Internet Coordinate Systems Tutorial

Marcelo Pias
Computer Laboratory
marcelo.pias@cl.cam.ac.uk

Outline

- **Motivation**
- Problem statement
- General Approach
- Techniques
  - Global Network Positioning (GNP)
  - Practical Internet Coordinates (PIC)
  - Lighthouses
  - PCA-based techniques (Virtual Landmark and ICS)
- Conclusions
- Open Issues

Motivation

What's the closest server to a client in Brazil?

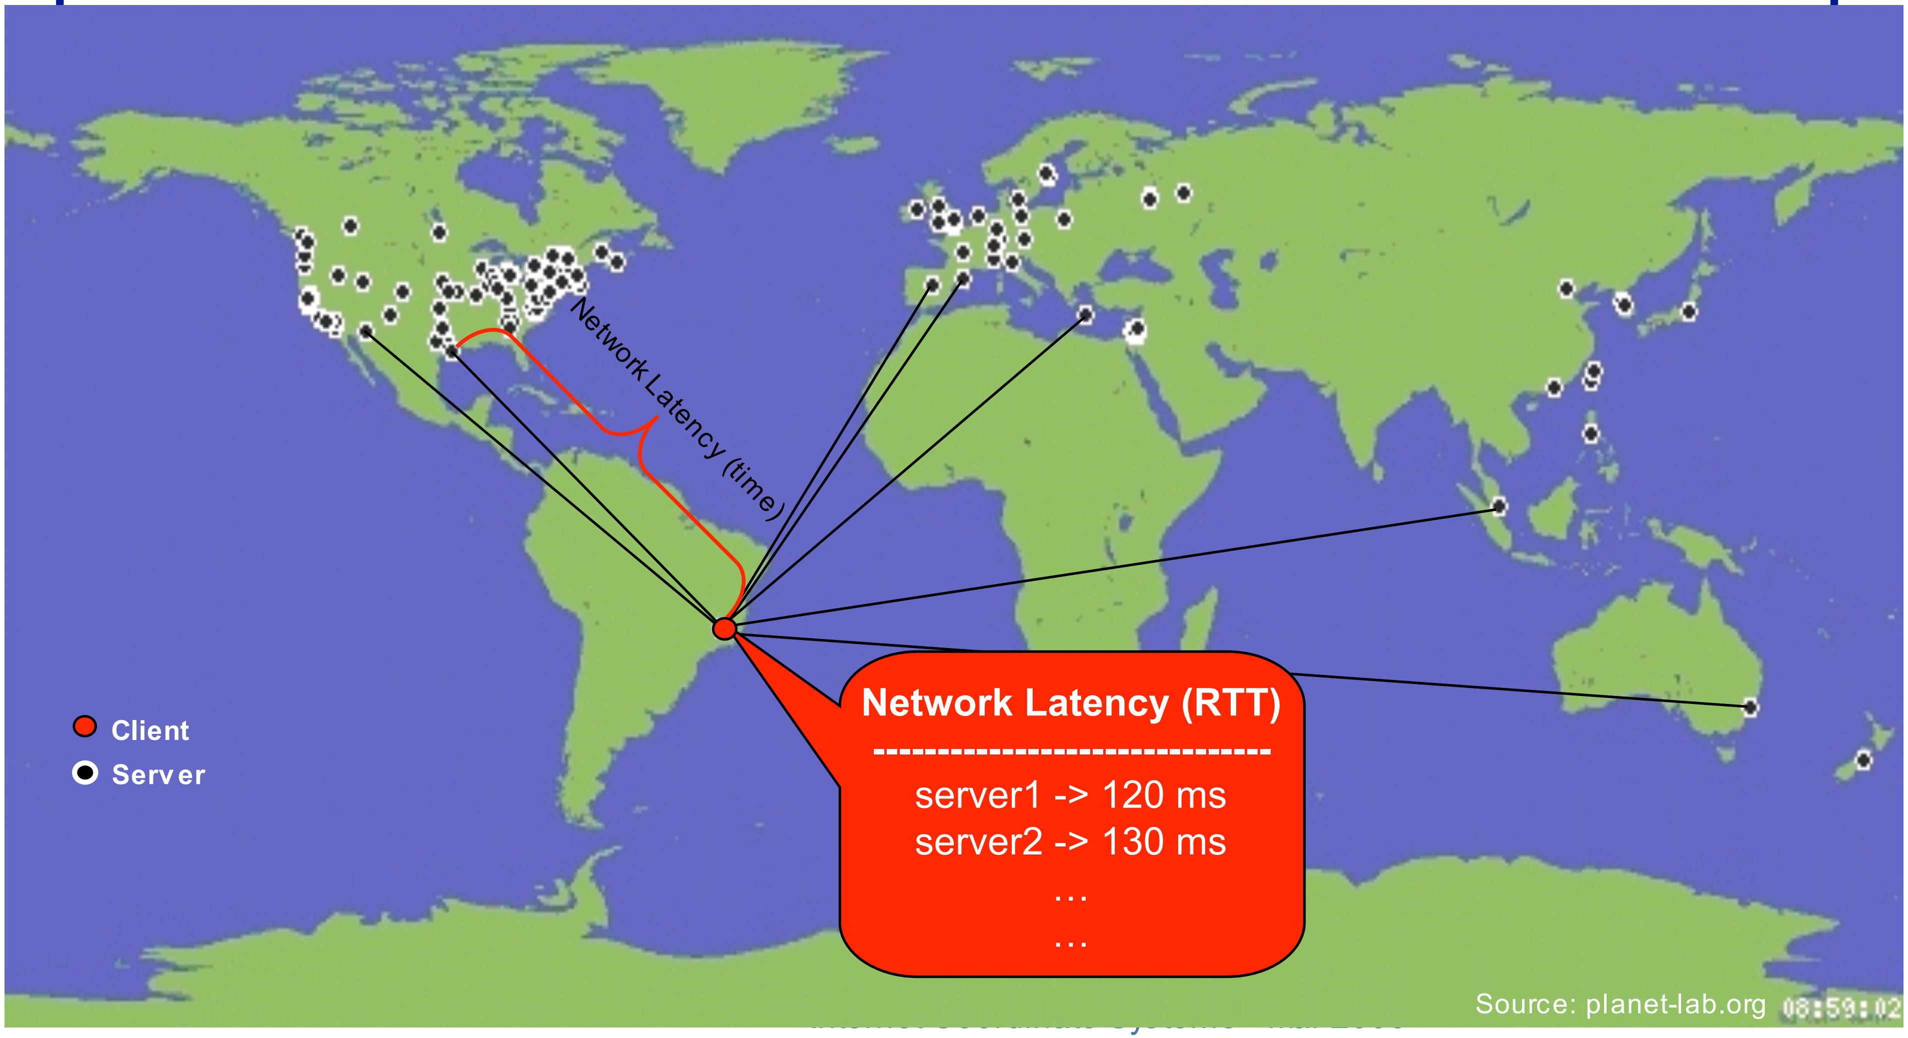

Motivation

- Network round-trip-time = network distance
  - E.g. ping measurements
- Issue
  - Client needs ‘N’ measurements to select the closest server
  - Update list of network distances (overhead)
- How do we solve this problem ?

Internet Coordinate System

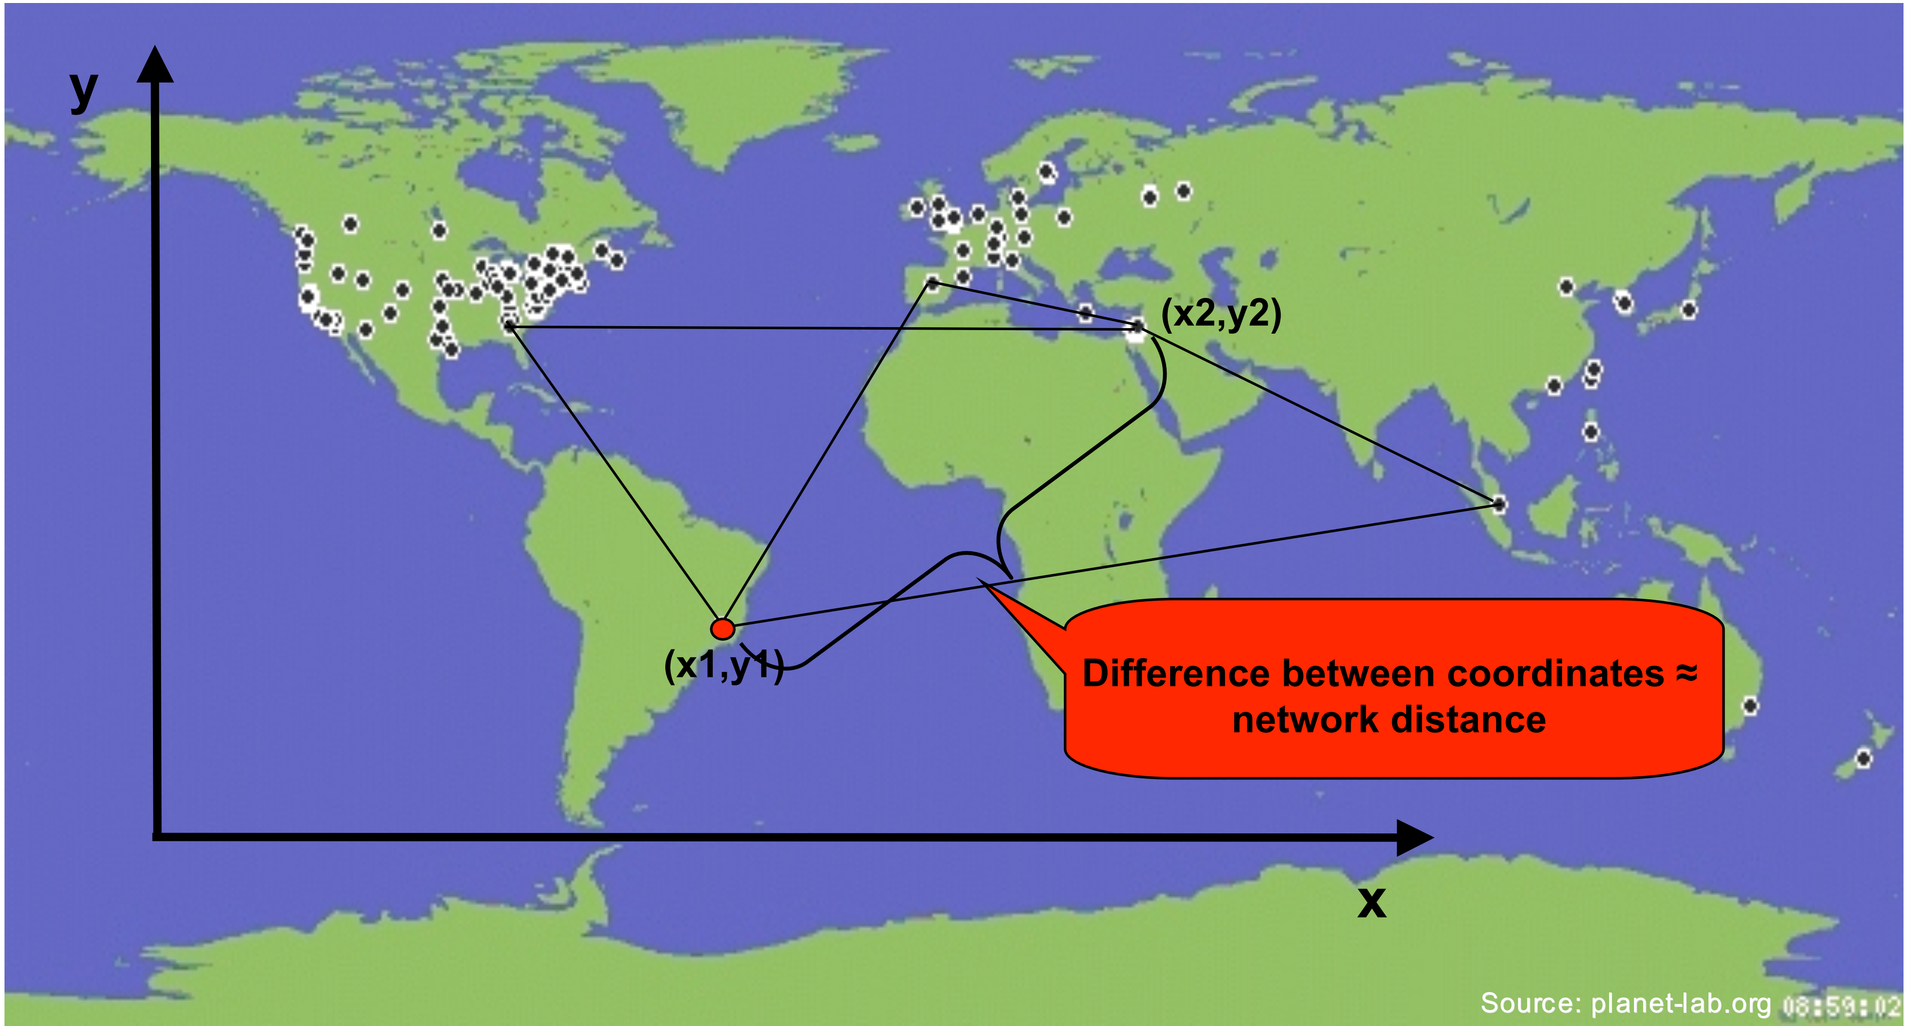

Outline

- Motivation
- **Problem statement**
- General Approach
- Techniques
  - Global Network Positioning (GNP)
  - Practical Internet Coordinates (PIC)
  - Lighthouses
  - PCA-based techniques (Virtual Landmark and ICS)
- Conclusions
- Open Issues

Problem Statement

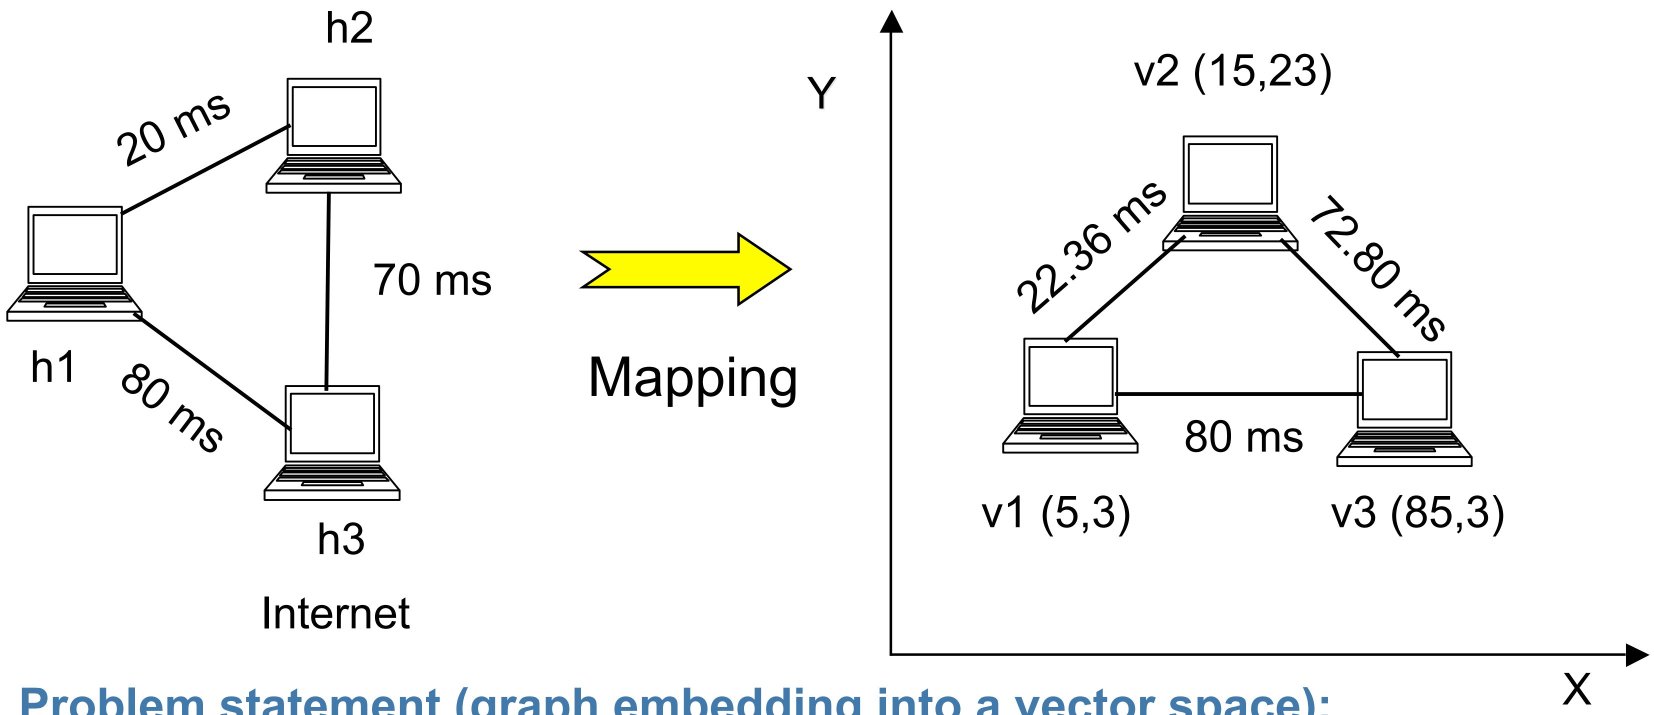

- **Problem statement (graph embedding into a vector space):**
  - *Find a scalable mapping α : H → V^k, such that d(hi,hj) ≈ D(vi,vj), where:*
  - *H is the original space (Internet graph)*
  - *V is the target vector space of dimensionality K*
  - Example: H = {h1,h2,h3}; V = {v1,v2,v3}; d(h1,h2) ≈ D(v1,v2)

Outline

- Motivation
- Problem statement
- **General Approach**
- Techniques
  - Global Network Positioning (GNP)
  - Practical Internet Coordinates (PIC)
  - Lighthouses
  - PCA-based techniques (Virtual Landmark and ICS)
- Conclusions
- Open Issues

General Approach

- Steps:
  1) Select a subset of hosts for ‘reference points’ (RP)
     - Create the origin of the coordinate system
  2) Measure *round-trip-time (distance)* between RPs
  3) Calculate coordinates for each RP
  4) Measure RTT between host and RPs
  5) Calculate coordinates for the host
- Different proposed techniques for steps 1,3 and 5
- Reference points = landmarks, lighthouses, beacons

Outline

- Motivation
- Problem statement
- General Approach
- **Techniques**
  - Global Network Positioning (GNP)
  - Practical Internet Coordinates (PIC)
  - Lighthouses
  - PCA-based techniques (Virtual Landmark and ICS)
- Conclusions
- Open Issues

Global Network Positioning (GNP)

- Pioneering work: T.S.E. Ng, H. Zhang [ACM IMW'01, INFOCOM'02]
- Landmark coordinates

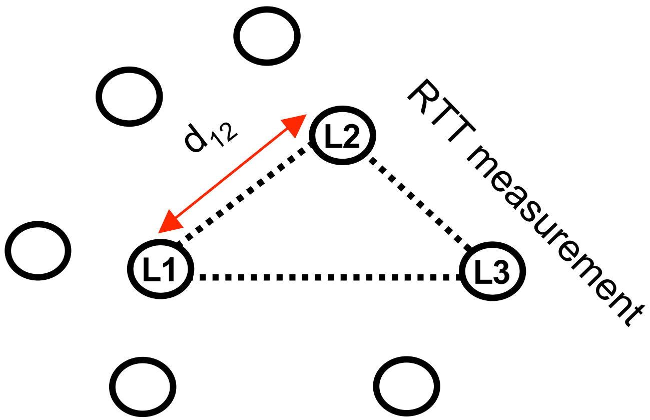

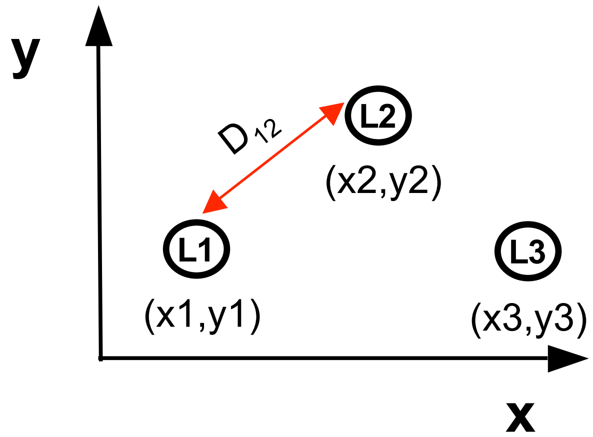

1) Landmark Selection (fixed set)
2) 'L' landmarks measure mutual network RTT/distance (ping)
3) Landmarks computes coordinates by minimizing the overall error between the measured and the estimated distances
Multi-dimensional global minimisation problem

minimise: error(d_{ij},D_{ij})

Global Network Positioning (GNP)

- Host coordinates

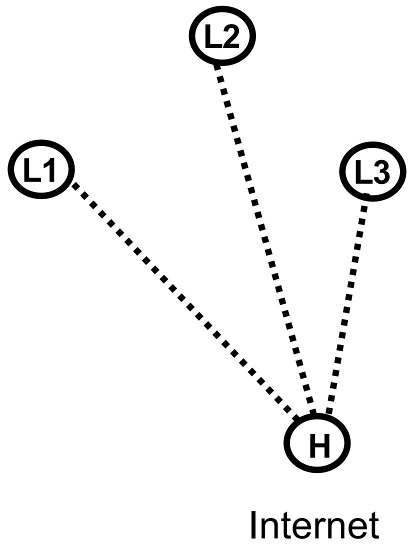

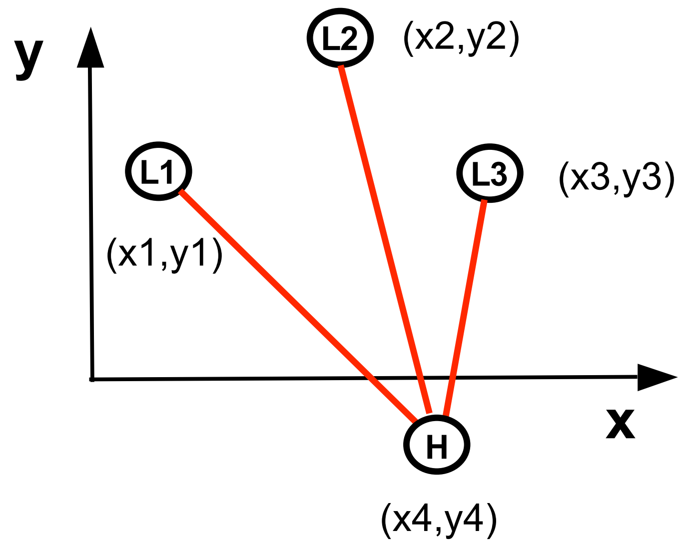

1) Host measures its network distances (RTT) to the 'L' landmarks
2) Host computes its own coordinate relative to the Landmarks
3) Multi-dimensional global minimisation problem
   minimise: error(d_{ij},D_{ij})

Global Network Positioning (GNP)

- Issues:
  - Landmark selection
    - Fixed set
    - Landmark failures and overload
    - What's the optimal selection ?
  - Technique (Simplex downhill)
    - Unique coordinates are not guaranteed
    - Depends on the starting point of the algorithm

Lighthouses

- Pias, M. et al [IPTPS'03]
- Host selects random reference points (lighthouses)
- Coordinates computed through linear transformations

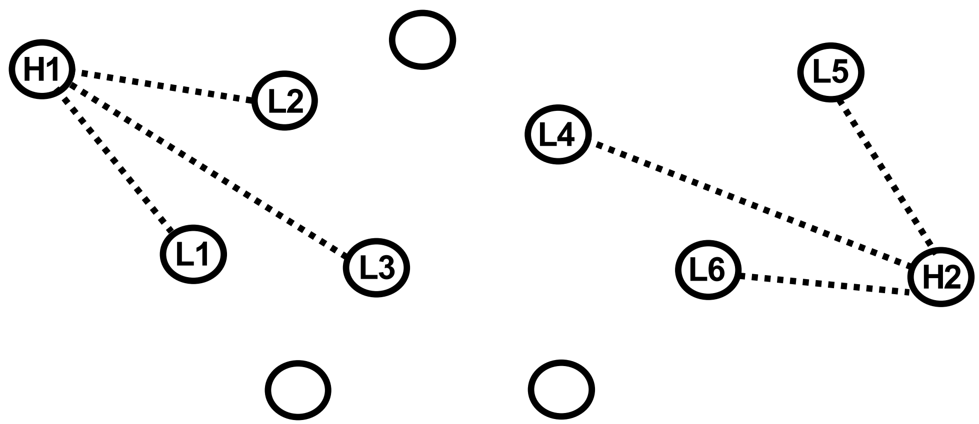

Lighthouses

- Lighthouses coordinates

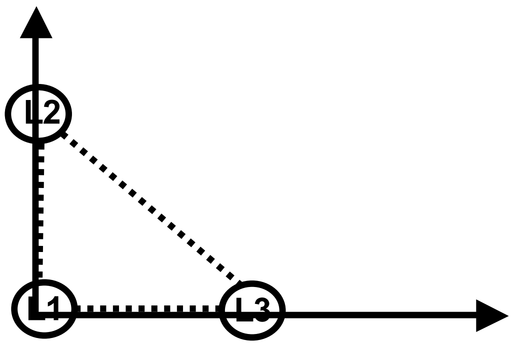

Derive a distance matrix D = $\begin{bmatrix} 0 & d_{12} & d_{13} \\ 0 & 0 & d_{23} \\ 0 & 0 & 0 \end{bmatrix}$

Computes an orthogonal basis using QR decomposition: D = Q . R
Q is the orthogonal basis that creates the coordinate system
Each lighthouse is assigned a column vector of Q

Lighthouses

- Host coordinates

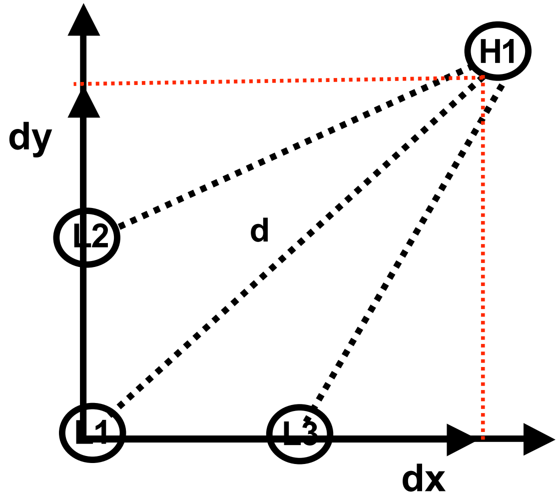

1) Host measures its network distances to the ‘L’ lighthouses
2) Distances of the host are projected onto the orthogonal basis
3) Host coordinates H = Q . B, where B is the matrix with RTTs between the host and lighthouses

Lighthouses

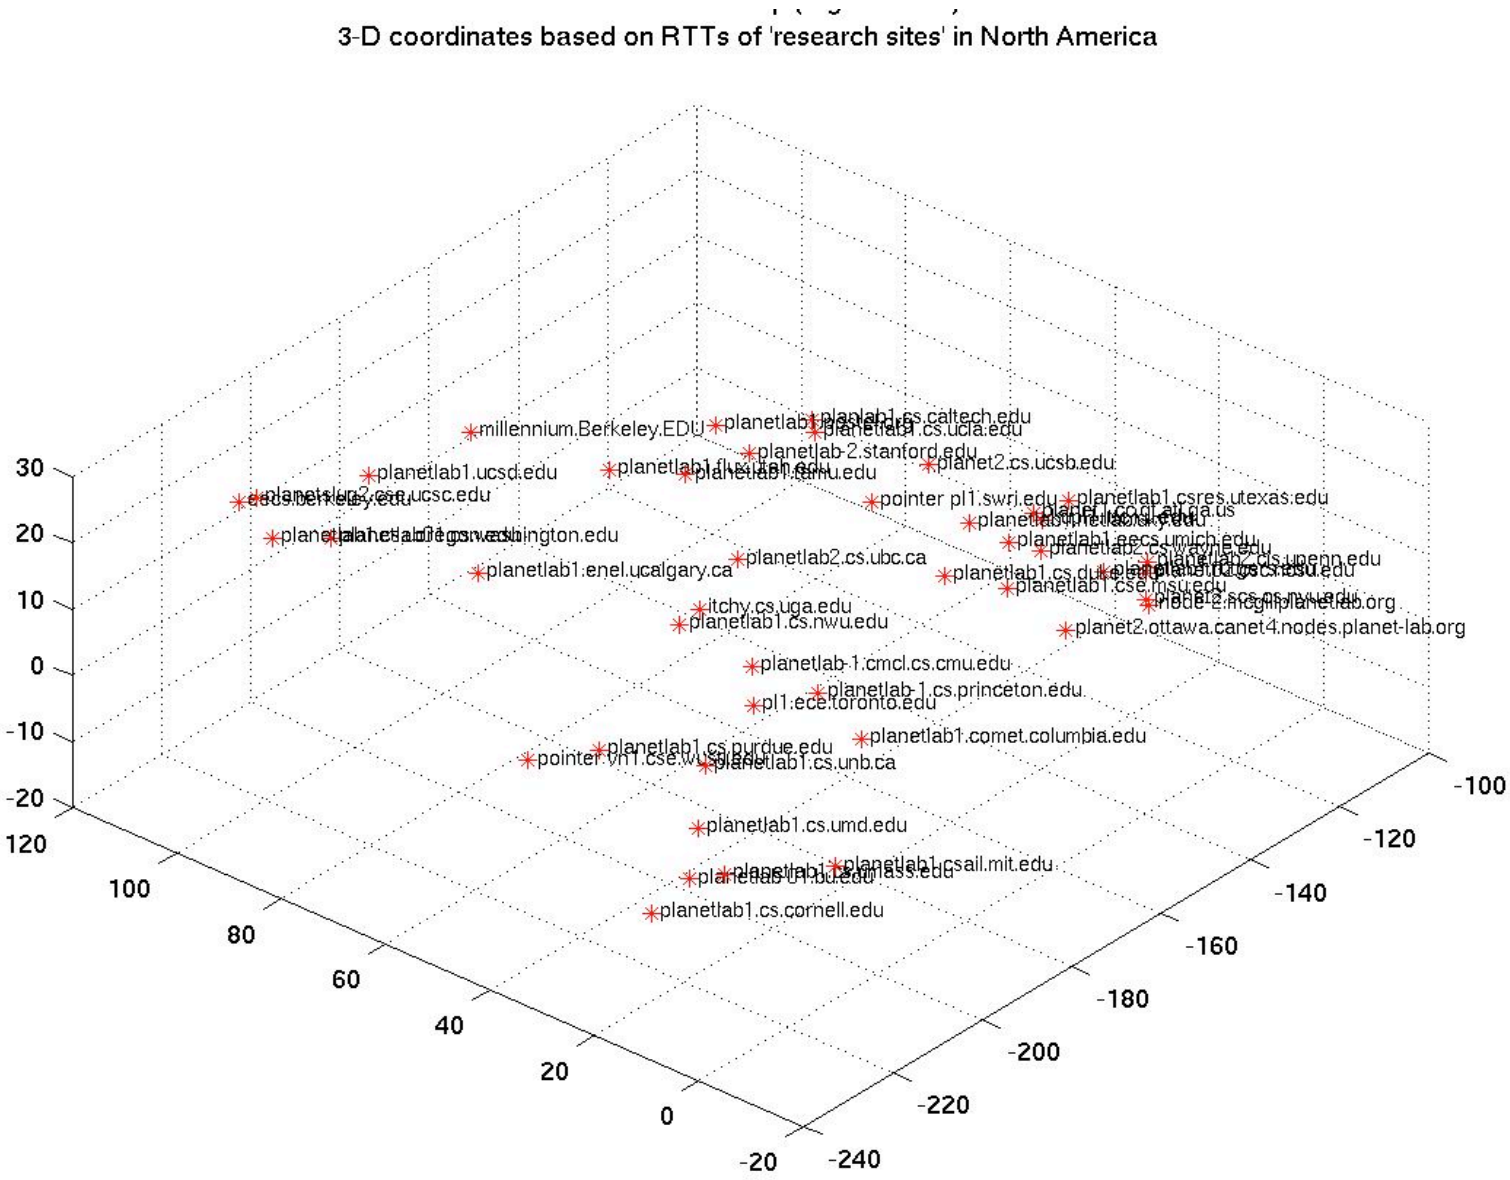

Practical Internet Coordinates (PIC)

- Costa, M. et al [IEEE ICDCS'04]
- Host coordinates

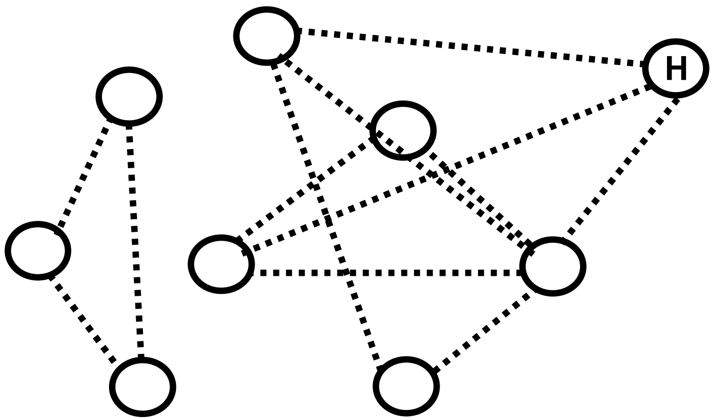

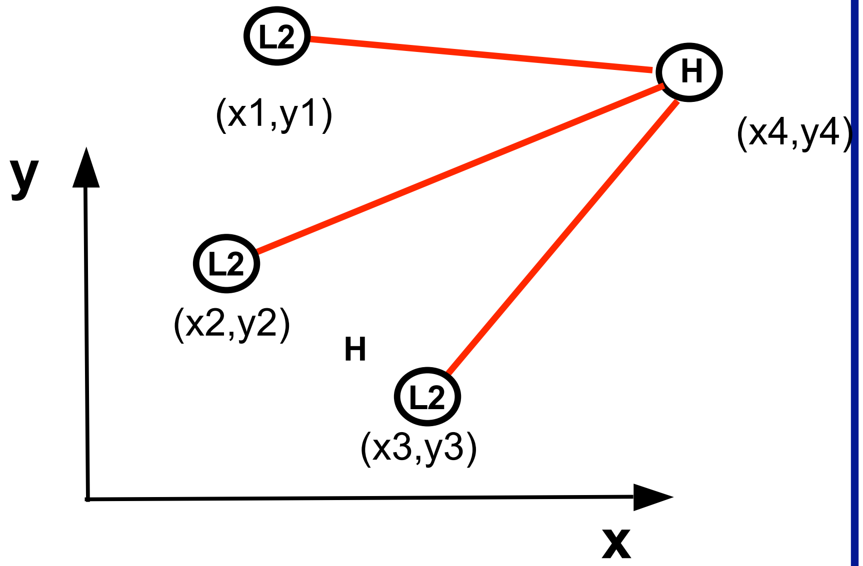

1) Any host with computed coordinates can *serve* as a landmark
2) Host measures RTTs to other 'L' landmarks
3) Host computes its own coordinate relative to the Landmarks
4) Multi-dimensional global minimisation problem

minimise: error(d_{ij},D_{ij})

Practical Internet Coordinates (PIC)

- PIC was tested in Pastry (Structured P2P system):
  - Each node maintains a routing table with distances to closest nodes
  - Without coordinates: a joining node measures 297 RTT distances in a p2p system of 20,000 nodes
  - Using coordinates: joining node measures 32 RTTs
- Selection strategy
  - Random: pick landmarks randomly
  - Closest: pick landmarks ‘closest’ to the host
  - Hybrid: pick landmarks as in random and others as in closest

PCA-based techniques (Virtual Landmarks and ICS)

- Tang, L, Crovella, M. [ACM IMC'03]: "Virtual Landmarks"
- Lim, H, Hou, J.C, Choi, C-H [ACM IMC'03]: "ICS"

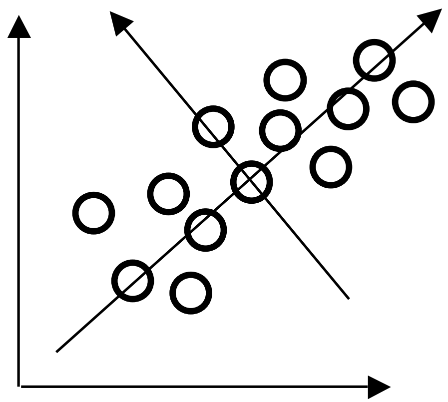

1) Larger number of landmarks/beacons (m) - high dimensionality
2) Derive a landmark distance matrix m x m
3) Use Principal Component Analysis to derive an optimal basis

PCA-based techniques (Virtual Landmarks and ICS)

- Optimal basis through Singular Value Decomposition: $D = U . W . V^T$
- Where columns of U are the principal components and form an orthogonal basis
- U has ‘m’ columns (components)
- Use the first ‘k’ principal components that allow ‘good’ projections

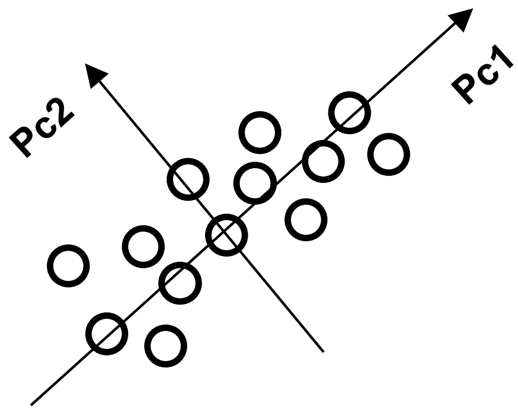

PCA-based techniques (Virtual Landmarks and ICS)

- Host Coordinates
  - Linear projections on the first 'k' principal components
  - $H_i = U^T . d_i$

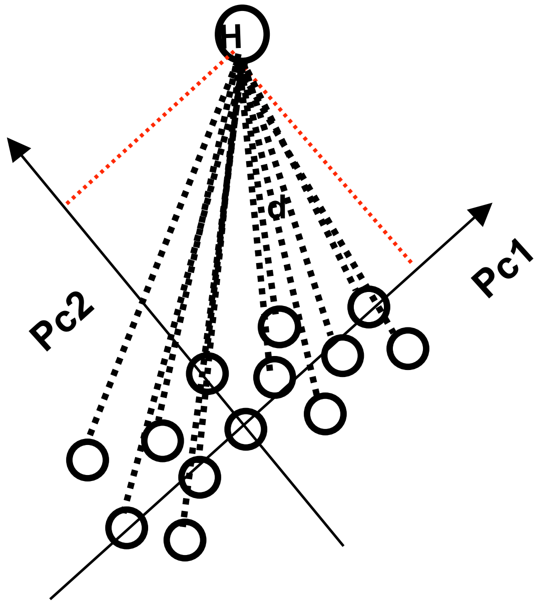

Outline

- Motivation
- Problem statement
- General Approach
- Techniques
  - Global Network Positioning (GNP)
  - Practical Internet Coordinates (PIC)
  - Lighthouses
  - PCA-based techniques (Virtual Landmark and ICS)
- **Conclusions**
- Open Issues

Conclusions

- Techniques explored:
  - Minimisation of error functions: GNP and PIC
  - Linear matrix transformation: Lighthouses, Virtual Landmarks and ICS
- Applications
  - Closest server selection (e.g. distributed network games)
  - Network-aware construction of peer-to-peer systems
  - Routing in mobile ad-hoc networks
  - Network distance estimation
- Internet Coordinate System is promising but …

Open Issues/Future work

- Landmark placement
- How many dimensions do we need to create an "Internet Coordinate System" ?
    - Some of the research suggested 6-9 dimensions
    - However, different datasets give different values
- Routing policies x dimensionality x error
- Future work
    - Visualisation tools (network topology/dynamics)
    - Refine the Usable Coordinate System (UCS) on PlanetLab

BACKUP Slides

Motivation

What's the closest server to a client in Brazil?

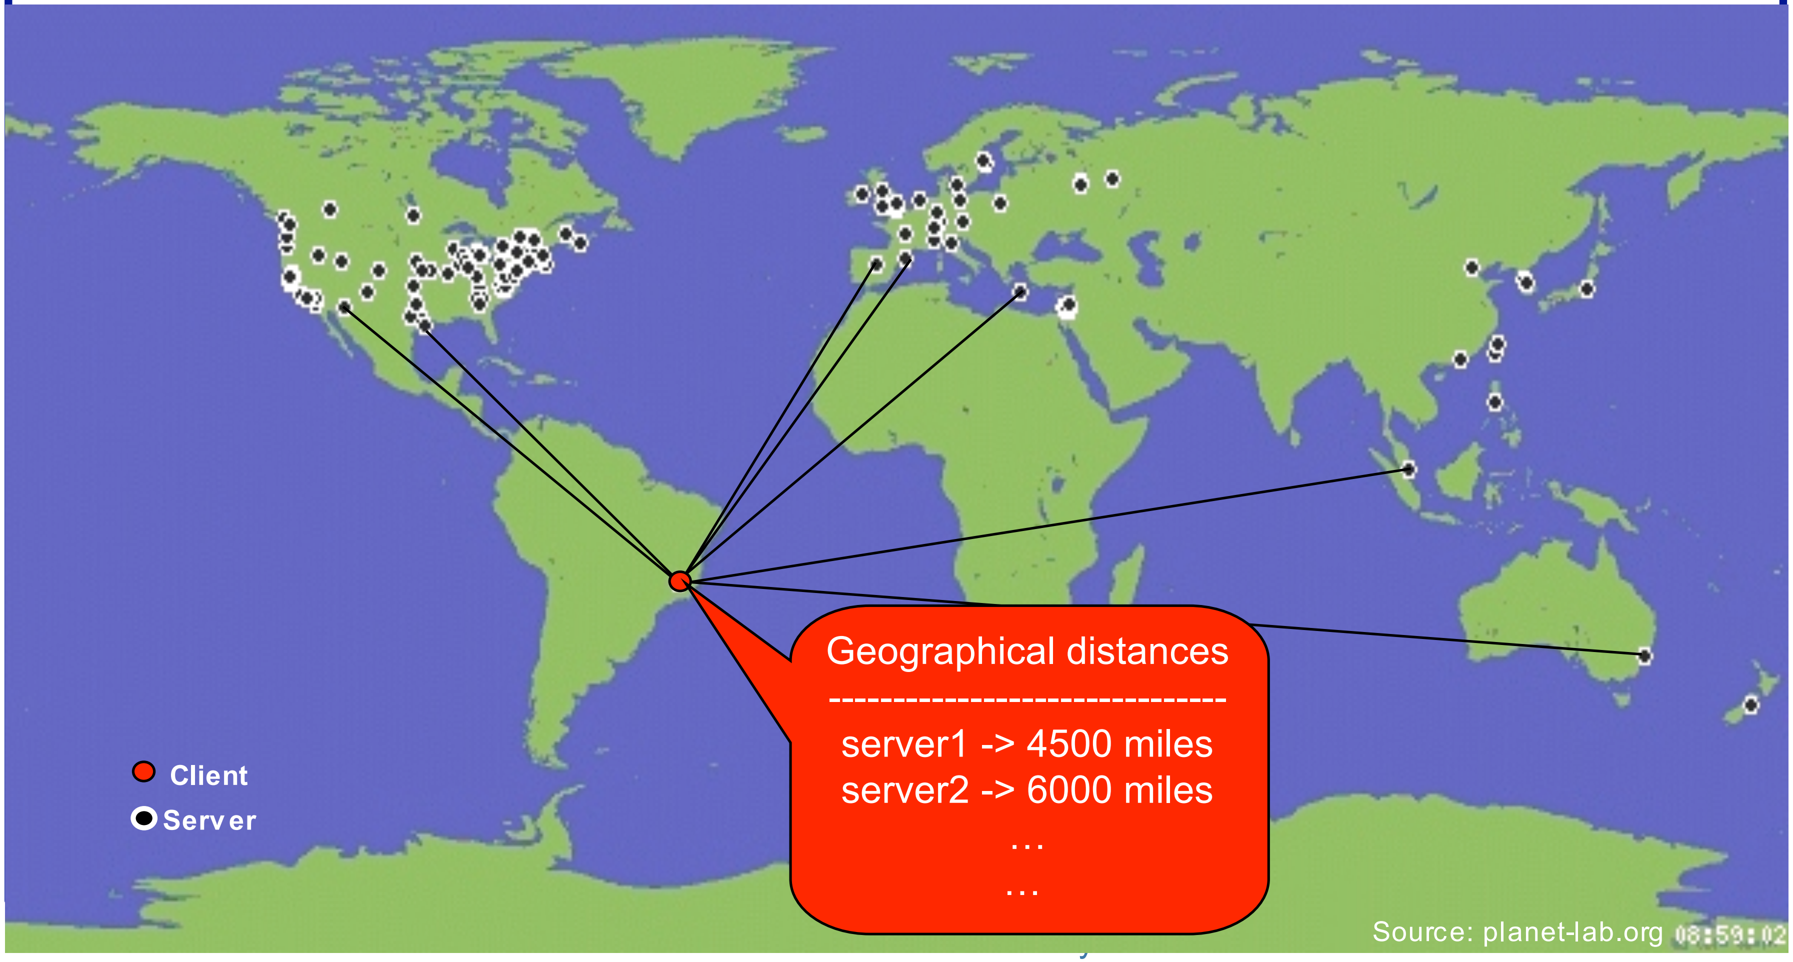

Motivation

- Difficulties:
  - Geographical distance ≠ network distance
    - Routing policies/connectivity
    - GPS is not wide available
  - Client needs 'N' distances to select the closest server

Practical Internet Coordinates (PIC)

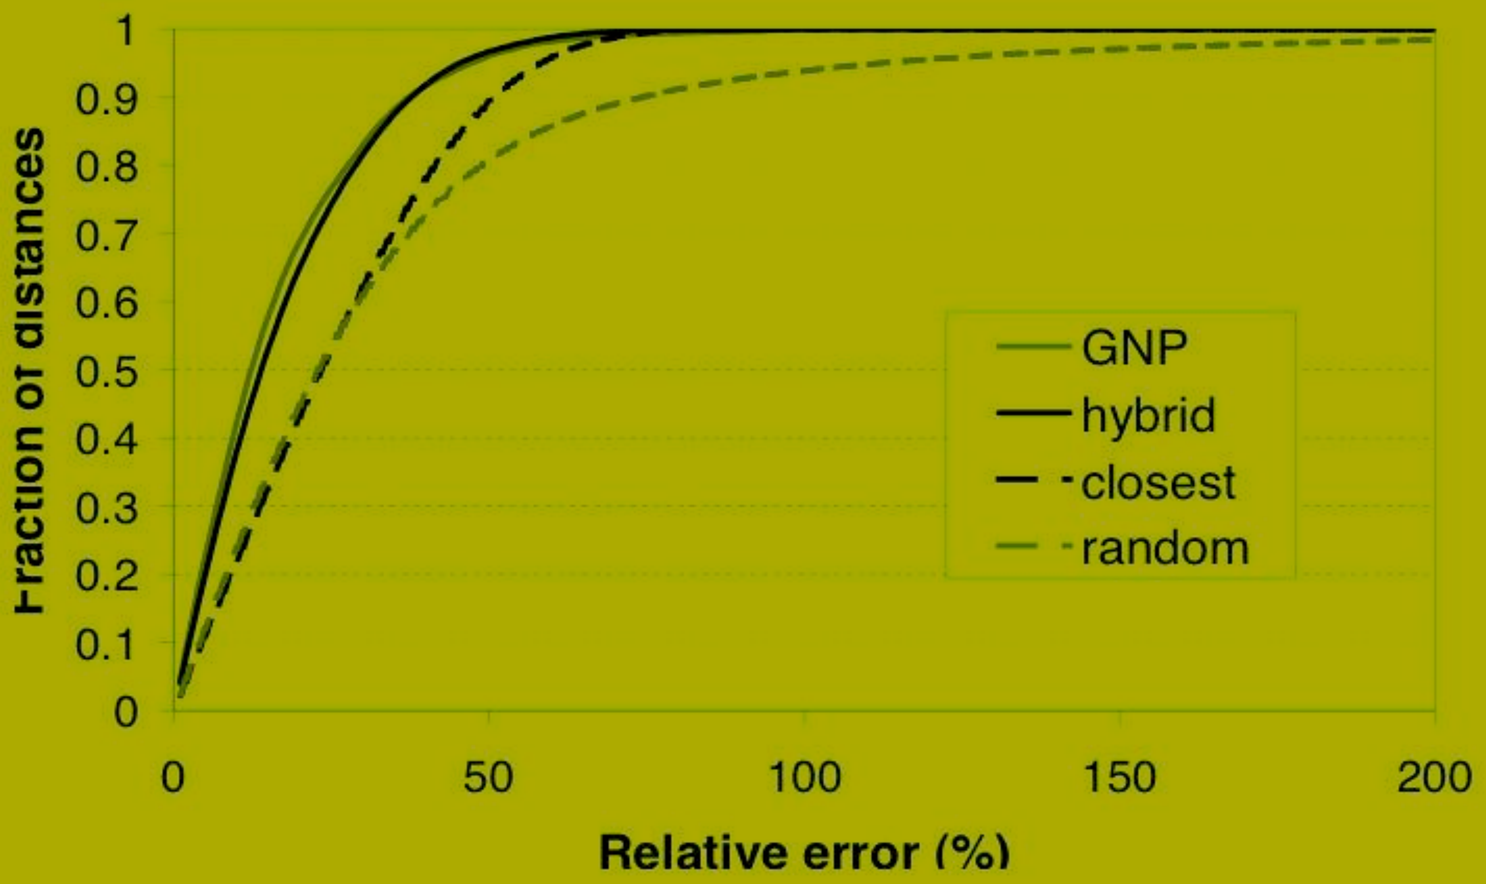

Thanks to Manuel Costa (MSR)

Distortion vs. Dimensionality

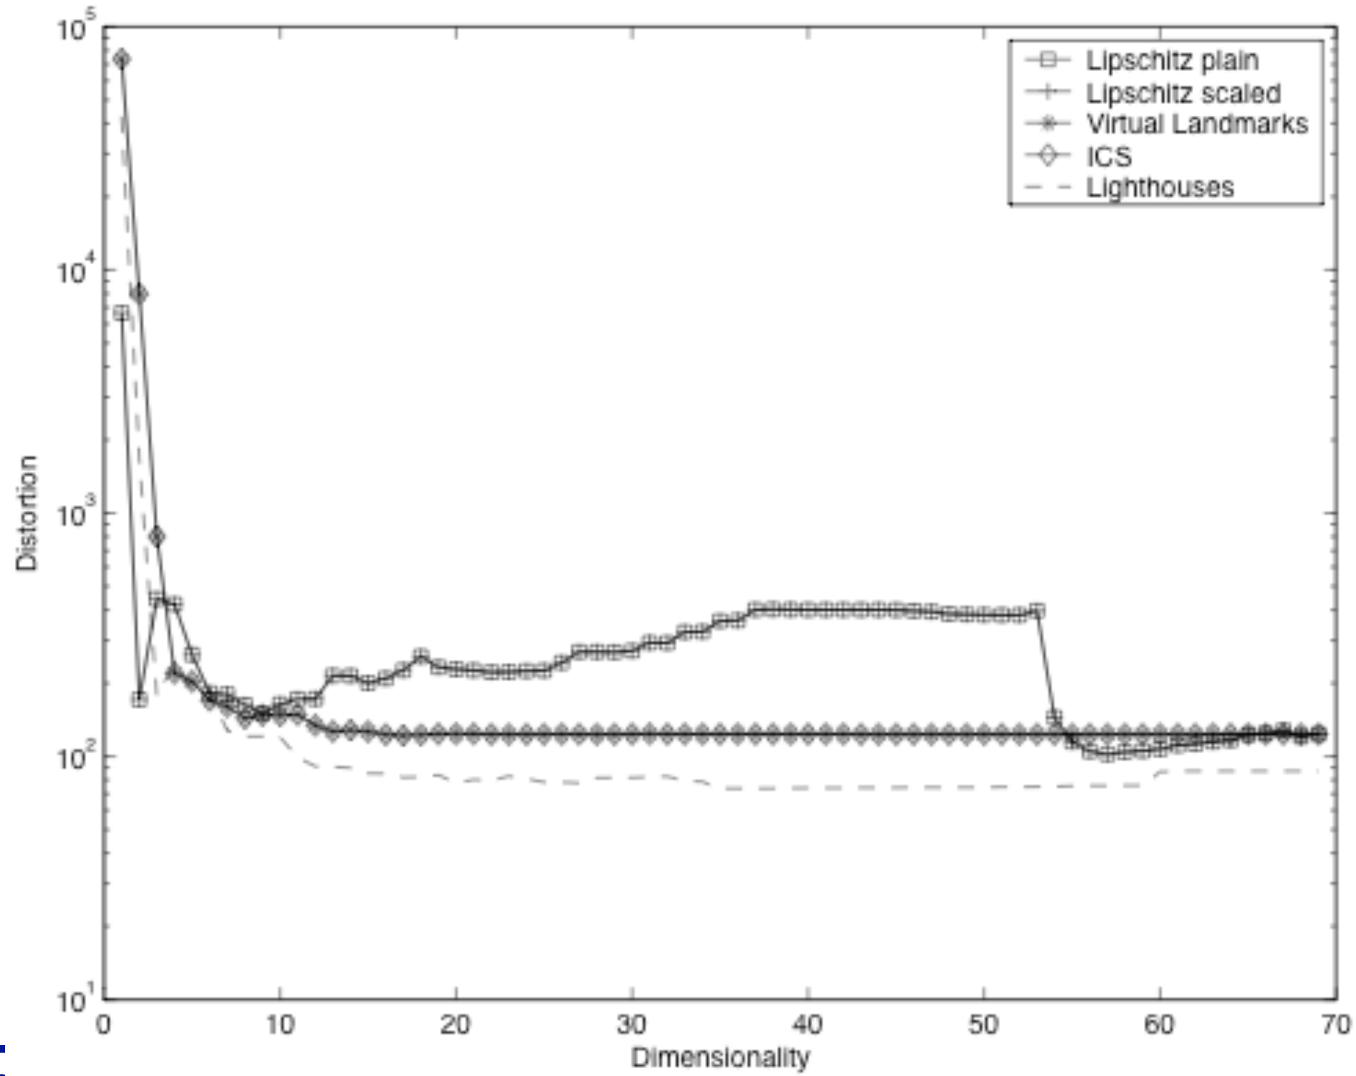

Distortion vs. Dimensionality

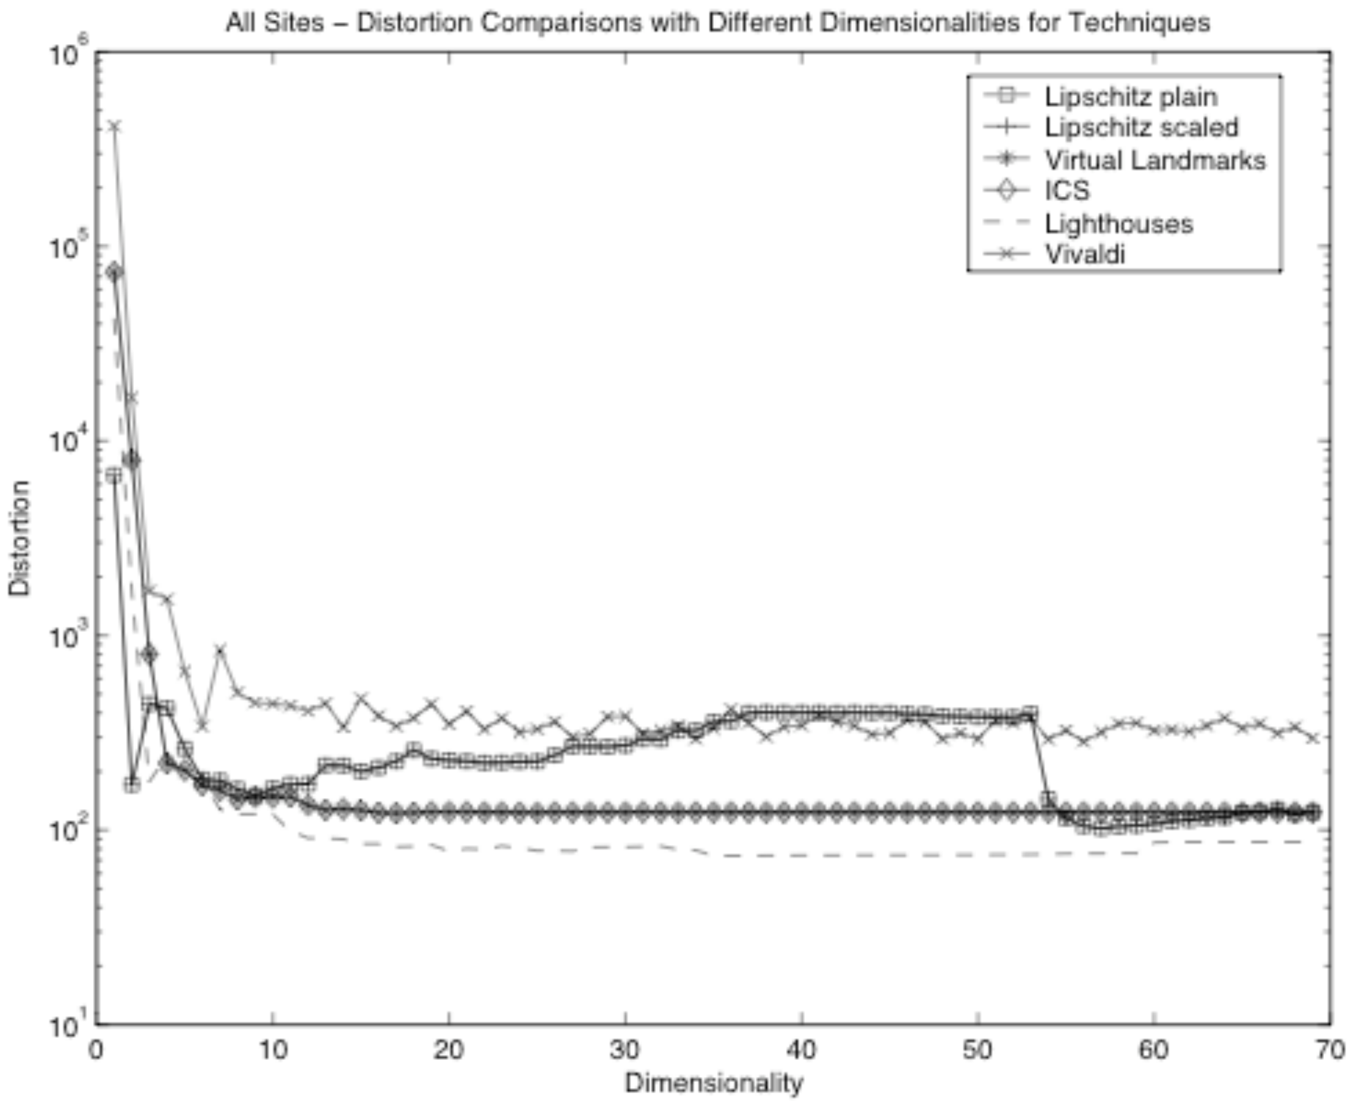

Triangle Inequality