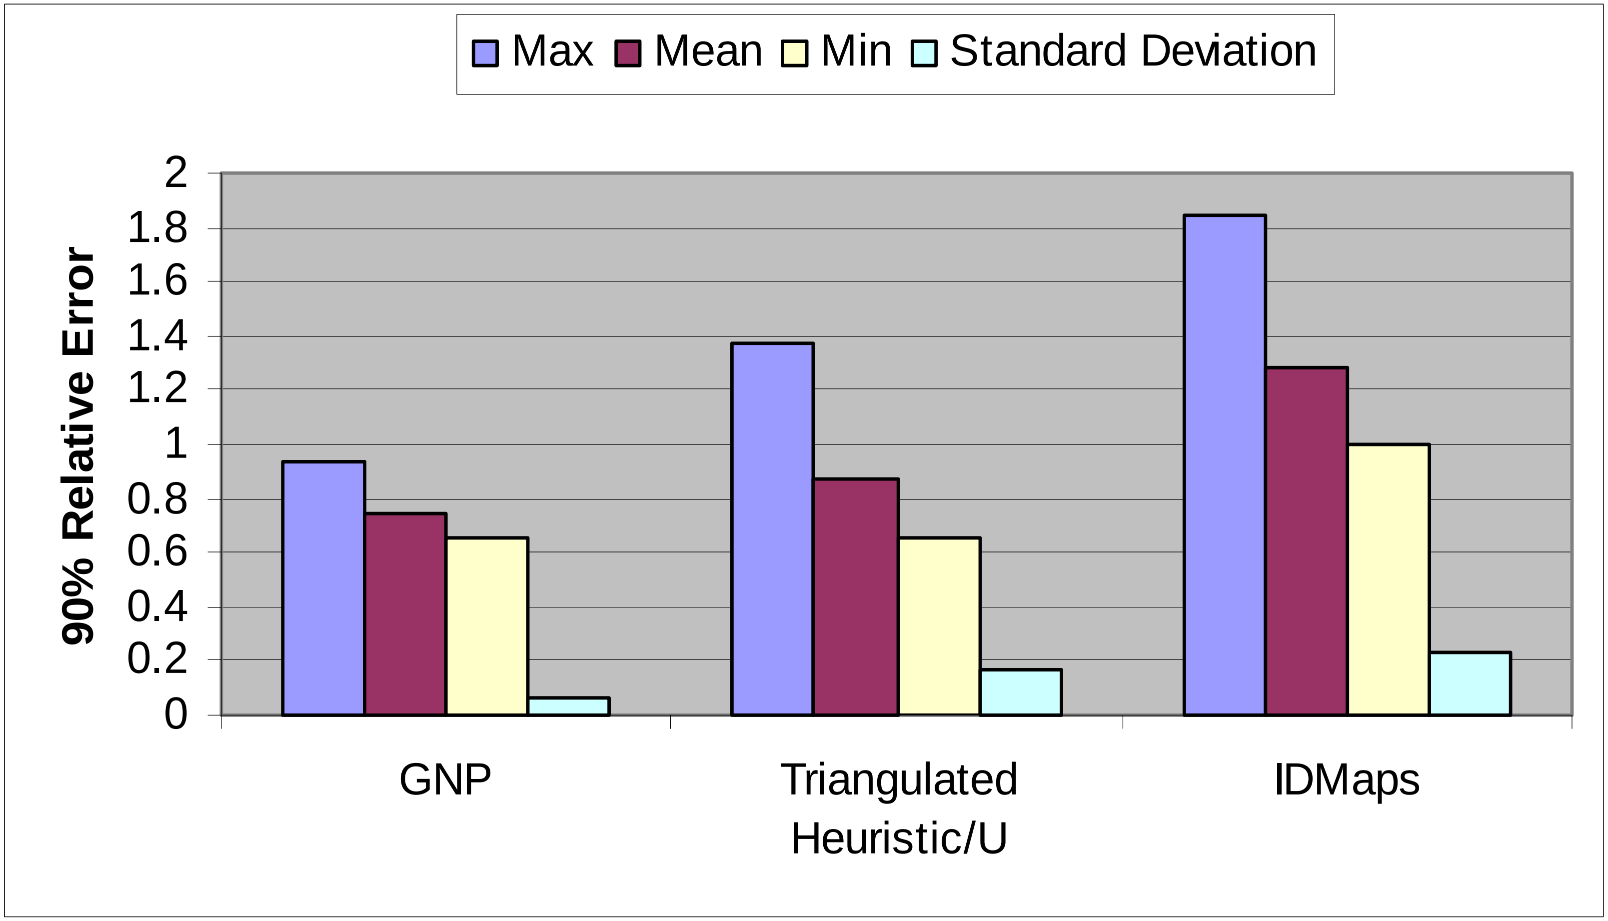

Conclusions

- Coordinates-based approaches represent a new class of solutions
- These solutions fit well with the peer-to-peer architecture
  - Potentially better performance and scalability than the client-server architecture
- Careful Internet evaluation shows that coordinates-based approaches are more accurate than IDMaps
- GNP is the most accurate and robust solution

Internet as an Euclidean Space? Remarkable!

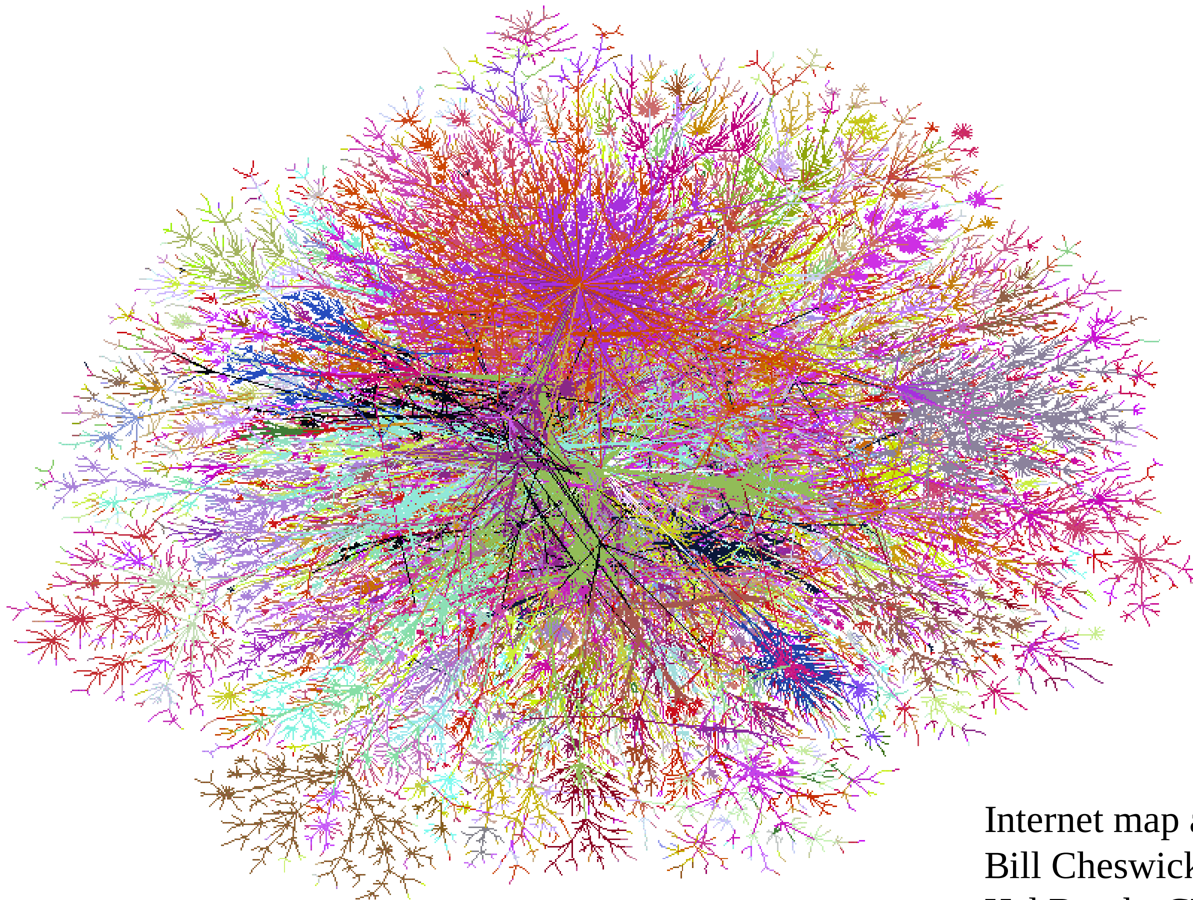

Internet map as of 1998 by
Bill Cheswick, *Bell Labs*
Hal Burch, *CMU*