Towards Global Network Positioning

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

New Challenges

- Large-scale distributed services and applications
  - Napster, Gnutella, End System Multicast, etc
- Large number of configuration choices
- K participants ⇒ $O(K^2)$ e2e paths to consider

New Challenges

- Large-scale distributed services and applications
  - Napster, Gnutella, End System Multicast, etc
- Large number of configuration choices
- K participants ⇒ $O(K^2)$ e2e paths to consider

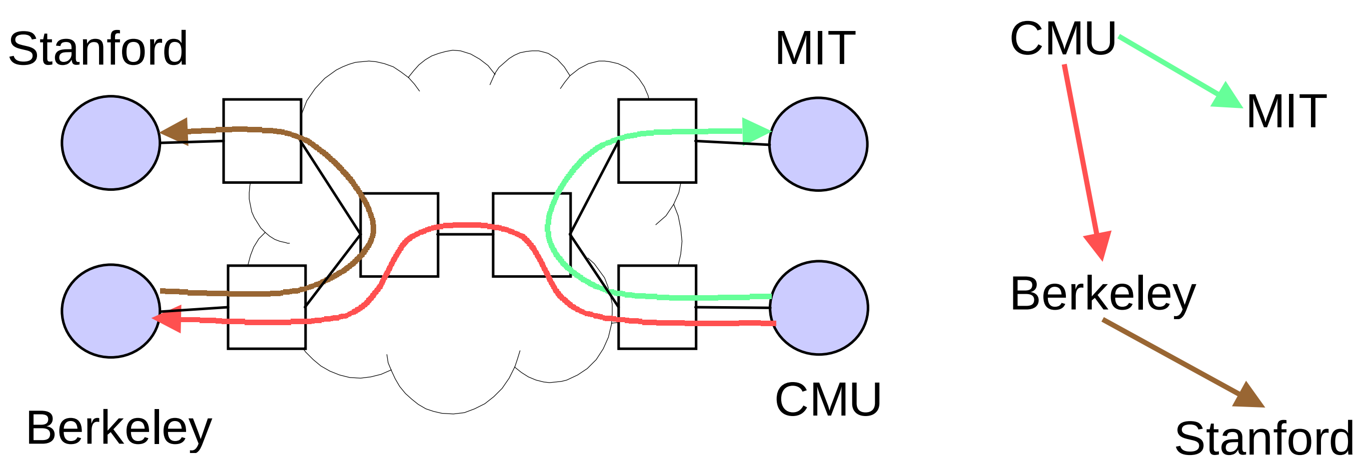

New Challenges

- Large-scale distributed services and applications
  - Napster, Gnutella, End System Multicast, etc
- Large number of configuration choices
- K participants ⇒ $O(K^2)$ e2e paths to consider

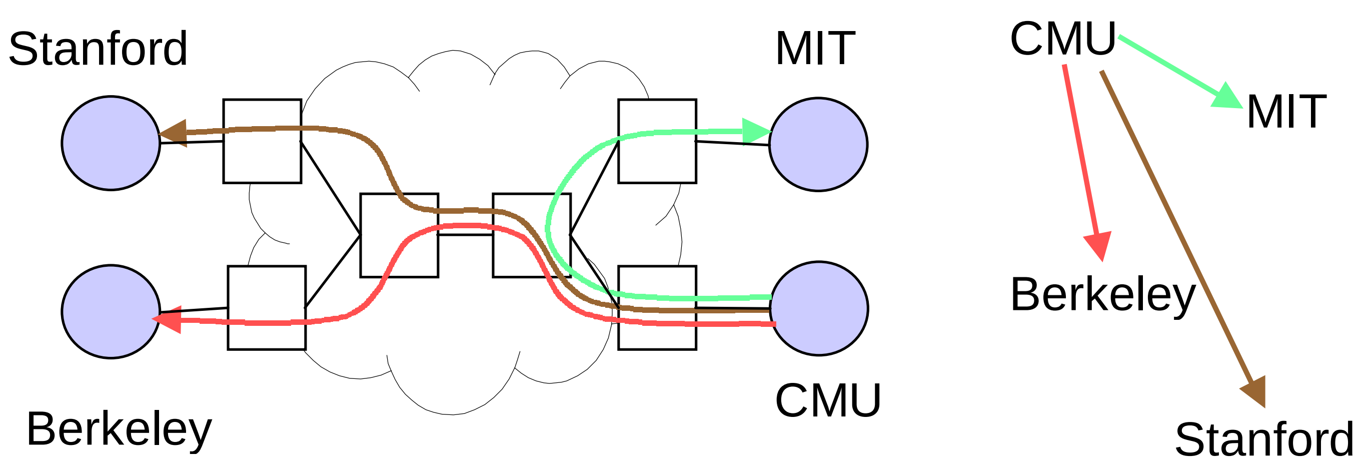

New Challenges

- Large-scale distributed services and applications
  - Napster, Gnutella, End System Multicast, etc
- Large number of configuration choices
- K participants ⇒ $O(K^2)$ e2e paths to consider

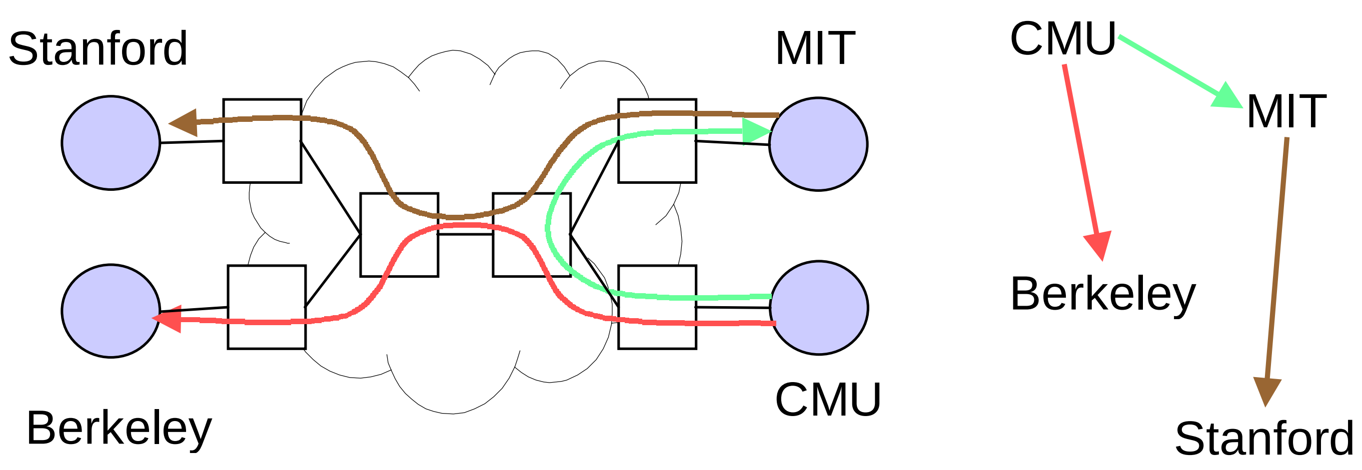

Role of Network Distance Prediction

- On-demand network measurement can be highly accurate, but
  - Not scalable
  - Slow
- Network distance
  - Round-trip propagation and transmission delay
  - Relatively stable
- Network distance can be predicted accurately without on-demand measurement
  - Fast and scalable first-order performance optimization
  - Refine as needed

State of the Art: IDMaps [Francis et al ‘99]

- A network distance prediction service

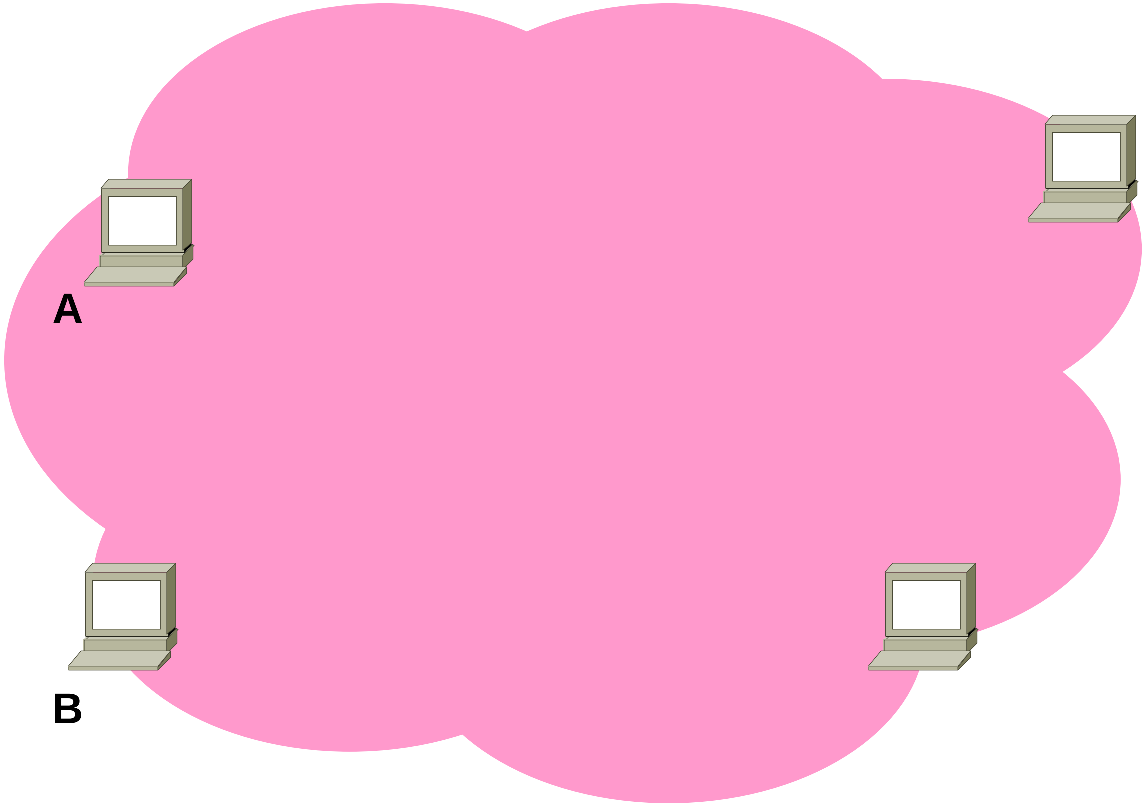

State of the Art: IDMaps [Francis et al '99]

- A network distance prediction service

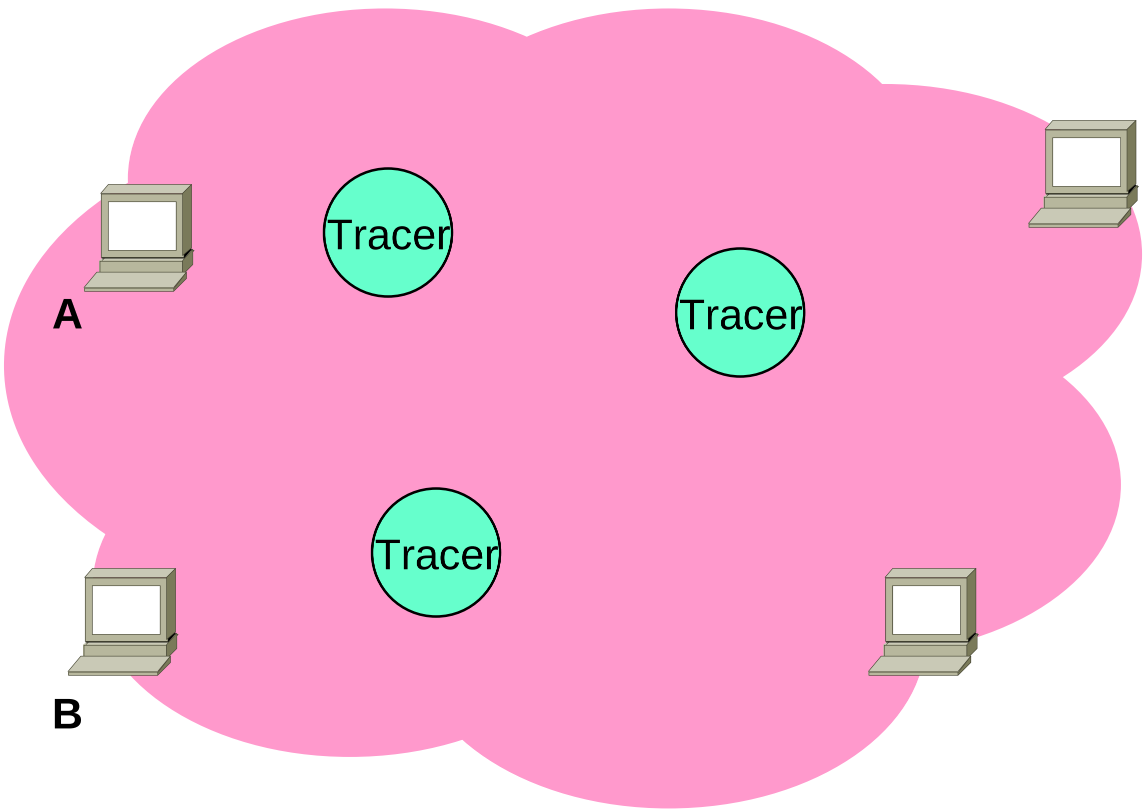

State of the Art: IDMaps [Francis et al '99]

- A network distance prediction service

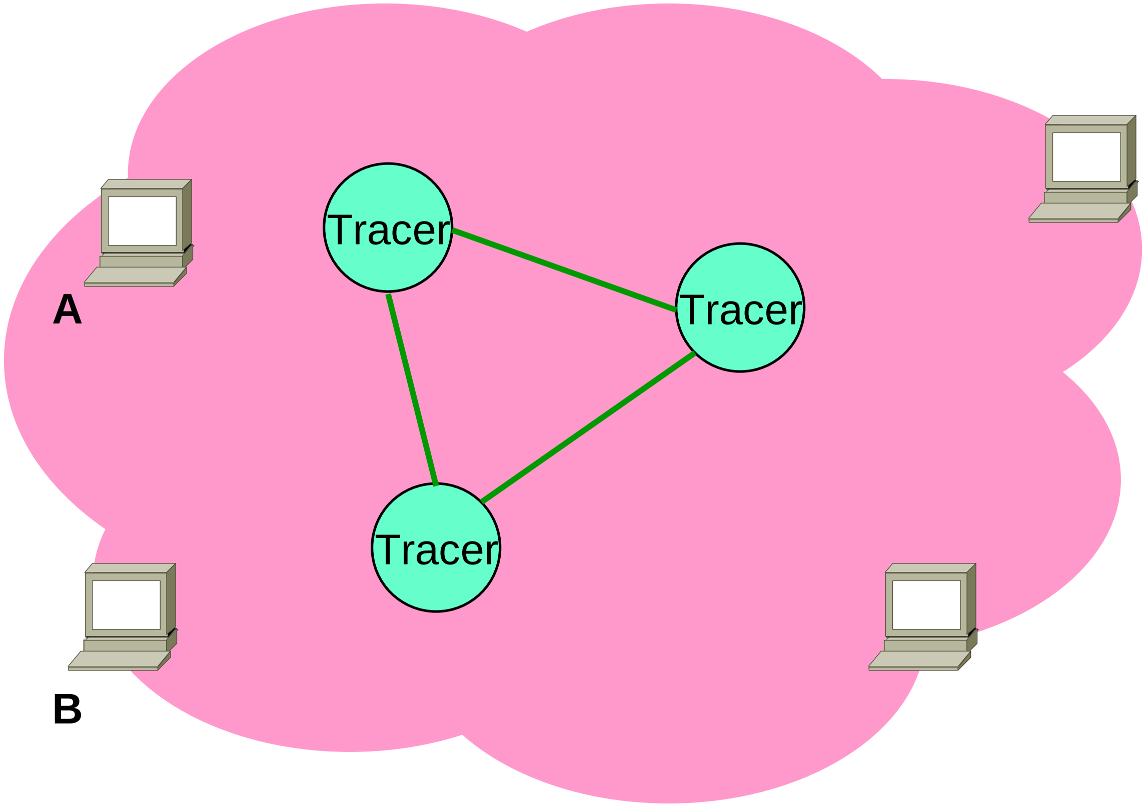

State of the Art: IDMaps [Francis et al ‘99]

- A network distance prediction service

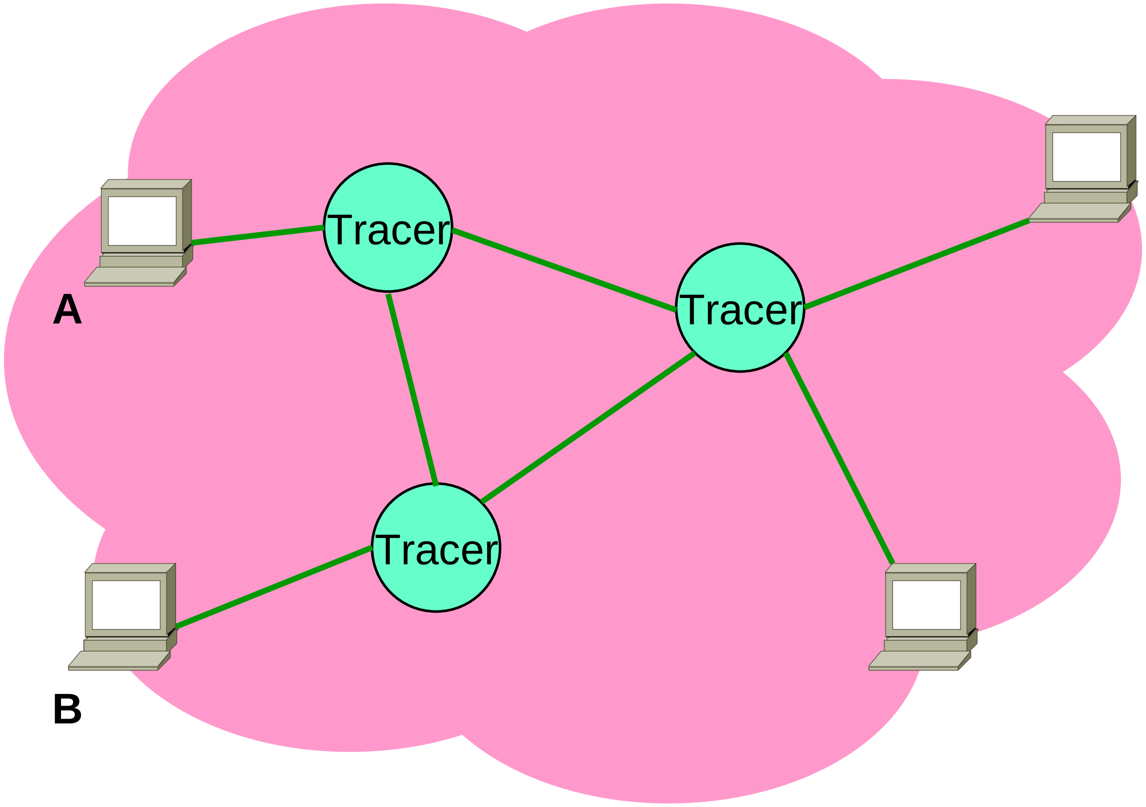

State of the Art: IDMaps [Francis et al ‘99]

- A network distance prediction service

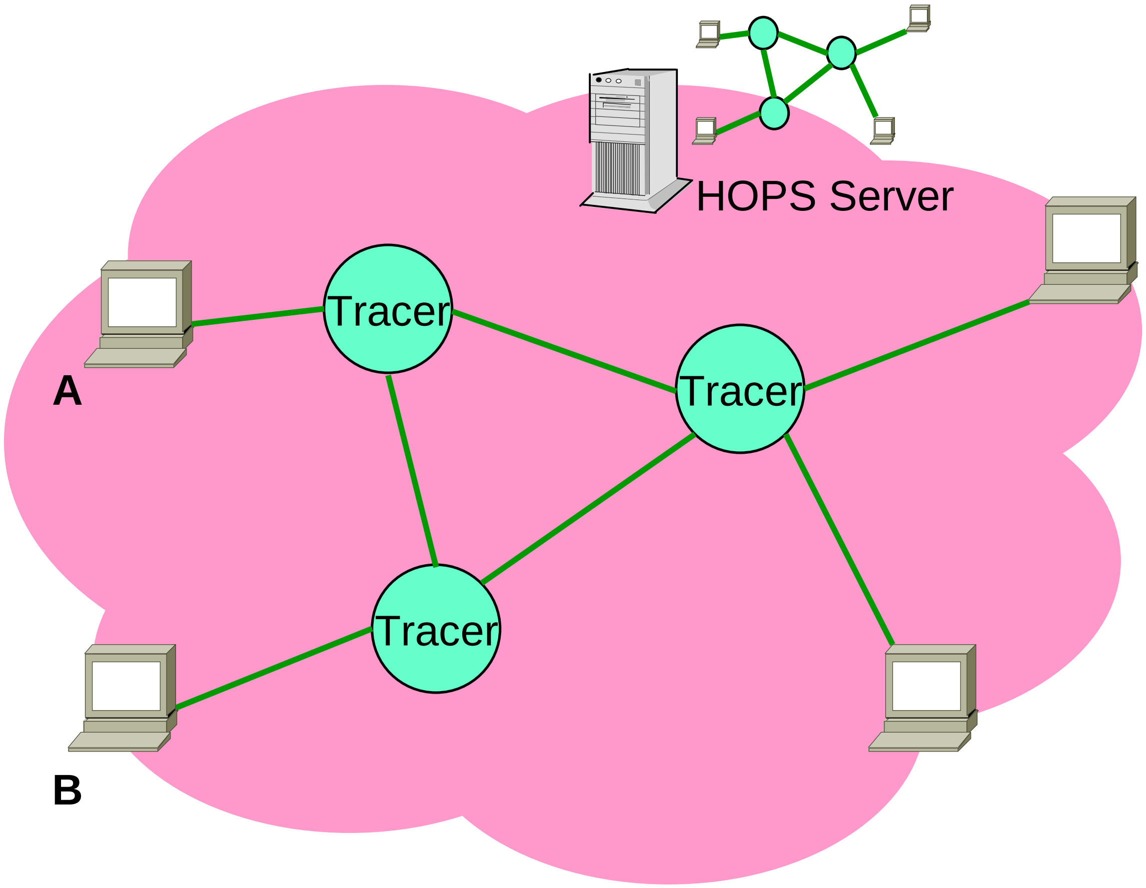

State of the Art: IDMaps [Francis et al '99]

- A network distance prediction service

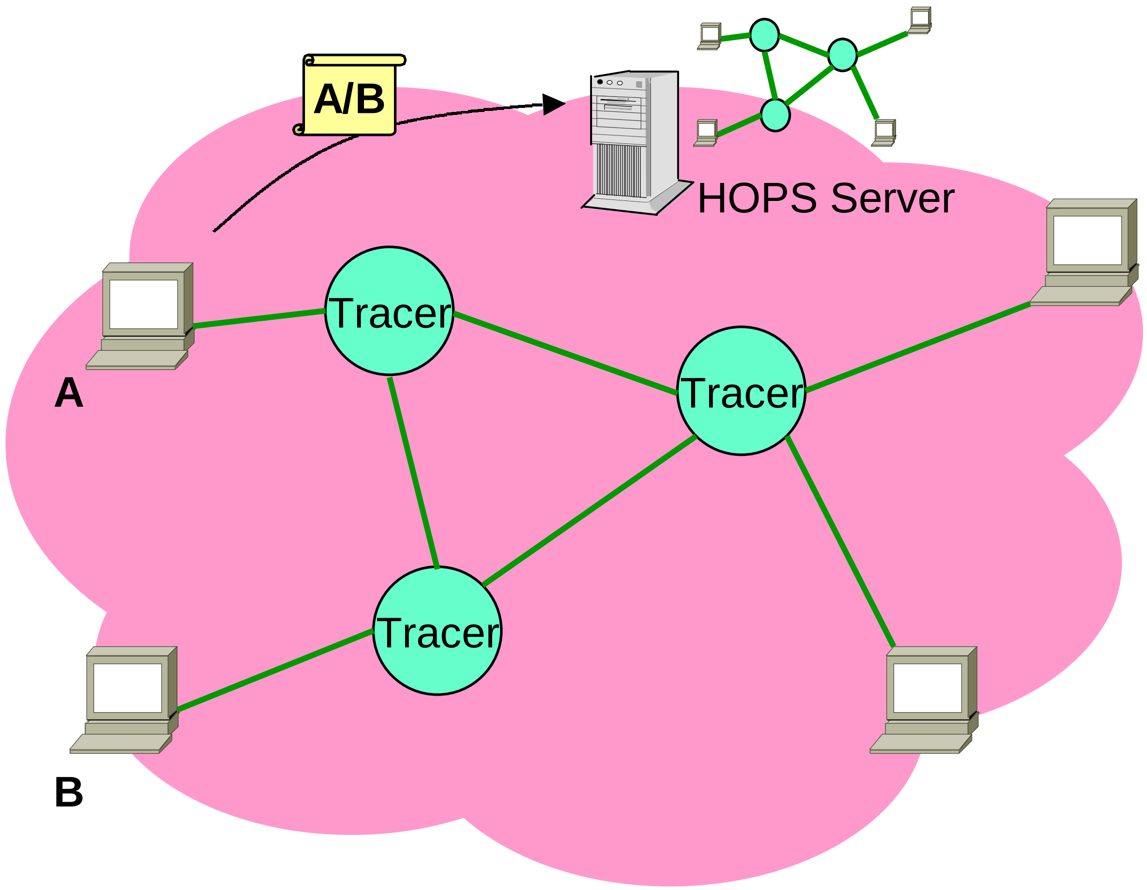

State of the Art: IDMaps [Francis et al '99]

- A network distance prediction service

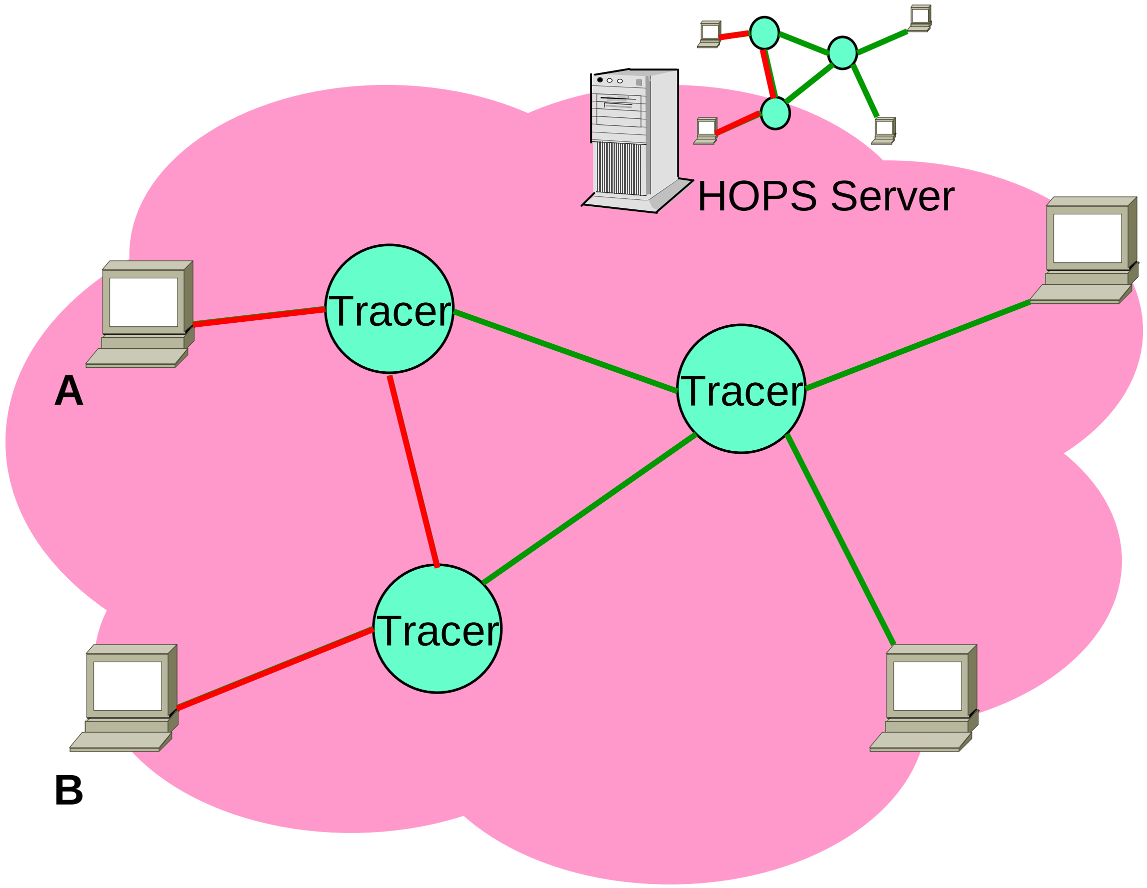

State of the Art: IDMaps [Francis et al '99]

- A network distance prediction service

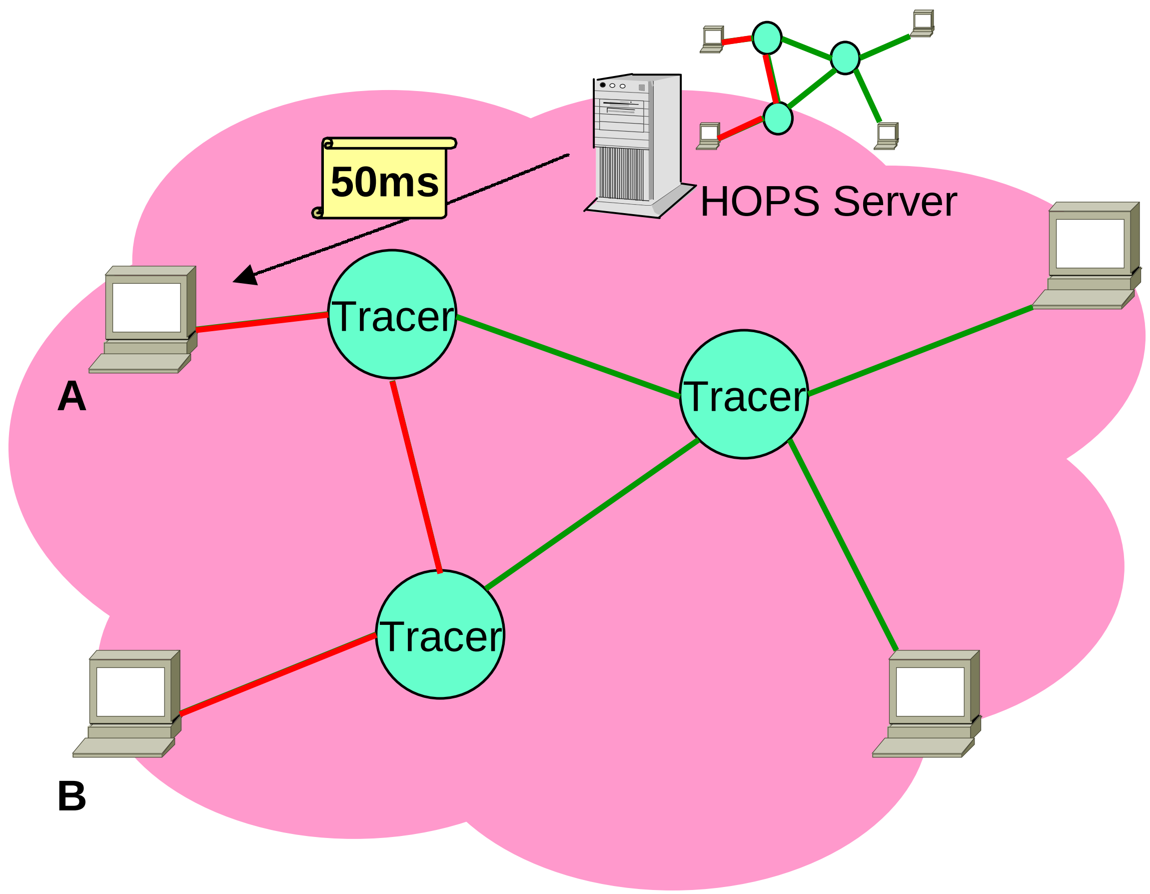

What Can be Improved?

- Scalability
- Speed
- Accuracy

Global Network Positioning (GNP)

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

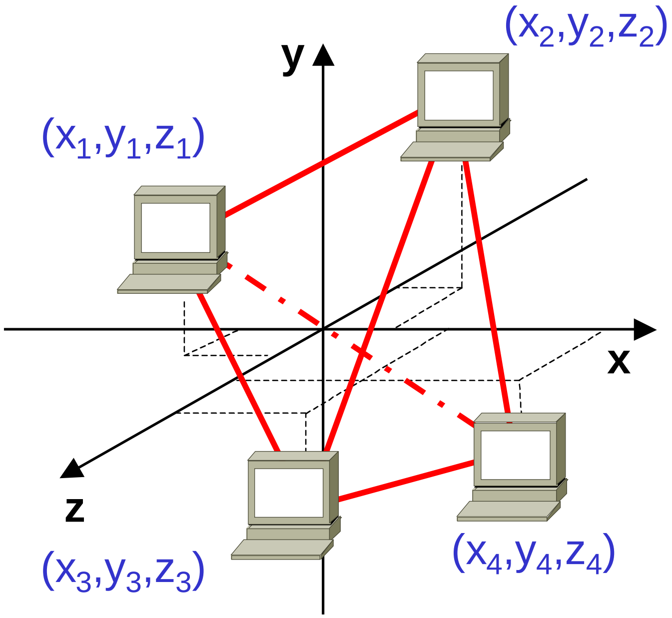

Landmark Operations

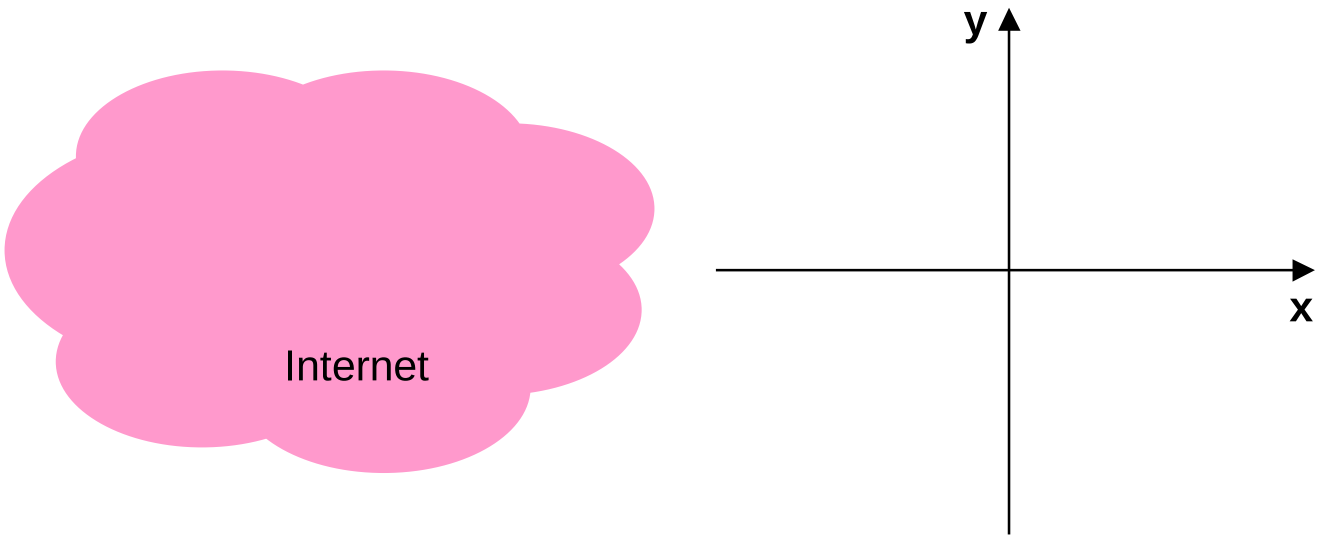

Landmark Operations

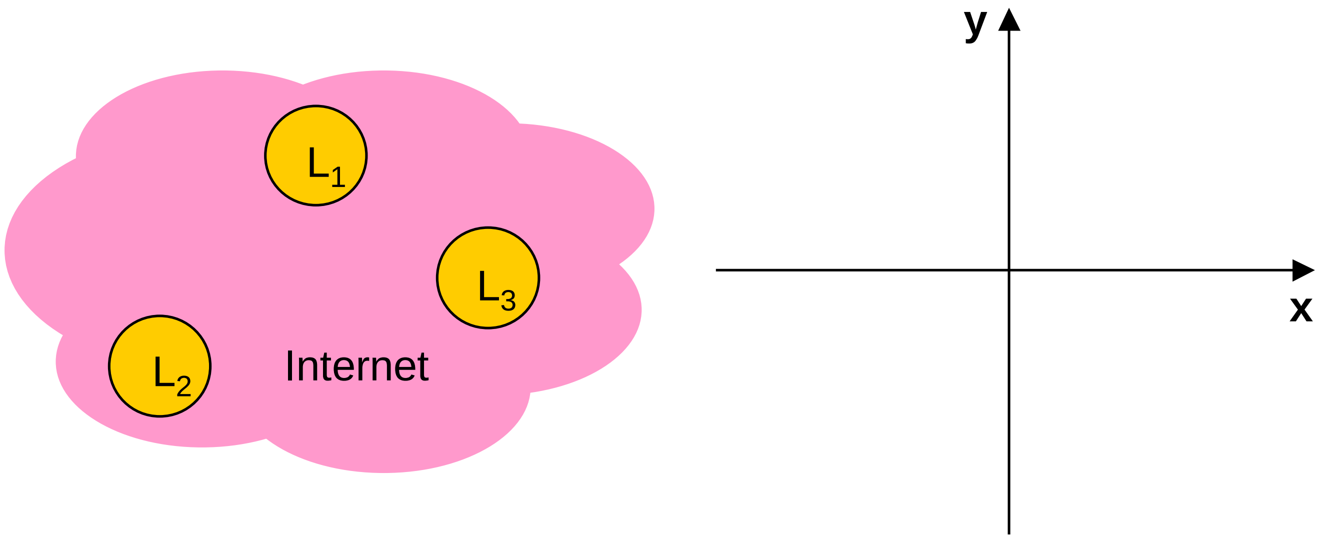

- Small number of distributed hosts called Landmarks measure inter-Landmark distances

Landmark Operations

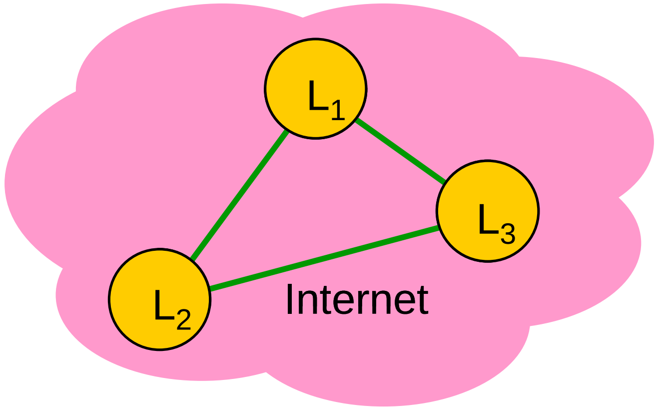

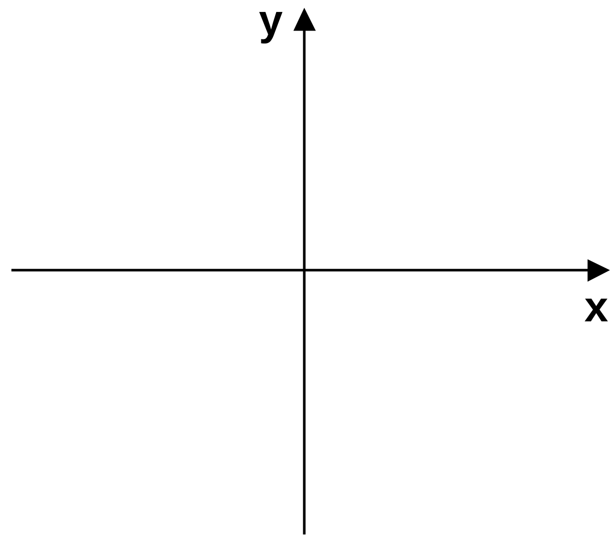

- Small number of distributed hosts called Landmarks measure inter-Landmark distances

Landmark Operations

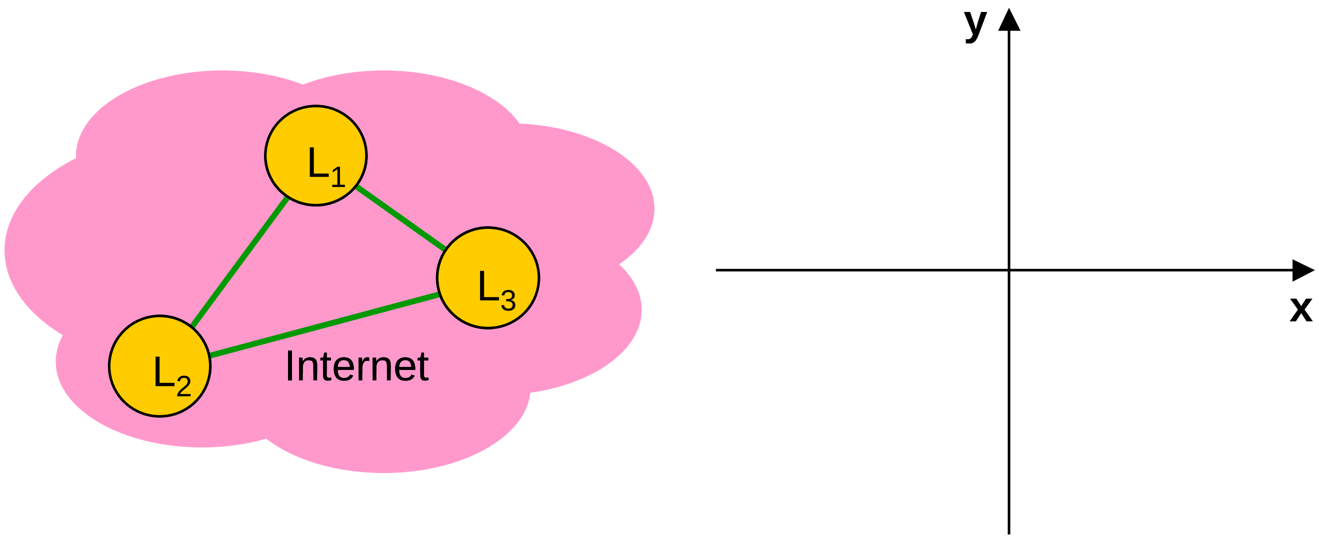

- Small number of distributed hosts called Landmarks measure inter-Landmark distances
- Compute Landmark coordinates by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

Landmark Operations

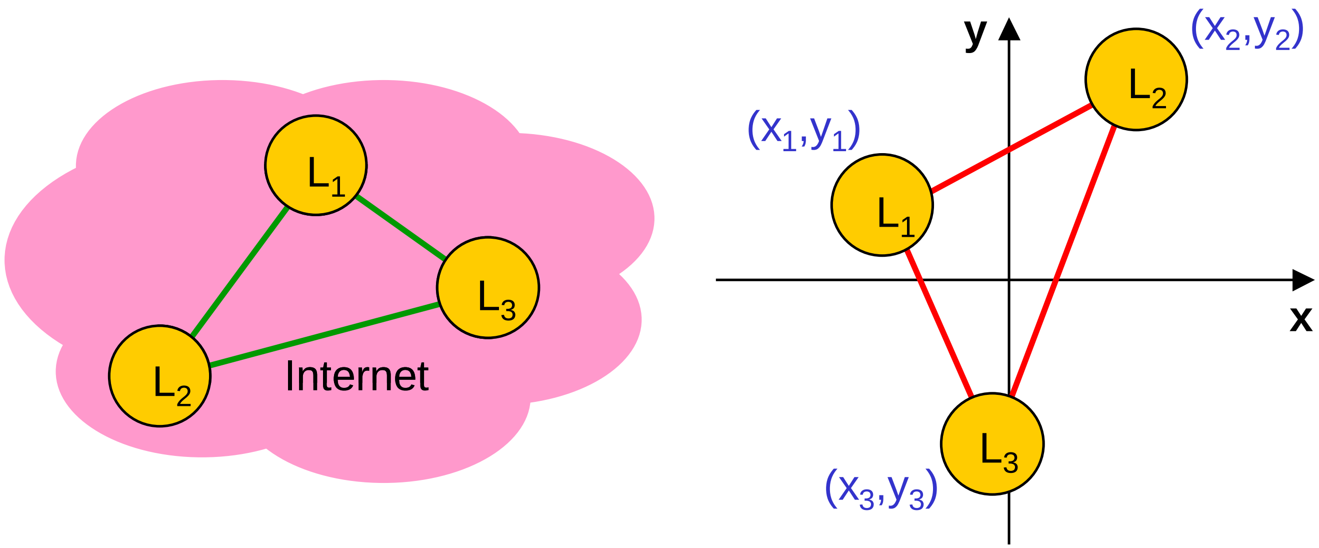

- Small number of distributed hosts called Landmarks measure inter-Landmark distances
- Compute Landmark coordinates by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization problem

Ordinary Host Operations

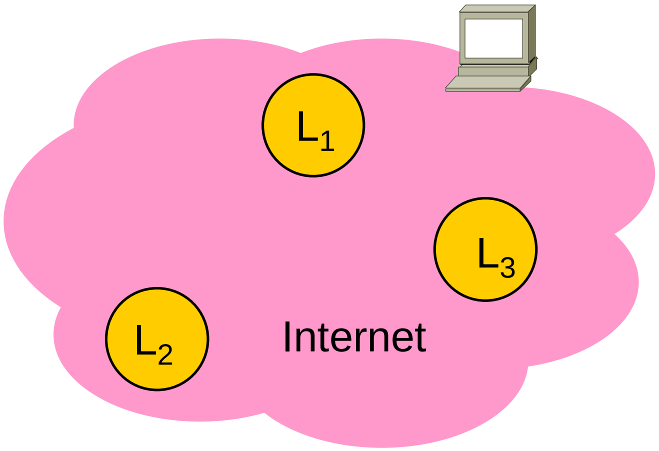

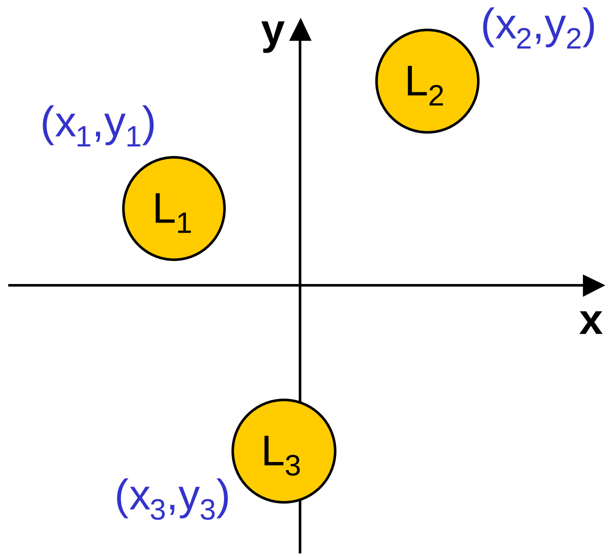

Ordinary Host Operations

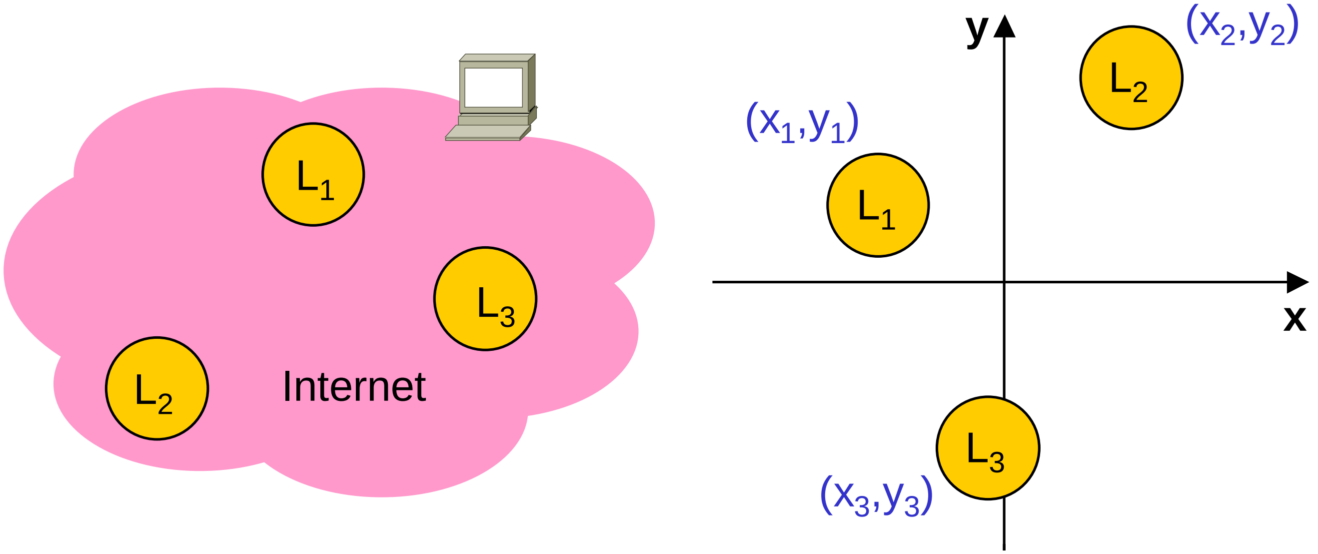

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings

Ordinary Host Operations

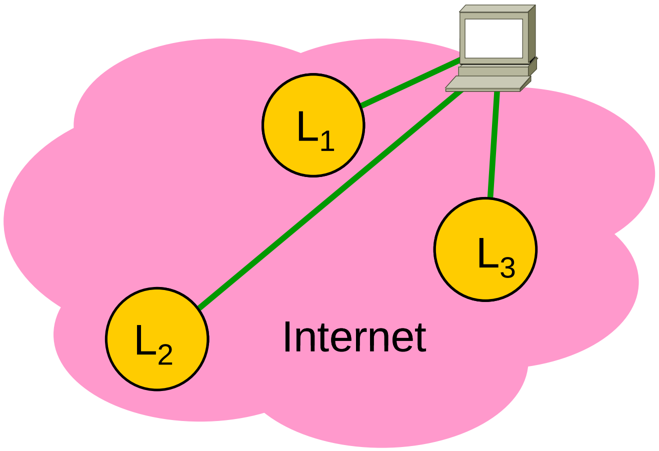

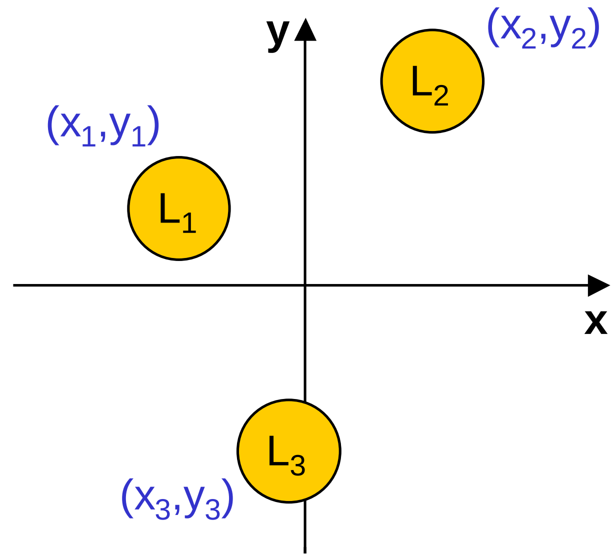

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings

Ordinary Host Operations

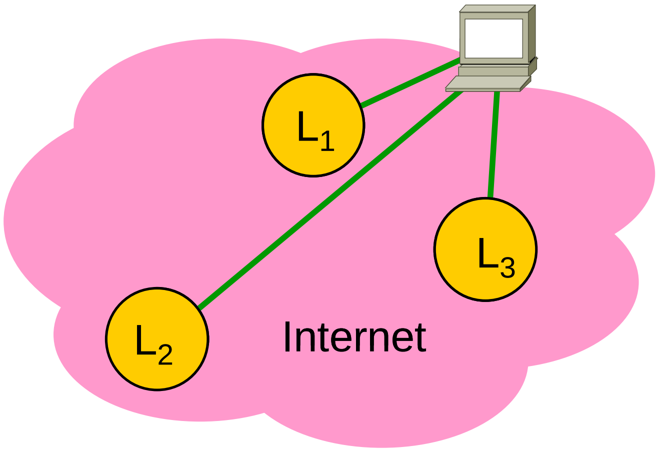

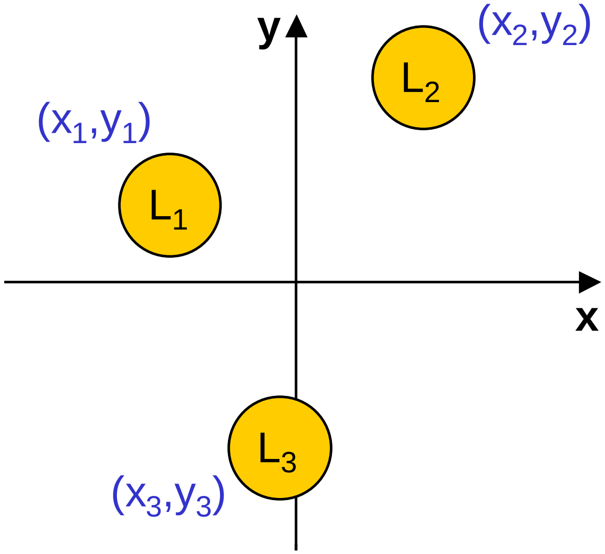

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings
- Ordinary host computes its own coordinates relative to the Landmarks by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization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

Ordinary Host Operations

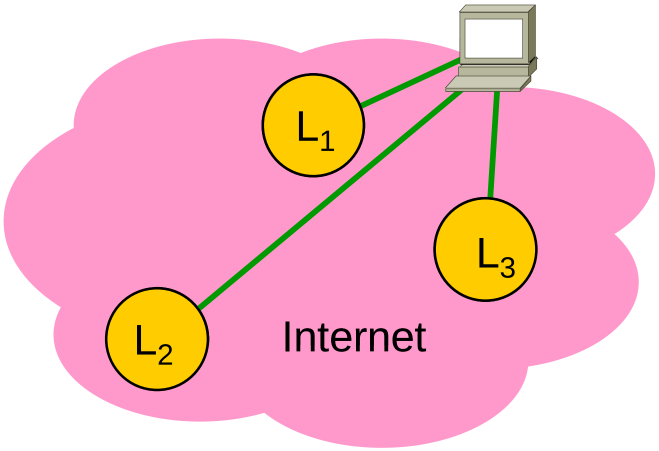

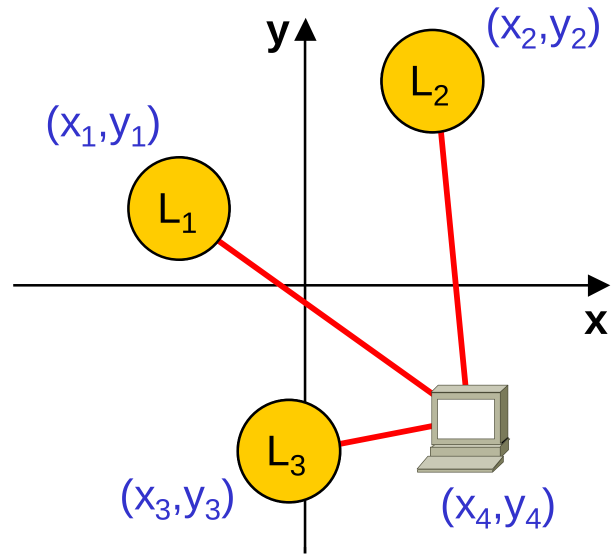

- Each ordinary host measures its distances to the Landmarks, Landmarks just reflect pings
- Ordinary host computes its own coordinates relative to the Landmarks by minimizing the overall discrepancy between measured distances and computed distances
  - Cast as a generic multi-dimensional global minimization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

GNP Advantages Over IDMaps

- High scalability and high speed
  - End host centric architecture, eliminates server bottleneck
  - Coordinates reduce $O(K^2)$ communication overhead to $O(K*D)$
  - Predictions are locally and quickly computable by end hosts
- Enable new applications
  - Structured nature of coordinates can be exploited
- Simple deployment
  - Landmarks are simple, non-intrusive (compatible with firewalls)

Evaluation Methodology

- 19 Probes we control
  - 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
  - Each distance is the minimum RTT of 220 pings

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

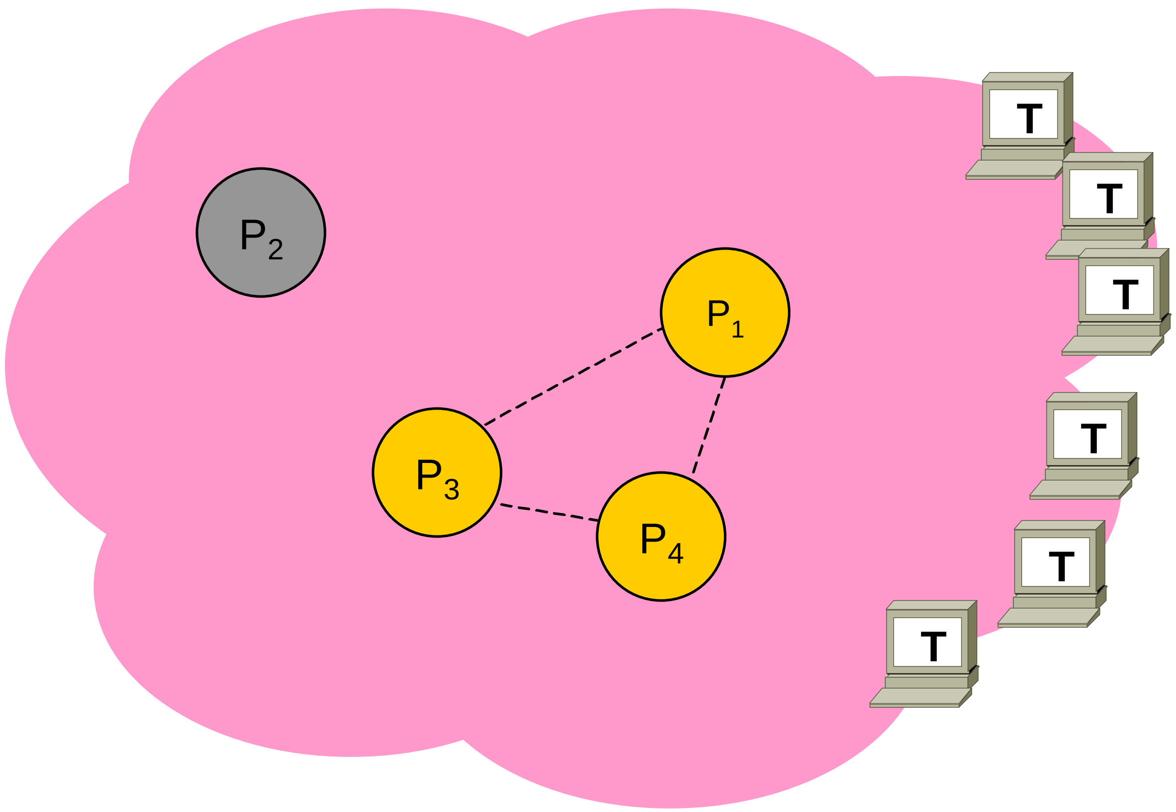

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

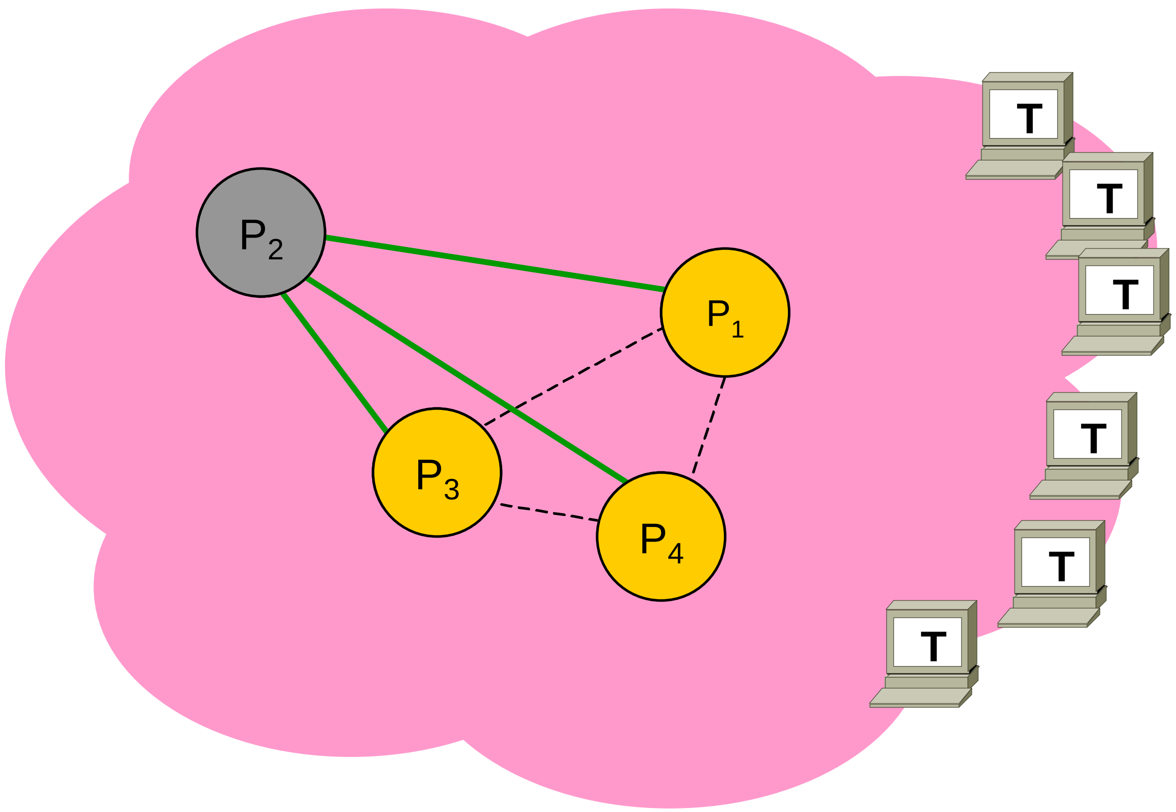

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

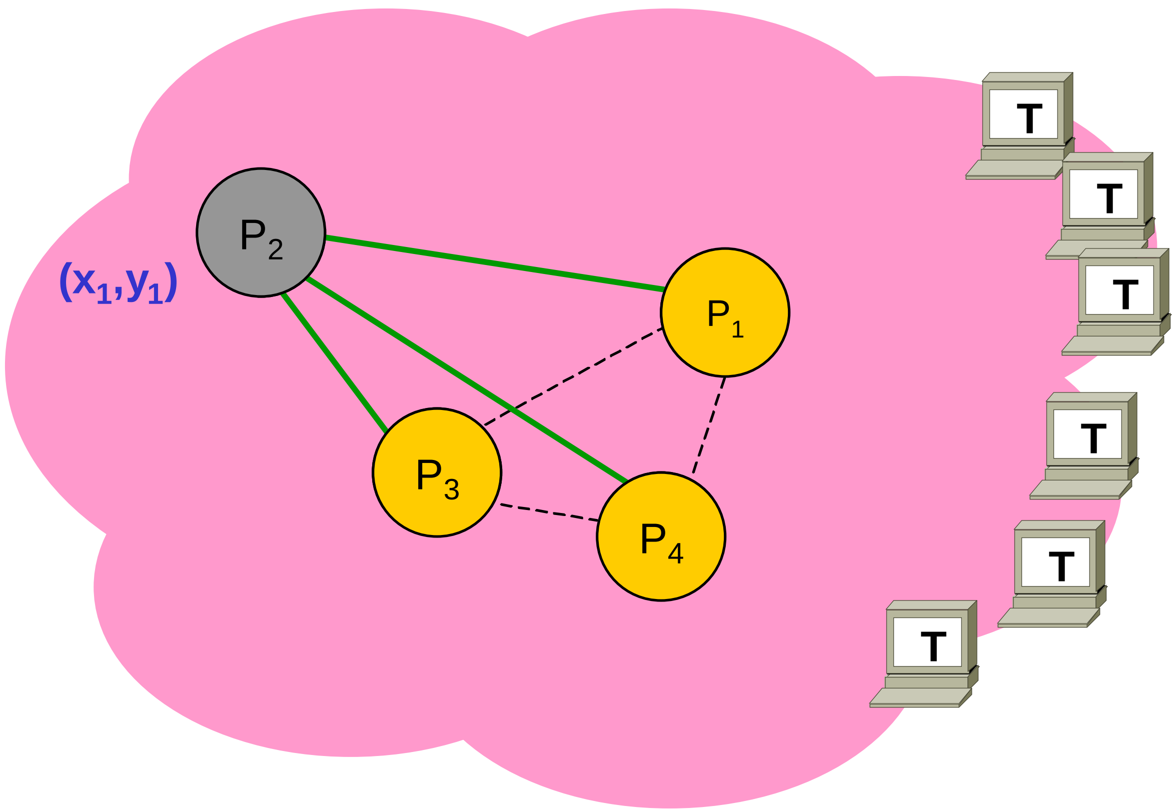

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

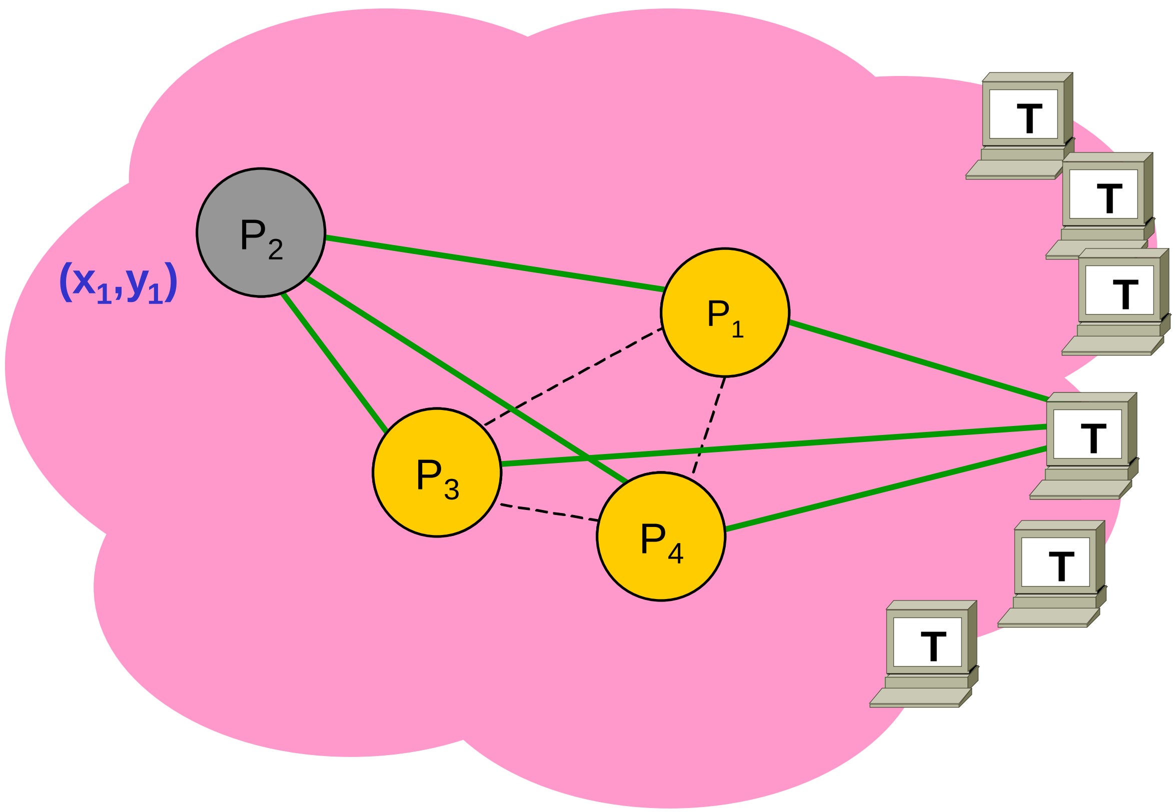

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

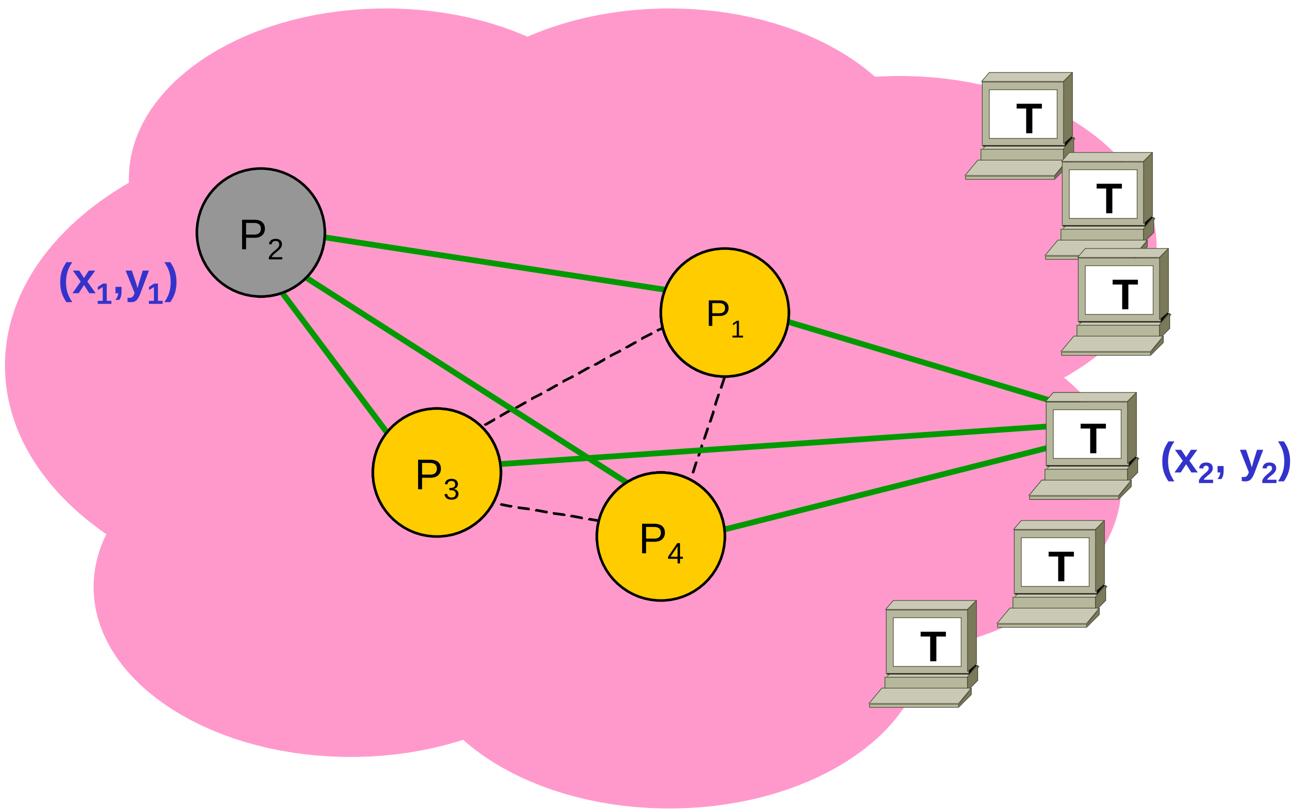

Evaluation Methodology (Cont'd)

- Choose a subset of well-distributed Probes to be Landmarks, and use the rest for evaluation

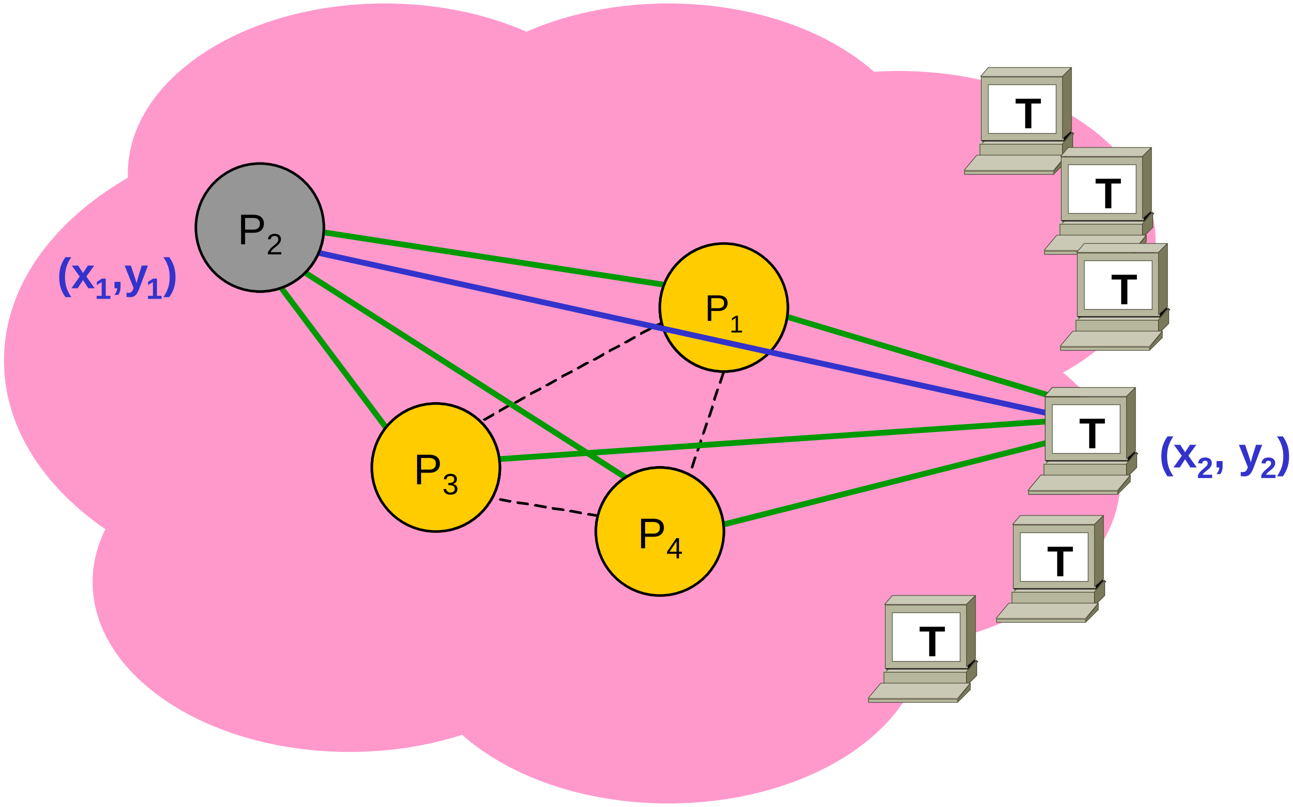

Performance Metric

- Relative error
  - Symmetrically measure over and under predictions

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

GNP Accuracy

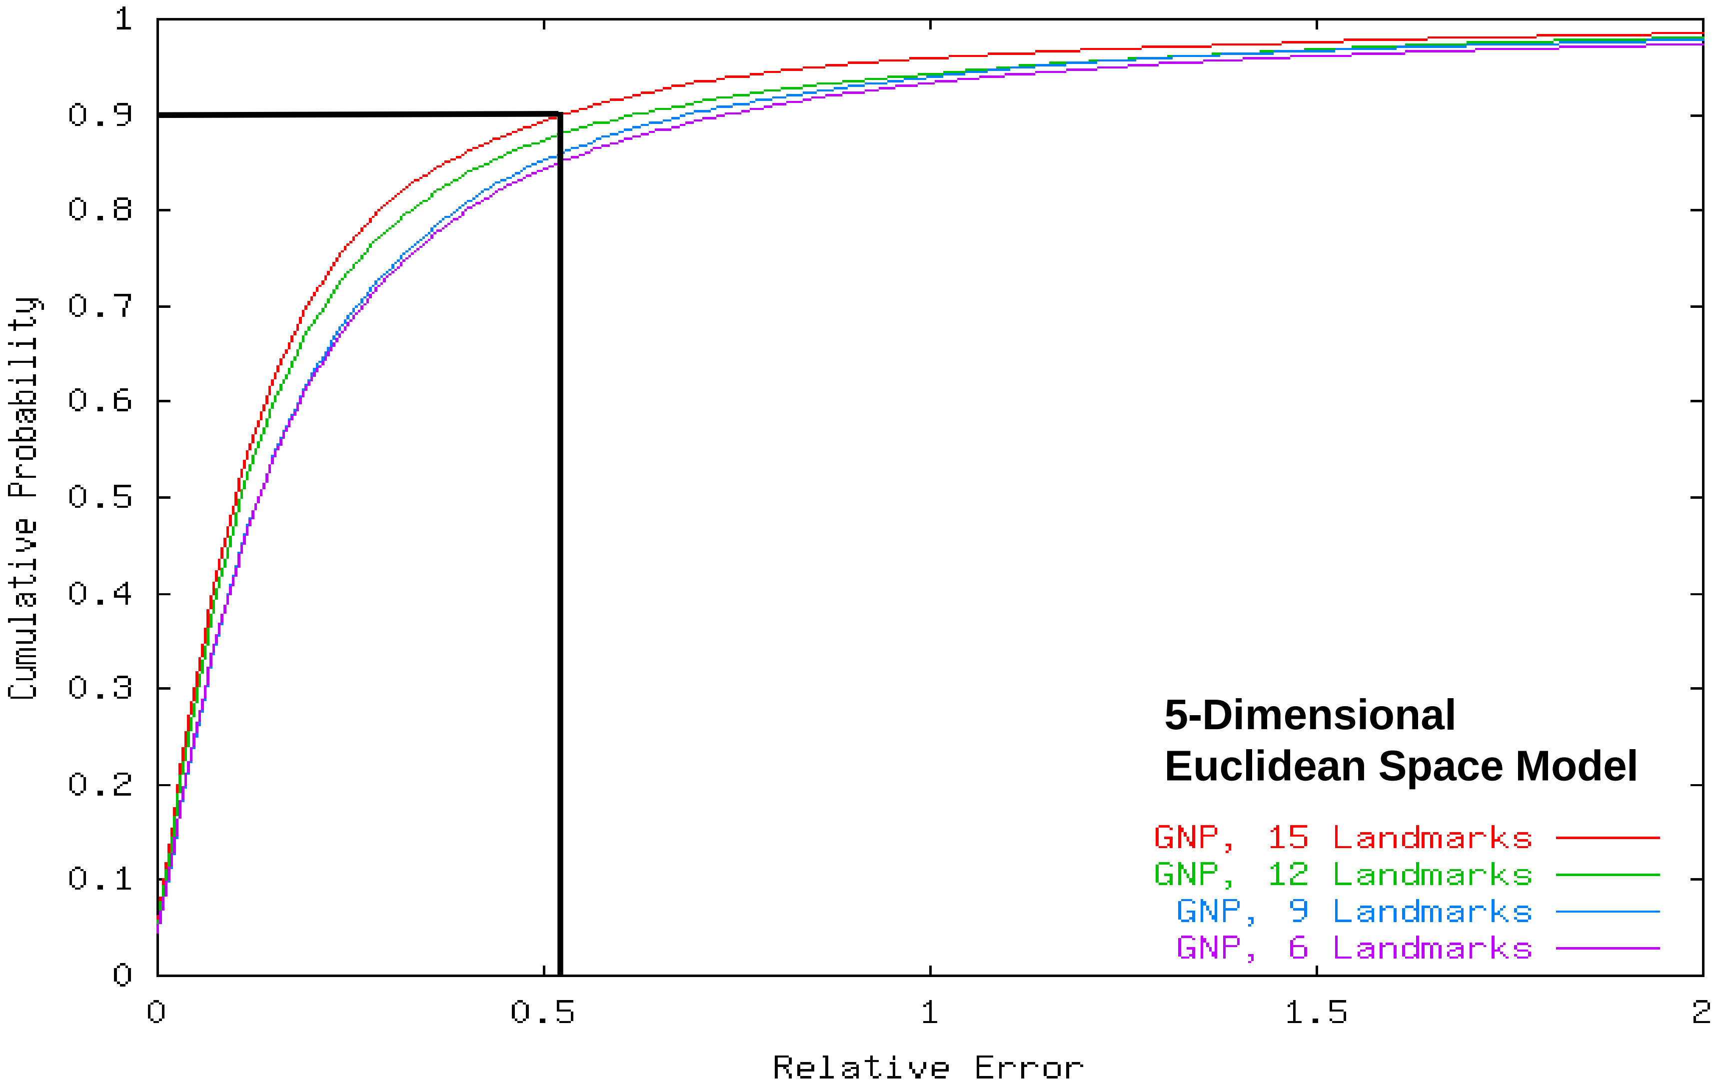

GNP vs IDMaps

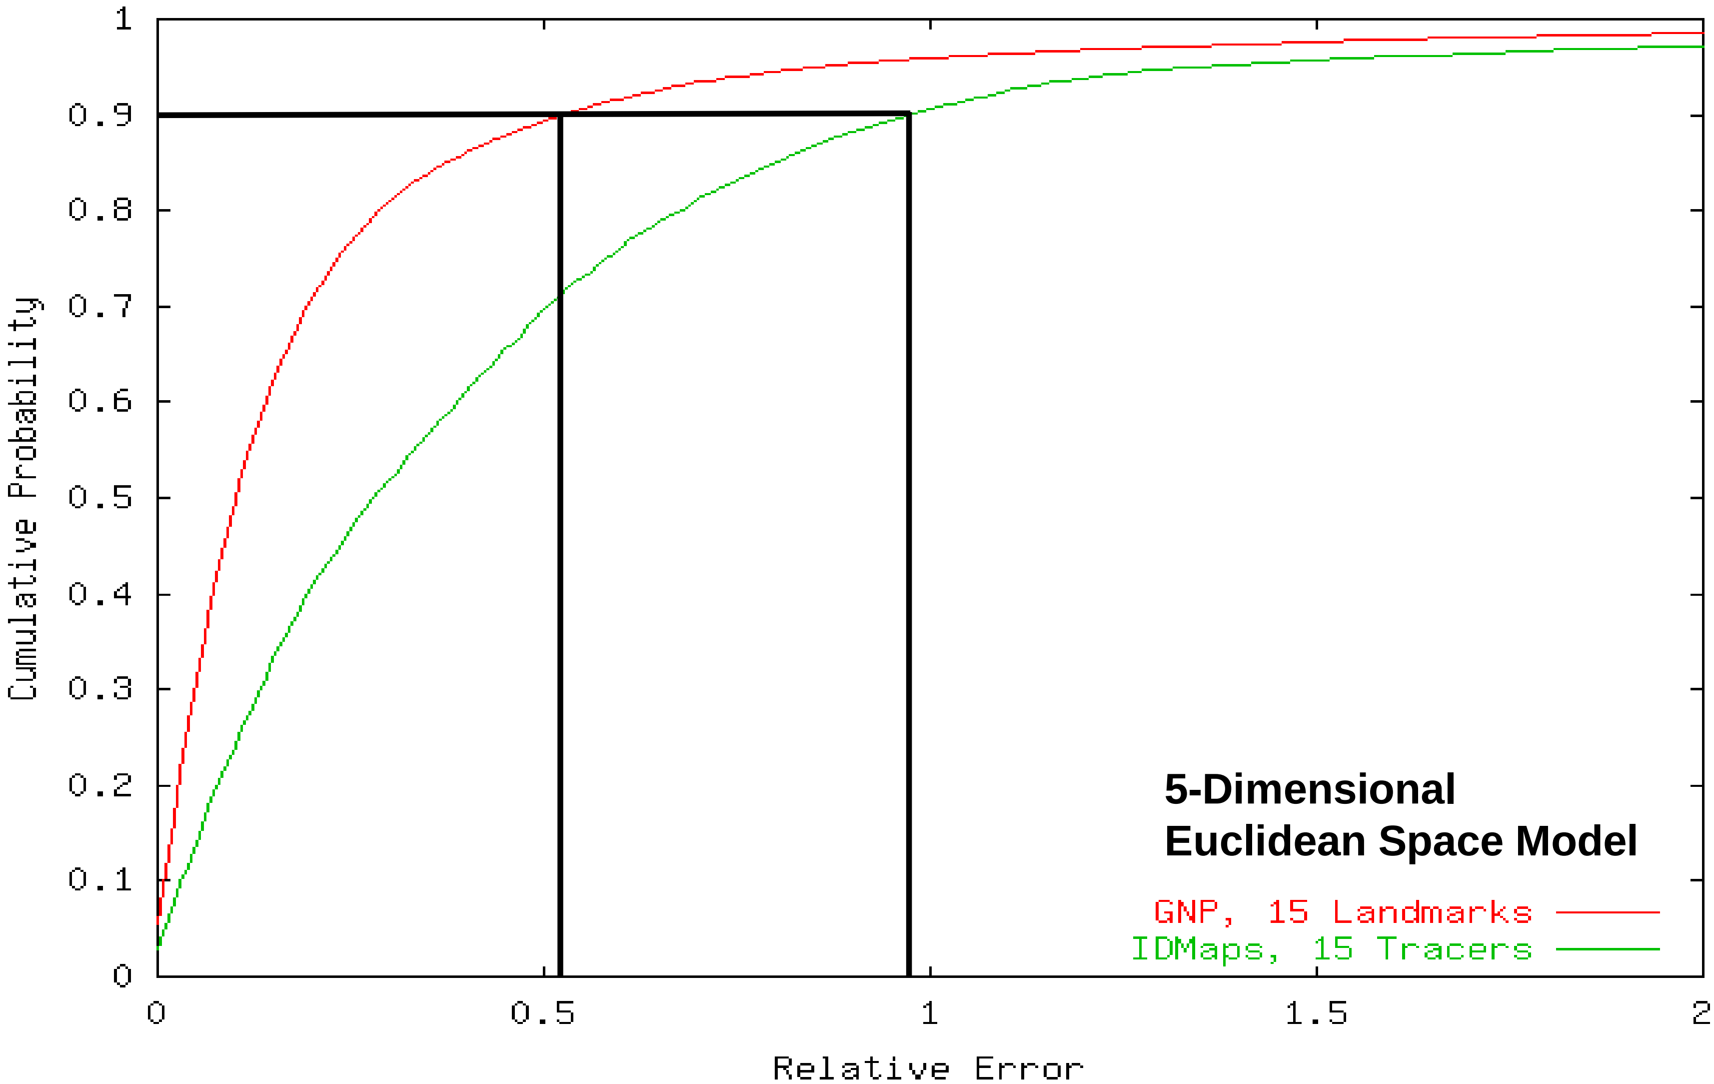

Why the Difference?

- IDMaps tends to heavily over-predict short distances
- Consider (measured ≤ 50ms)
  - 22% of all paths in evaluation
  - IDMaps on average over-predicts by 150 %
  - GNP on average over-predicts by 30%

Why the Difference?

- IDMaps tends to heavily over-predict short distances
- Consider (measured ≤ 50ms)
  - 22% of all paths in evaluation
  - IDMaps on average over-predicts by 150 %
  - GNP on average over-predicts by 30%

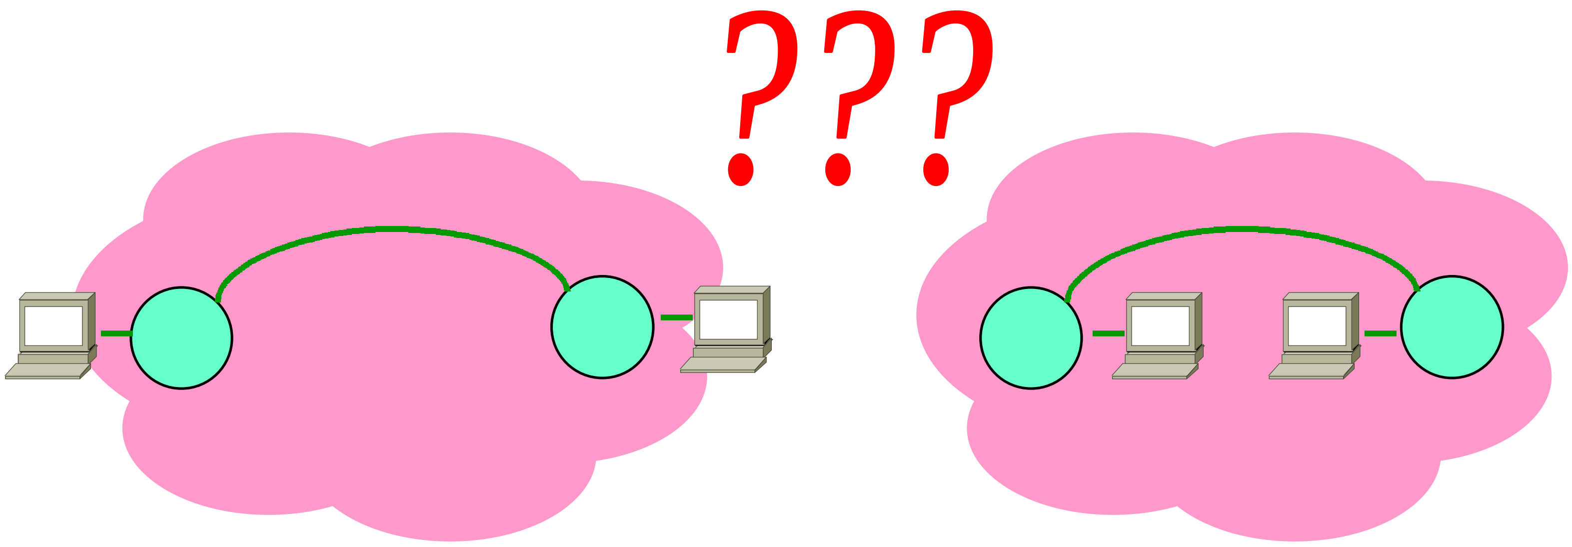

Summary

- Network distance prediction is key to performance optimization in large-scale distributed systems
- GNP is scalable
  - End hosts carry out computations
  - O(K*D) communication overhead due to coordinates
- GNP is fast
  - Distance predictions are fast local computations
- GNP is accurate
  - Discover relative positions of end hosts