On the Placement of Internet Instrumentation

Sugih Jamin
EECS Department
University of Michigan
jamin@eecs.umich.edu

Joint work with:

C. Jin (UM), Y. Jin (UCLA),
Y. Shavitt (Bell Labs), D. Raz (Bell Labs), and L. Zhang (UCLA)

http://idmaps.eecs.umich.edu/

Take Home Slide

Nearest mirror selection problem: Given n mirrors of a web server, which mirror is the closest to a client?

Knowing Internet distances makes closest mirror selection more accurate than random selection.

The placement of distance measurement boxes on the Internet effects the accuracy of distance-based nearest mirror selection.

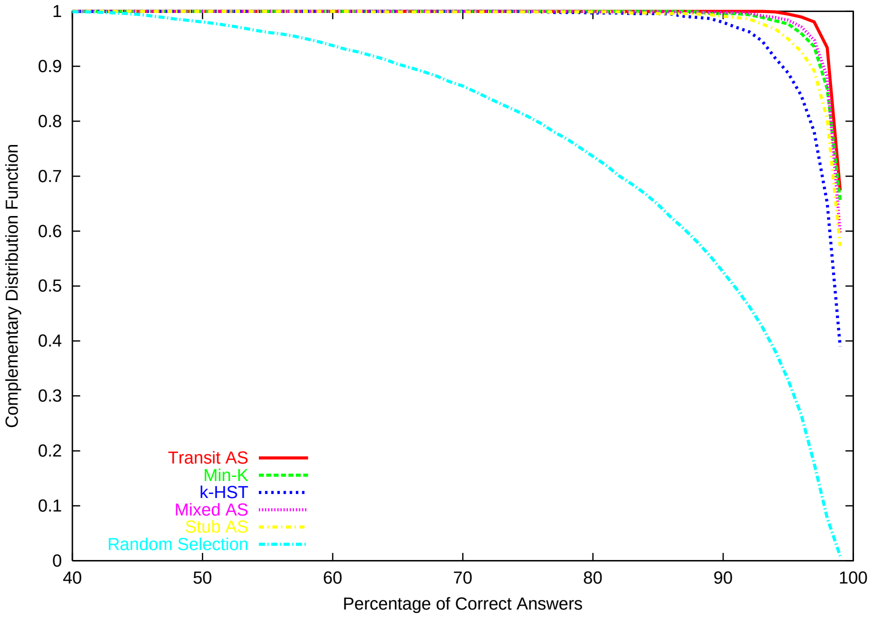

Outline

- IDMaps and distance measurement
- IDMaps performance metric
- Graph theoretic approaches to the placement problem
- Placement heuristics
- Performance metric computation
- Topology generation
- Simulation results

What is IDMaps?

Internet-wide *infrastructure* to collect distance information.

Distance metrics: hop count, round-trip time, minimum bandwidth, etc.

Uses of distance information:

- Nearest mirror selection
- Construction of host-level multicast trees
- Anycasting
- etc.

Sugih Jamin (jamin@eecs.umich.edu)

Goals and Non-goals of IDMaps

IDMaps provides:

- long-term approximate distances,
- distance estimation between any 2 points on the Internet.

IDMaps does *not* provide:

- end-to-end application-level performance
- available bandwidth or current delay
- characteristics of any specific path

IDMaps Component

- Tracers: autonomous instrumentation boxes.
- APs (Address Prefixes): regions of the Internet. All hosts within an AP are considered equidistance from the rest of the Internet.

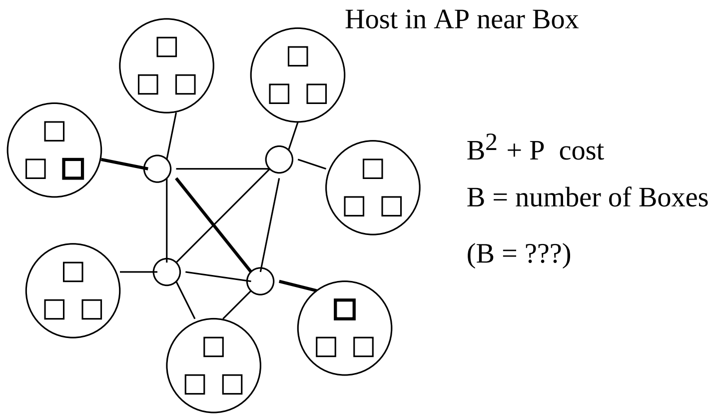

Sugih Jamin (jamin@eecs.umich.edu)

Distance Measurement

- Tracers measure distances between themselves
- Tracers measure distances to APs
- Each distance measurement is called a *virtual link* (VL).
- A distance map is a graph consisting of Tracers, APs, and VLs.

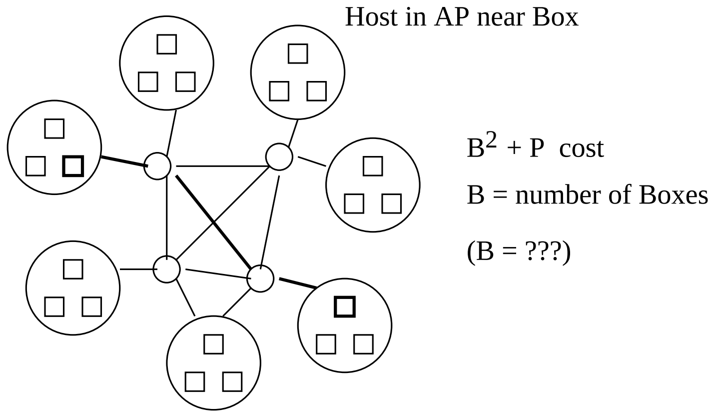

Sugih Jamin (jamin@eecs.umich.edu)

Distance Estimation

- Tracers advertise distance information
- IDMaps clients collect advertised distances and build a distance map
- Approximate distance between two points is the shortest path between them on the distance map

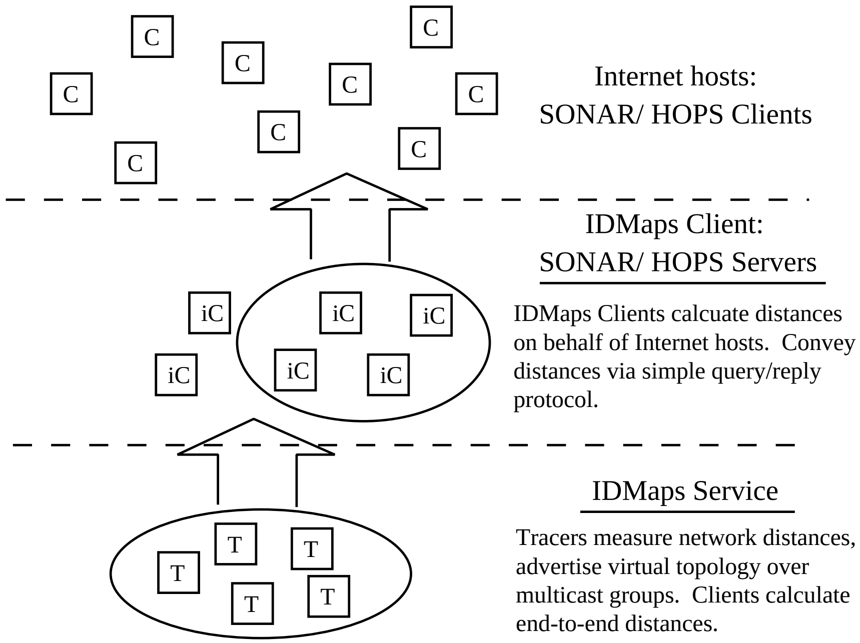

Some Research Issues

1. How many Tracers are necessary?
2. Where should Tracers be placed?
3. How many VLs are necessary?

IDMaps Performance Metric

Trade-off: scalability vs. accuracy

The number of Tracers and VLs necessary depends on the accuracy demanded of IDMaps.

Performance metric for nearest mirror selection:

- Clients are correctly directed to the nearest mirror.
- Accuracy of distance measurement need not be precise.

Graph Theoretic Approaches to Placement

Performance goal: minimize maximum distance from a node to its nearest Tracer (Note: not the same as the nearest mirror selection goal).

Basic idea: given a graph and a target maximum distance ($\mathcal{D}$), partition graph into sub-graphs of diameter $\leq \mathcal{D}$.

- k-hierarchically well-separated trees (k-HST) [1]:
  - o Top down: nodes further apart are partitioned first.
  - o Recursively partition the graph into sub-graphs with smaller diameters such that each sub-graph has a diameter $\leq \mathcal{D}$.
- minimum K-center [2]:
  - o Bottom up: nodes closer together are incorporated first.
  - o Form sub-graphs by incorporating lower cost edges up to edges with cost at most $\mathcal{D}/2$.

Both assumes known network topology.
Useful as yard sticks to compare placement heuristics.

Placement Heuristics

Internet topology unknown and constantly changing.

Placement heuristics:

- Stub-AS: initial deployment on end-hosts
- Transit-AS: Tracers deployed on/near backbone routers
- Mixed: both of the above

Number of Tracers determined by budgetary and administrative constraints.

Simulate the placement algorithms and compare the performance of IDMaps with varying number of Tracers.

Performance Metric Computation

1. Place 3 mirrors on a network.
2. For each node on the network, determine the mirror closest to it based on known topology and based on distance map.
3. Compute % of correct answer.
4. Relocate the 3 mirrors and repeat steps 2 and 3.
5. Repeat steps 1 to 4 1000 times and compute $F^{-1}(x), x$: percent correct answer.

An answer is “correct” if the distance to selected mirror
is not more than λ (= 2) times the distance to actual nearest mirror.

Want: Prob[% incorrect answer large] $\leq \epsilon$

Instead: show the relative performance of various placement algorithms.

Topology Generation

Consider 3 models: Waxman [4], Tiers [5], ASconn.

Place N nodes uniformly on an sxs plane.

Differences:

- Node's degree of connectivity: how many neighbors per node?
- Which nodes should be neighbors?

Observations from NLANR data:

- Node degree distribution on the Internet follows power law [6].
- The top most connected ASs belong to ISPs with topologically dispersed networks.
- The top most connected ASs form a full mesh.
- The top most connected ASs do not connect directly with stub-ASs.

ASconn Topology Generation

ASconn consists of 2 phases:

Phase 1:

- Place $N - \eta\kappa$ nodes uniformly on an sxs plane,
- assign node degree based on power law,
- connect the τ (= 5) highest degree nodes into a full mesh, and
- connect the rest of the nodes in order of degree of connectivity.

Phase 2:

- Take the κ (= 10) highest degree nodes and replace each with η (= 20) nodes,
- for each η nodes, place them uniformly on the plane,
- divide the connectivity of the original node among them, and
- connect them to each other using the Waxman model.

Topology Characteristics: Node Degree

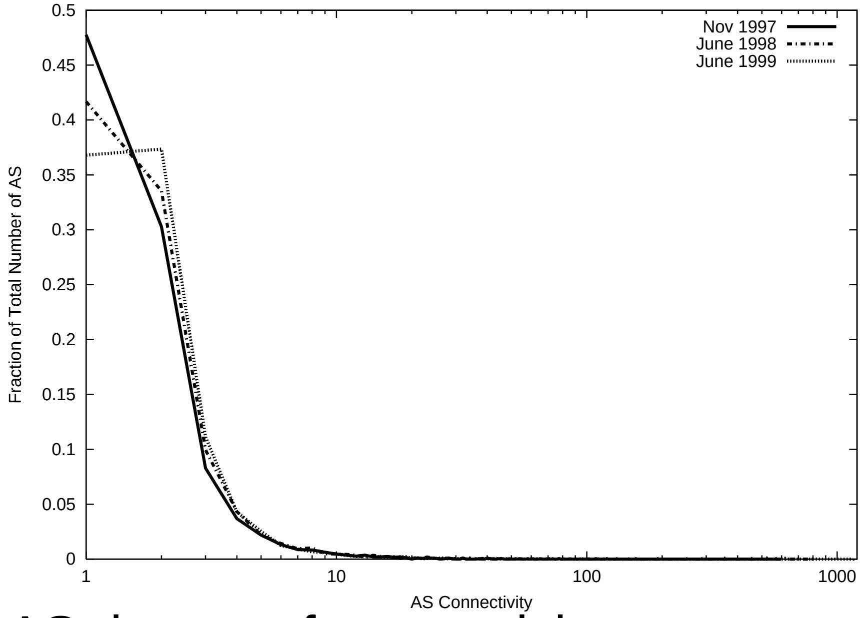

AS degree of connectivity on the Internet.

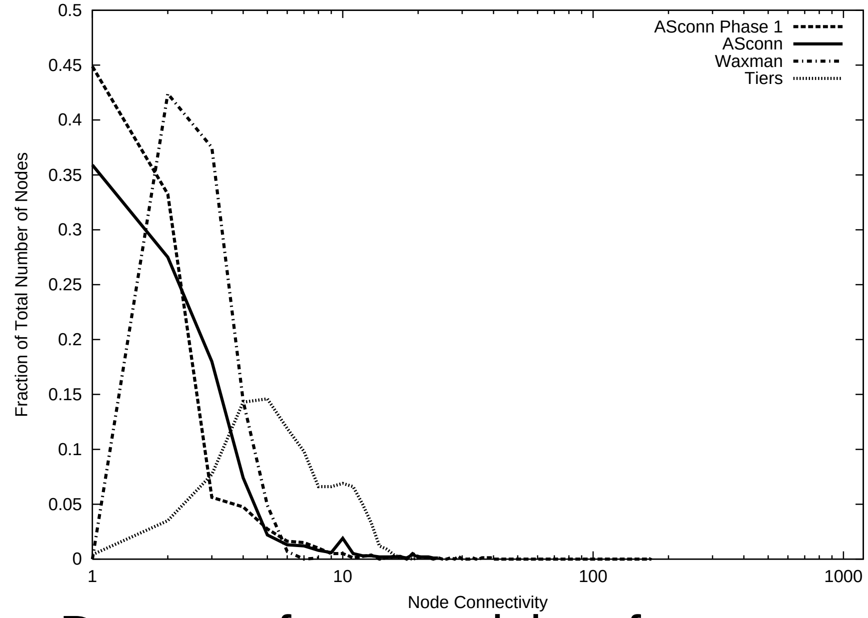

Degree of connectivity of generated networks.

Simulation Results

- Knowing distance map improves accuracy of nearest mirror selection.
- Placement algorithm effects accuracy of distance map

Accuracy of nearest mirror selection on 1,000-node ASconn network:

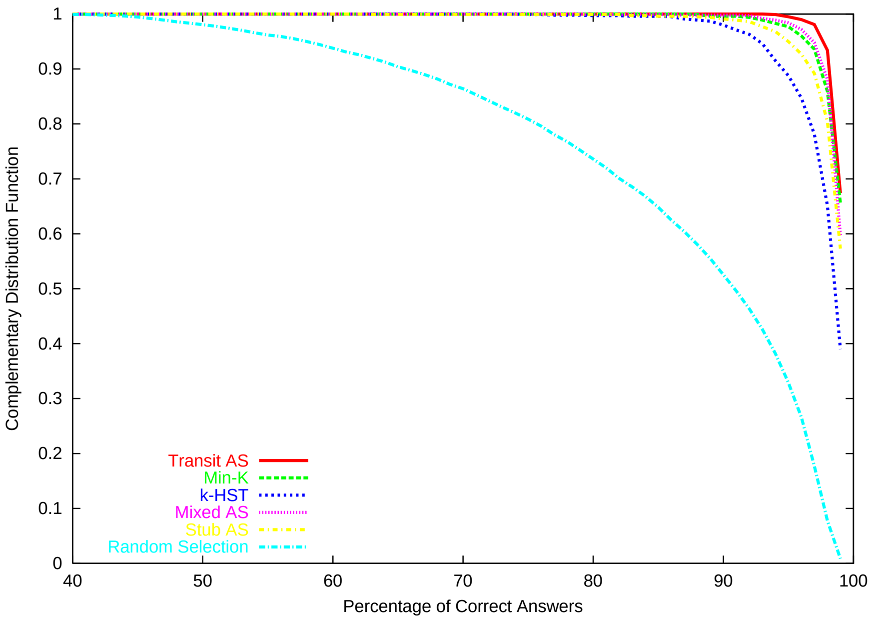

Sugih Jamin (jamin@eecs.umich.edu)

Effect of Network Topology

Accuracy of nearest mirror selection on 1,000-node Waxman network:

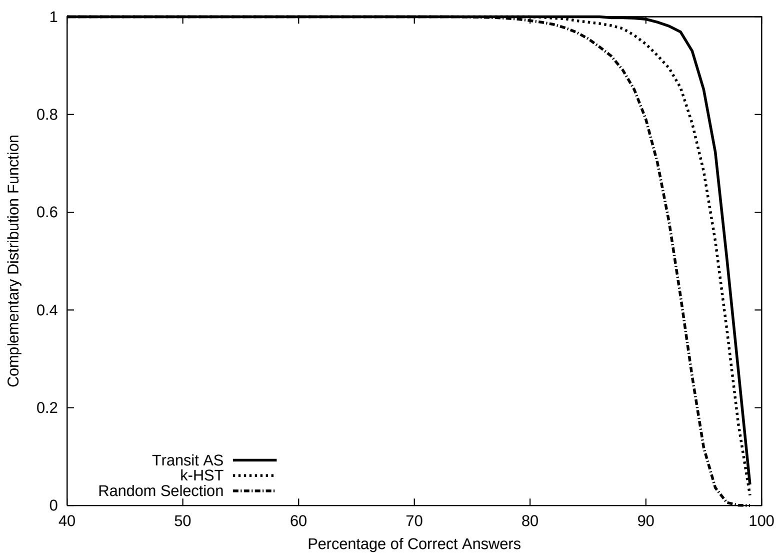

Accuracy of nearest mirror selection on 1,000-node Tiers network:

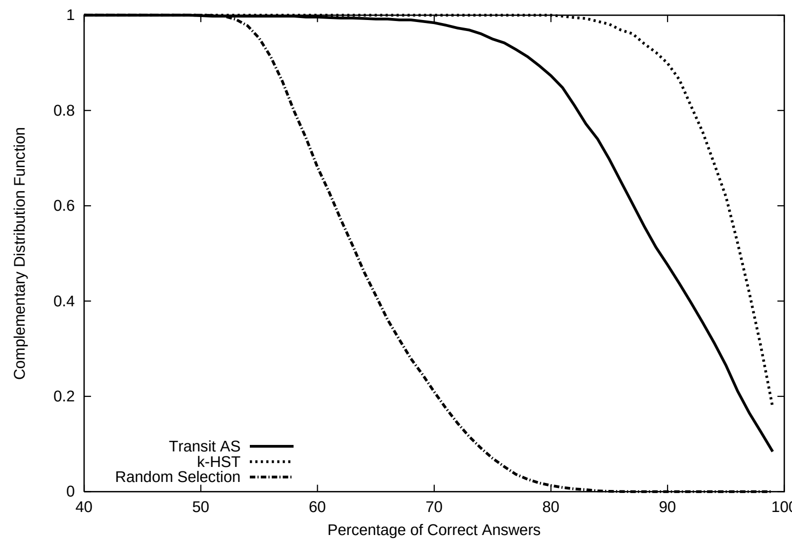

Sugih Jamin (jamin@eecs.umich.edu)

Effect of Having More Tracers

Accuracy of nearest mirror selection on 1,000-node Tiers network with varying number of Tracers:

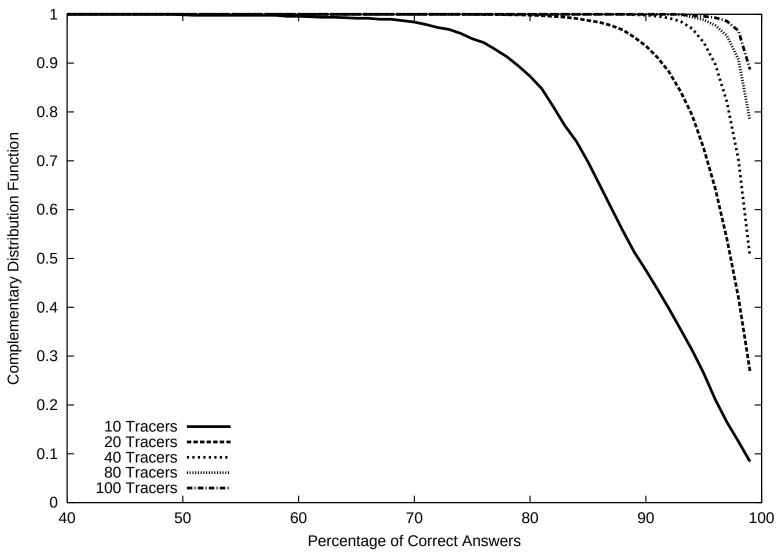

Accuracy of nearest mirror selection on 4,200-node ASconn network with varying number of Tracers:

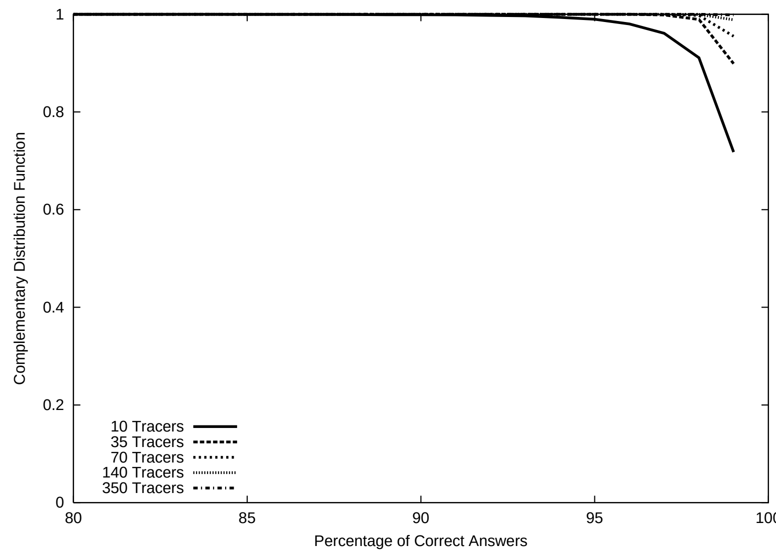

Sugih Jamin (jamin@eecs.umich.edu)

References

1 Yair Bartal, “Probabilistic Approximation of Metric Space and Its Algorithmic Applications,” in *FOCS’96*, Oct. 1996.

2 Vijay Vazirani, *Approximation Methods*, Springer-Verlag, 1999.

3 I. Althöfer et al., “On Sparse Spanners of Weighted Graphs,” *Discrete and Computational Geometry*, vol. 9, pp. 81–100, 1993.

4 Bernard M. Waxman, “Rougint of multipoint connections,” *IEEE JSAC*, 6(9):1617–1622, Dec. 1998.

5 K. Calvert, M.B. Doar, E.W. Zegura, “Modeling Internet Topology,” *IEEE Communications Magazine*, June 1997.

6 C. Faloutsos, M. Faloutsos, P. Faloutsos, “On Power-Law Relationships of the Internet Topology,” *Proc. of ACM SIGCOMM*, Aug. 1999.