



Triangulation and Embedding using Small Sets of Beacons

Alex Slivkins
Cornell University

Joint work with Jon Kleinberg and Tom Wexler
(to appear in FOCS 2004)



Motivation: Internet latencies

- Latency (round-trip time) is a notion of distance in Internet
 - distance matrix defined by latencies is almost a metric
- Two lines of research:
 - networking community studies the matrix of latencies
 - in FOCS/STOC community, lots of research on metrics
- Can we make a connection? fundamental disconnect:
 - theoretical studies: assume full access to distance matrix
 - distributed settings: only small #latencies can be measured
- We try to correct this discrepancy, in the context of ...





Estimating Internet latencies

- Many networking papers try to estimate latencies for many or most node pairs **using only a small #measurements**
 - one reliable measurement = e.g. min of 220 pings
- Our setting: large overlay network in Internet
 - P2P network, file-sharing system, distributed gaming app
- Applications:
 - find nearest neighbors (servers, file replicas, game buddies)
 - find farther neighbors (fault-tolerant storage)
 - proximity routing in P2P networks
 - two nodes can quickly estimate their latency

How to estimate latencies?

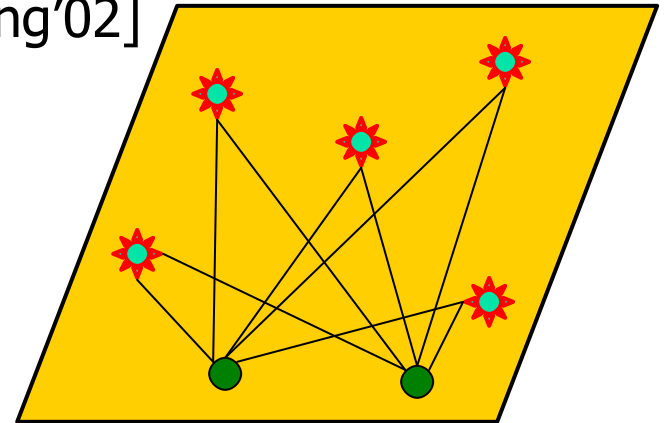
- IDMaps [Francis+ '01]: select small #nodes as “beacons”
 - users measure latency to beacons, **not** to one another
 - two nodes estimate latency via triangle inequality
- **magic**: relative error <1 on 90% node pairs
 - 900 random nodes, 15 beacons
 - $\text{relative error}(x,y) = |x-y| / \min(x,y)$
- Similar approaches: [Hotz '94], [Guyton+ '95], [Kommareddy+ '01]



Estimate latencies via embedding


Global Network Positioning (GNP) [Ng+Zhang'02]

- select small #nodes as “beacons”
 - users measure latencies to beacons
- embed into low-dim Euclidian space
 - embed the beacons
 - embed non-beacons **one by one**
- **magic**: 90% node pairs are embedded with relative error $<.5$
 - 900 random nodes, 15 beacons, 7 dimensions
- lots of follow-up work: **NPS** [Ng+ '03], **Vivaldi** [Cox+ '04], **Lighthouses** [Pias+ '03], **BigBang** [Tankel+ '03], **ICS** [Lim+ '03], **Virtual Landmarks** [Tang+ '03], **PIC** [Costa+ '04], **PALM** [Lehman+ '04]



How can we explain the magic?

- How can we explain **IDMaps** and **GNP** theoretically?
 - assume that latencies form a metric
- Classical problem: embed into Euclidean space (or into L_p , $p \geq 1$)
 - [Bourgain '85], [Linial+ '95], [Rao '99], [Gupta '02], ...
 - full access to distance matrix (not suitable for networking apps)
- In beacon-based model, we get **first provable guarantees** for:



GNP-style embeddings into Euclidean space (and L_p , $p \geq 1$)

- constant distortion for most edges



IDMaps-style distance reconstruction

- upper and lower bounds that are close for most edges



Extra: L_p space

d -dimensional L_p space = d -dimensional space of reals with norm


$$\|x - y\|_p = \sum_{i=1}^d |x_i - y_i|^p$$

- e.g. L_2 is the usual Euclidian norm

Relaxed guarantees (1/2)

Cannot estimate **all** distances with small #beacons

- partition nodes into **k** equal-size clusters, $k > \# \text{beacons}$
- distance 0 within each cluster, 1 between clusters
- many clusters do not contain any beacons
- in these clusters distance to every beacon is 1
⇒ no way to tell which node lies in which cluster





Relaxed guarantees (2/2)


- Allow “ ϵ -slack”: guarantees for all but ϵ -fraction of node pairs
 - empirical results of GNP and IDMaps have ϵ -slack, too
 - new angle for theoretical work on embeddings
- This seems similar to property testing
 - similar goal: get global info using small #probes
 - guarantees with some form of slack
- However, there are differences:
 - “does property hold?” vs “embed/ reconstruct distances”
 - “change small fraction of distances” vs “ignore them”



Doubling metrics

- In our framework, no good guarantees for arbitrary metrics
 - latencies are not arbitrary! need some restrictions ...
- Consider **s-doubling metrics**
 - Def any ball can be covered by **s** balls of half the radius
 - generalize low-dim Euclidean space: dimension $\approx O(\log s)$
 - suggested as good model for latencies, e.g. [Plaxton+ '97]
- Theoretical work: Assouad'83, [GKL'03], [KL'04], [LMN'04], ...
 - embeddings, nearest neighbor, TSP, distance labeling
 - embed into L_p , $p \leq 2$ with distortion at least $\Omega(\sqrt{\log n})$
 - recall: if $.5 < d_{\text{true}} / d_{\text{estimate}} < 2$ then **distortion** < 4

Our results for doubling metrics



GNP-style embeddings into Euclidean space (and L_p , $p \geq 1$)

- assume s -doubling metric, allow ϵ -slack
- $O(s^5/\epsilon)$ beacons \Rightarrow distortion $O(\log s/\epsilon)$
- guarantees independent of #nodes
- recall that we cannot do this **without** ϵ -slack!



IDMaps-style distance reconstruction

- estimate using triangle inequality
- $s^{O(\log 1/\delta)}/\epsilon$ beacons \Rightarrow upper and lower bounds that are within factor $1+\delta$ for all but ϵ -fraction of edges




Overview

- Introduction
- IDMaps-style distance reconstruction
- GNP-style metric embedding
- “Fully distributed” approaches: low load on **every** node
- Extensions and open questions


IDMaps-style distance reconstruction

- Rules: randomly select small #nodes as “beacons”
 - each node measures distances to all beacons
 - nodes (u,v) exchange their data and estimate $d(u,v)$
- **Triangulation**: a way to form such estimate via triangle inequality
 - (u,v) , beacon b : $|d(u,b)-d(v,b)| \leq d(u,v) \leq d(u,b)+d(v,b)$
 - lower bound: $d^-(u,v) = \max_{\text{beacons } b} |d(u,b)-d(v,b)|$
 - upper bound: $d^+(u,v) = \min_{\text{beacons } b} d(u,b)+d(v,b)$




Triangulation: proof (1/3)

- Want the bounds to be close, for most node pairs
- Def node pair (u,v) is δ -good if some beacon lies within distance $\delta d(u,v)$ from u or v
 - then $d^+(u,v) / d^-(u,v) \leq 1+3\delta$
- It suffices to show that most node pairs are δ -good



Triangulation: proof (2/3)



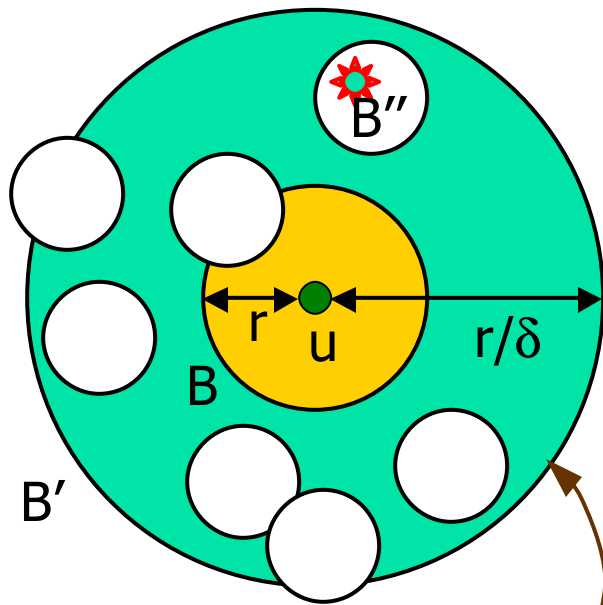
Thm Use $s^{O(\log 1/\delta)}/\epsilon$ beacons \Rightarrow w.h.p.
all but ϵ -fraction of node pairs are δ -good

- assume s -doubling metric

Fix node u , consider pairs (u,v)
 B = smallest ball around u with ϵn nodes
 r = radius of B
 B' = ball around u of radius r/δ

Since $|B| \geq \epsilon n$, B contains a beacon w.h.p.
 \Rightarrow pairs (u,v) , $v \notin B'$ are δ -good
What about pairs (u,v) , $v \in B'$?

Triangulation: proof (3/3)



B = smallest ball around u
that contains ϵn nodes

- Doubling property $\Rightarrow B'$ can be covered by $k = f(s, \epsilon, \delta)$ small balls of radius δr
 - Fix some small ball B'' , consider pairs $(u, v), v \in B''$
 - if B'' contains $> \epsilon n / k$ nodes, then w.h.p. it contains a beacon, so these pairs are δ -good
 - else **ignore** these pairs
 - only $\leq \epsilon n$ node pairs ignored for each u
- Theorem proved**

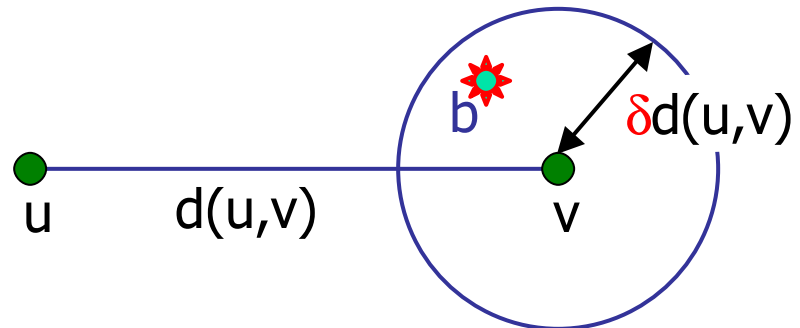


Overview

- Introduction
- IDMaps-style triangulation
- GNP-style metric embedding
- “Fully distributed” approaches: low load on **every** node
- Extensions and open questions

GNP-style metric embedding (1/3)

- Rules: embed s -doubling metric into L_p , $p \geq 1$
 - randomly select k nodes as “beacons”
 - coordinates of node = fn(its distances to beacons)
 - allow ϵ -slack; want k , dimension and distortion to be fn(s, ϵ)
- Pick k so that all but ϵ -fraction of node pairs are $\Theta(1/\log k)$ -good
 - recall: node pair (u, v) is δ -good if some beacon lies within distance $\delta d(u, v)$ from u or v
 - here $\delta = \Theta(1/\log k)$




GNP-style metric embedding (2/3)

Claim Suppose distortion on distances to beacons is $D=O(\log k)$.
Then distortion on every δ -good node pair is $O(D)$.

Proof sketch:

- set $\delta=1/3D$; it follows that $d_{\text{embed}}(v,b) \leq d_{\text{embed}}(u,b) / 2$
- $d(u,v) \approx d(u,b)$ both for true and embedded distances
- distortion on $uv = O(\text{distortion on } ub) = O(D)$

claim proved





GNP-style metric embedding (3/3)

- Thm distortion $O(\log k)$ on all but ϵ -fraction of node pairs
- It suffices to provide an embedding such that distortion on distances to beacons is at most $O(\log k)$
- Our algorithm:
 - let S be the set of all beacons
 - for each i select $\Theta(k)$ random subsets S_{ij} of S of size 2^i
 - ij -th coordinate of node u = distance from u to S_{ij}
- This is similar to the algorithm of [Linial+ '95], except
 - they use only $\Theta(\log k)$ sets S_{ij} for each i
 - their proof bounds distortion **only** between the beacons, whereas we bound distortion for **all** node-beacon pairs



Overview

- Introduction
- IDMaps-style triangulation
- GNP-style metric embedding
- “Fully distributed” approaches: low load on **every** node
- Extensions and open questions



Fully distributed approaches (1/4)

- Informal def fully distributed algorithm: low load on **all** nodes
- Recent trend in networking: fully distributed analogs of **GNP**
 - Vivaldi [Cox+ '04], Lighthouses [Pias+ '03], PIC [Costa+ '04], PALM [Lehman+ '04]
 - embed Internet latencies into low-dim Euclidean space
 - use distributed heuristics, with results similar to **GNP**
- Their approach:
 - each node selects **k** random nodes as **neighbors**
 - only distances to neighbors are measured
 - my set of neighbours is different from yours!



Fully distributed approaches (2/4)


- Basic question: **Vivaldi** only looks at distances to neighbors. Suppose these distances are embedded with small distortion. Any guarantees on the rest of the node pairs?
 - each node selects k neighbors uniformly at random
 - assume doubling metric, allow ϵ -slack, want small k
- **Thm** If distances to neighbors are embedded with distortion D then for all but ϵ -fraction of node pairs distortion is $O(D)$
 - essentially $k = (D \log n)^{\text{constant}}$

Fully distributed approaches (3/4)

Overview of the proof:

- assume distortion D on distances to neighbors
- all but ϵ -fraction of node pairs are δ -good, $\delta = 1/\Omega(D \log n)$
- def pair (u,v) is δ -good if ball B of radius $\delta d(u,v)$ around u or v has enough nodes so that every node has ≥ 3 neighbours in B
- distortion $O(D)$ on each δ -good node pair

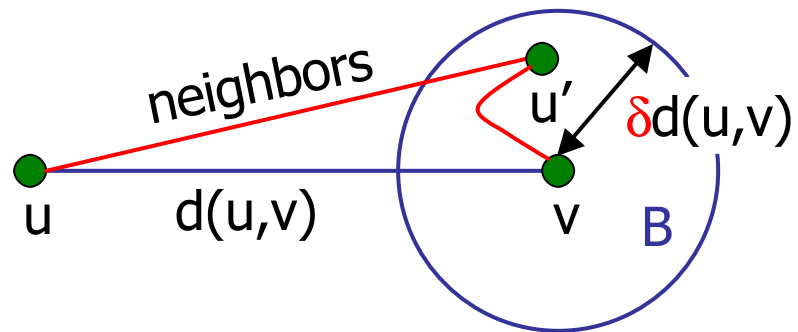
theorem proved



Fully distributed approaches (4/4)

Claim distortion $O(D)$ on each δ -good node pair (proof sketch)

- recall: every node has ≥ 3 random neighbours in B
- set of neighbour pairs induces an expander on B
 \Rightarrow this expander contains $u'v$ -path P with $O(\log n)$ hops
- in the embedding, each hop is distorted by at most D , so
 $\text{length}(P) \leq D \times O(\log n) \times \text{diam}(B)$, much shorter than edge uu'
- $d(u,v) \approx d(u,u')$
true and embedded
- distortion on uv
= $O(\text{distortion on } uu')$
= $O(D)$, **claim proved**





Overview

- Introduction
- IDMaps-style triangulation
- GNP-style metric embedding
- “Fully distributed” approaches
 - general justification
 - triangulation and embedding with provable guarantees
- Extensions and open questions



Fully distributed triangulation (1/4)


Overview of the algorithm:

- each node selects k neighbors, measures distances to them
 - my set of neighbors is different from yours!
- each node elects itself as a beacon with probability $\Theta(k/n)$
 - $\Theta(k)$ beacons total in the network
- if elected, beacon announces itself to network via gossiping
 - gossip to neighbors only; protocol is crucial (details coming ...)
- nodes receive announcements and estimate distances to beacons **without measuring them directly** (details coming ...)
- two nodes (u,v) exchange their estimates and estimate $d(u,v)$

Fully distributed triangulation (2/4)


How do nodes estimate distances to beacons?

- node u deduces that $d(u, \text{beacon})$ lies between $x-y$ and $x+y$
- node v deduces that $d(v, \text{beacon})$ lies between 0 and z
- nodes (u,v) exchange their bounds and deduce that $d(u,v)$ lies between $x-(y+z)$ and $x+(y+z)$
 - upper and lower bounds on $d(u,v)$, like in beacon-based triangulation



Fully distributed triangulation (3/4)

- Thm If $k = (D \log n)^{\text{constant}}$ then for all but ε -fraction of node pairs there will be a beacon such that $(y+z) < x/2D$
- then upper and lower bounds on $d(u,v)$ are within factor $1+1/D$

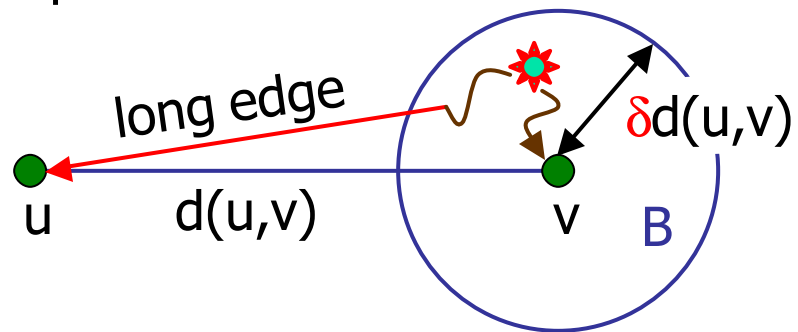


Fully distributed triangulation (4/4)

Proof sketch:

- all but ε -fraction of node pairs are δ -good, $\delta = 1/\Omega(D \log n)$
 - pair (u,v) is δ -good if ball B of radius $\delta d(u,v)$ around u or v has enough nodes so that every node has ≥ 3 neighbours in B
- ball B contains a beacon; desired neighbor-to-neighbor paths exist since the set of neighbour pairs induces an expander on B
- gossiping protocol finds such paths (lots of details omitted)
- \Rightarrow for δ -good pairs the bounds are OK

theorem proved





Fully distributed metric embedding

Overview of the algorithm [Slivkins SODA'05]:

- run our fully distributed triangulation algorithm
 - compute upper bound d^+ on distances to beacons
- run our GNP-style embedding on same beacon set
 - for distances to beacons, use d^+ instead of d
 - analysis is more complicated since d^+ is not a metric



Overview

- Introduction
- IDMaps-style triangulation
- GNP-style metric embedding
- “Fully distributed” approaches
 - general justification
 - triangulation and embedding with provable guarantees
- Extensions and open questions



Strongly doubling metrics

- Def ball of twice the radius has at most s times many points
 - generalize grids in low-dimensional Euclidean space
 - [Karger+Ruhl '02], [Krauthgamer+Lee '03], [Bartal+Mendel '04], ...
- Thm embedding into L_p with gracefully degrading distortion:
 - $\frac{3}{4}$ of node pairs embedded with distortion $\Delta = O(\log s)$,
 - $\frac{3}{4}$ of the remaining node pairs – with distortion 2Δ ,
 - $\frac{3}{4}$ of the remaining node pairs – with distortion 3Δ , ...
- Thm simple beacon-based embedding into L_p
 - select k nodes at random as “beacons”, embed them
 - embed each node at the nearest beacon




Triangulation in arbitrary metrics

- Lower bound d^- could be entirely useless
 - e.g. in uniform metric $d^- = 0$ for all node pairs
 - recall: $d^-(u,v) = \max_{\text{beacons } b} |d(u,b) - d(v,b)|$
- Upper bound $d^+(u,v) \leq 3d(u,v)$
 - with ϵ -slack, $f(\epsilon)$ beacons suffices
 - recall: $d^+(u,v) = \min_{\text{beacons } b} d(u,b) + d(v,b)$
- This factor of 3 is tight for triangulation
 - counterexample: complete bipartite graph
 - moreover, it is tight for any beacon-based estimate

Perfect triangulation: upper bound matches lower bound

- L_1 grid: $d(u,v) = X_{\text{offset}} + Y_{\text{offset}}$
- $d^- = |d(u, \star) - d(v, \star)| = d(u,v)$
- $d^+ = |d(u, \star) + d(v, \star)| = d(u,v)$
- Claim Select $f(\varepsilon)$ beacons u.a.r.
Then $d^-(u,v) = d^+(u,v) = d(u,v)$
for all but ε -fraction of pairs (u,v) .



- Thm This extends to **dense** point sets in d -dim L_1 :
 - minimal distance between any two points is 1
 - all coordinates lie in $[0, kn^{1/d}]$, k = density parameter
 - $f(\varepsilon, k, d)$ beacons suffices: #beacons independent of #nodes



Open questions

How can we extend our results to general metrics?

- Embed into L_p , $p \geq 1$ with ϵ -slack and distortion $f(\epsilon)$?
 - recall: works for doubling metrics
- How well can you embed an arbitrary metric
 - using only distances to beacons?
 - via a fully distributed algorithm?
- Converse questions on beacon-based triangulation
 - good guarantees (e.g. $d^+/d^- \leq 1 + \delta$) imply what?
- On-going follow-up work
 - [Chan, Dhamhere, Gupta, Kleinberg, Slivkins]