Spanners with Slack

Michael Dinitz
Joint work with Hubert Chan and Anupam Gupta

Computer Science Department
Carnegie Mellon University

Theory Lunch
September 27, 2006

Outline

Spanner Definition

- Main problem: small representation of metric space

Definition

Give a metric (V, d), a *t-spanner* $H = (V, E)$ is a weighted graph such that for *all* $u, v \in V$, $d(u, v) \leq d_H(u, v) \leq t \cdot d(u, v)$

- t is the *stretch* or the *distortion*
- $|E|$ measures how sparse or small the spanner is. *Really* want $|E| = O(n)$
- Want to minimize $|E|$ and t, i.e. create a low-stretch sparse spanner

Spanner example

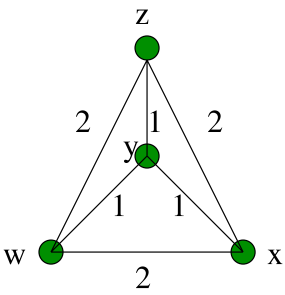

$$M = (V, d)$$

Spanner example

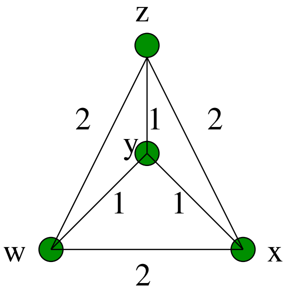

$M = (V, d)$

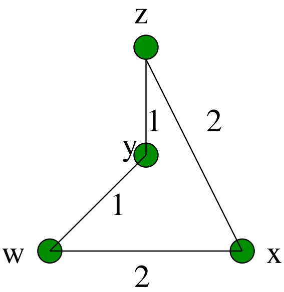

$H = (V, E)$

Spanner example

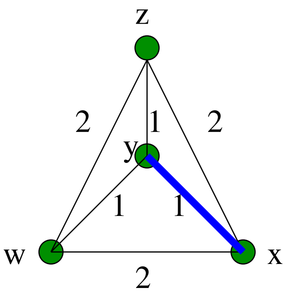

$M = (V, d)$

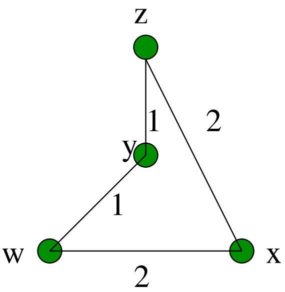

$H = (V, E)$

Spanner example

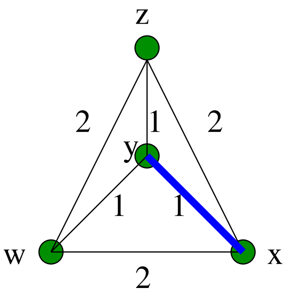

$M = (V, d)$

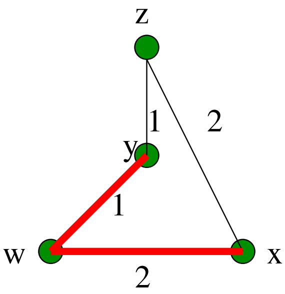

$H = (V, E)$

Spanner example

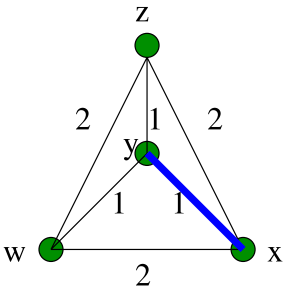

$M = (V, d)$

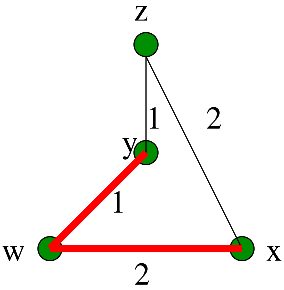

$H = (V, E)$

Stretch $= d_H(x, y)/d(x, y) = 3/1 = 3$

Research on Spanners

- Classic research:
    - Awerbuch '85: Inspired study of spanners
    - Peleg & Schaffer '89
    - Althofer, Das, Dobkin, Joseph, & Soares: Sparse spanners for weighted graphs
    - Euclidean spanners

Research on Spanners

- Classic research:
  - Awerbuch '85: Inspired study of spanners
  - Peleg & Schaffer '89
  - Althofer, Das, Dobkin, Joseph, & Soares: Sparse spanners for weighted graphs
  - Euclidean spanners
- New research
  - Baswana et al: Sparse *additive* spanners
  - Lower bounds for additive and Euclidean spanners

Simple Algorithm

Theorem (Althofer et al.)

For any integer k, a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges can be constructed efficiently

Simple Algorithm

Theorem (Althofer et al.)

For any integer k, a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges can be constructed efficiently

Use a Kruskal-like algorithm:

- Initialize H to be the empty graph
- Let $\{u, v\}$ be shortest edge we haven't looked at yet
- If $d_H(u, v) > (2k-1)d(u, v)$, put $\{u, v\}$ in H
- Otherwise discard $\{u, v\}$ and repeat

Althofer Example ($k = 2$)

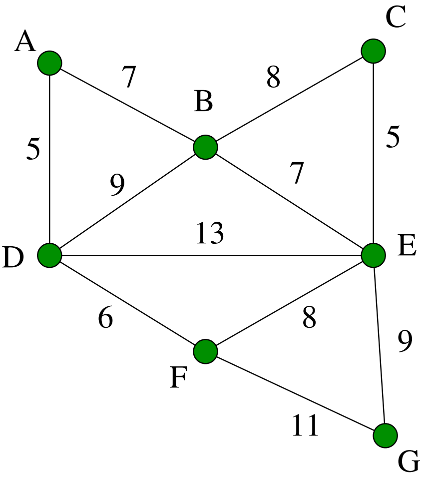

Althofer Example ($k = 2$)

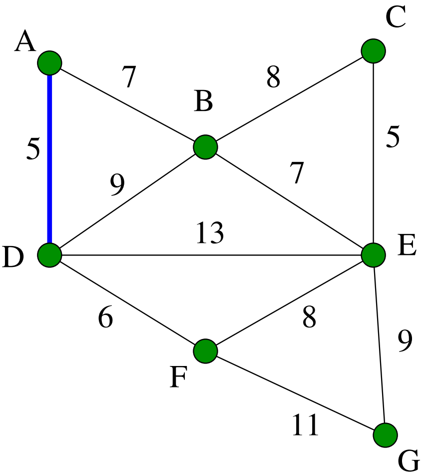

Althofer Example ($k = 2$)

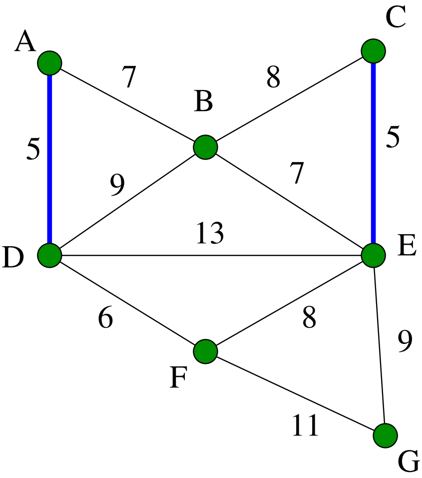

Althofer Example ($k = 2$)

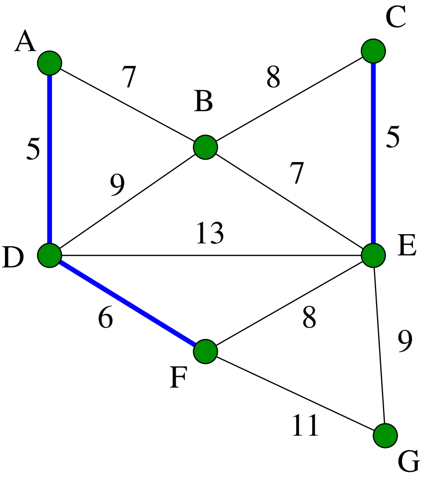

Althofer Example ($k = 2$)

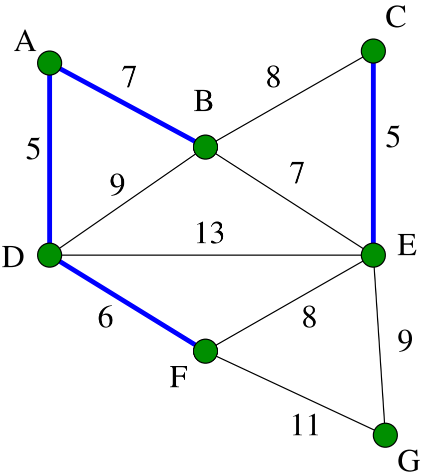

Althofer Example ($k = 2$)

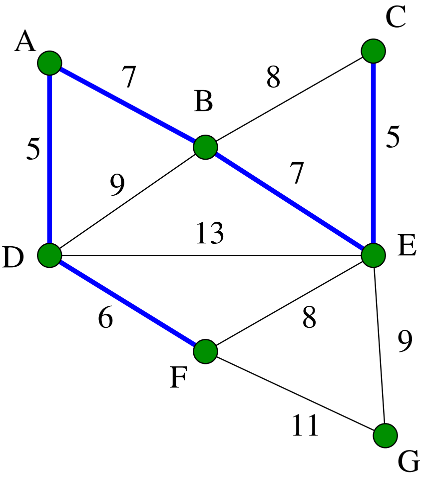

Althofer Example ($k = 2$)

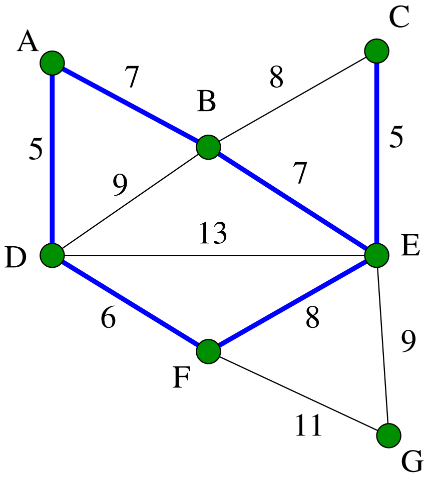

Althofer Example ($k = 2$)

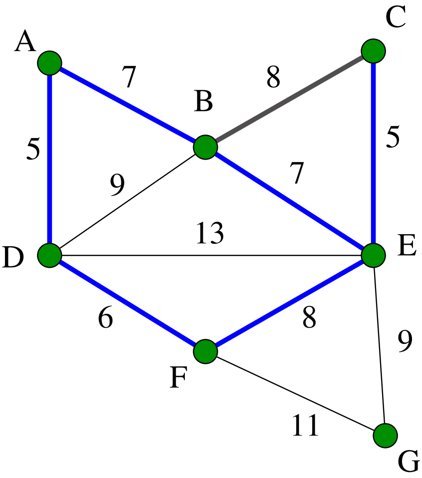

Althofer Example ($k = 2$)

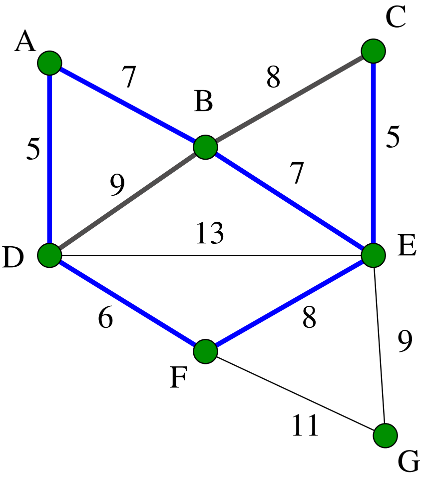

Althofer Example ($k = 2$)

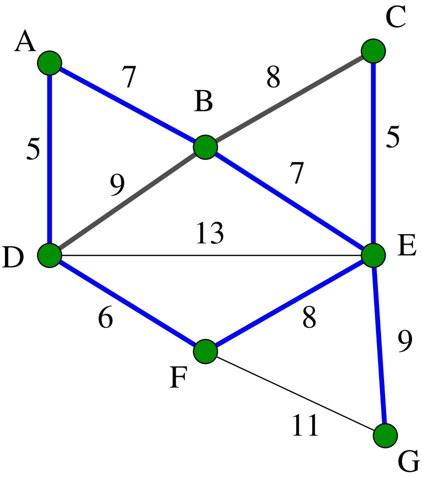

Althofer Example ($k = 2$)

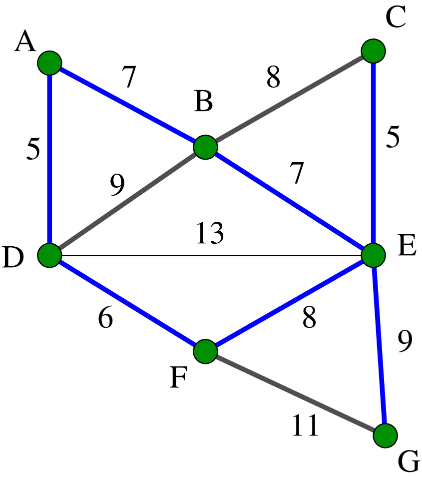

Althofer Example ($k = 2$)

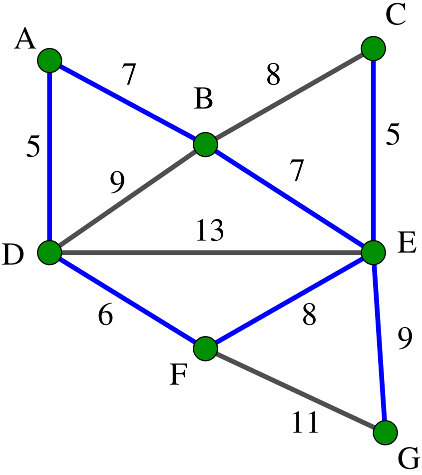

Althofer Example ($k = 2$)

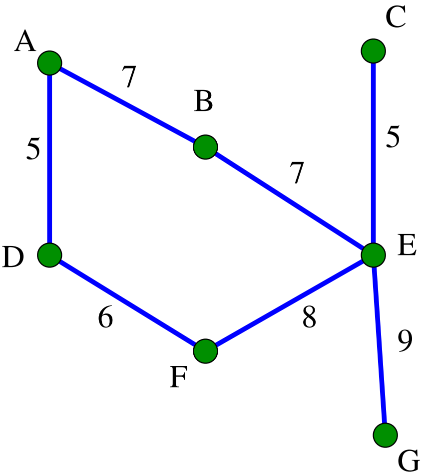

Correctness

Need to prove that stretch is at most $2k - 1$ and that $|E| = O(n^{1+1/k})$

Correctness

Need to prove that stretch is at most $2k - 1$ and that $|E| = O(n^{1+1/k})$

- Stretch: by construction.

Correctness

Need to prove that stretch is at most $2k-1$ and that $|E| = O(n^{1+1/k})$

- Stretch: by construction.
- Sparse:
  - Suppose edge $e = \{u, v\}$ creates a cycle C
  - Every other edge on C shorter than e
  - Without e, the distance between u and v was more than $(2k-1)\text{length}(e)$
  - So at least $2k-1$ other edges on C
  - Girth at least $2k-1$
- Well-known graph theory theorem: Girth of $2k-1$ implies $|E| = O(n^{1+1/k})$

Althofer is optimal

- Erdos girth conjecture: For every k, there is a graph with $\Omega(n^{1+1/k})$ edges and girth $2k-1$
- Implies that the Althofer spanner is tight (well, at least for subgraph spanners...)
- So if we want $O(n)$ edges, we need stretch of $\Omega(\log n)$!

What Now?

- In practice:
  1. $\log n$ stretch is too large
  2. Don't need low stretch for *all* pairs
- Use 2 to fix 1
- How well can we do? Ignore 5% of pairs and get $O(\sqrt{\log n})$ stretch on the rest? $O(\log \log n)$? $O(1)$?

What Now?

- In practice:
  1. $\log n$ stretch is too large
  2. Don't need low stretch for *all* pairs
- Use 2 to fix 1
- How well can we do? Ignore 5% of pairs and get $O(\sqrt{\log n})$ stretch on the rest? $O(\log \log n)$? $O(1)$?
- Ignoring a constant fraction of pairs lets us prove constant distortion on the rest!

ϵ-Neighborhoods

- Basic idea: ignoring small distances helps with large distances

Definition

Given ϵ, for any point $v \in V$, the ϵ-neighborhood $N_\epsilon(v)$ consists of the closest ϵn points to v

- $R(v, \epsilon) = \min\{r : |B(v, r)| \geq \epsilon n\}$
- v is ϵ-*far* from u if $d(u, v) \geq R(u, \epsilon)$

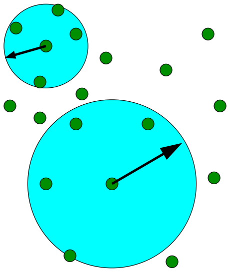

Slack definitions

- Original work on slack was on slack embeddings into ℓ_p spaces
- KSW '04, ABCDGKNS '05, ABN '06

Slack definitions

- Original work on slack was on slack embeddings into ℓ_p spaces
- KSW '04, ABCDGKNS '05, ABN '06
- Basic definition:

Definition (Slack Spanner)

Given a metric (V, d), a $t-$spanner $H = (V, E)$ has ϵ-slack if $d(u, v) \leq d_H(u, v) \leq t \cdot d(u, v)$ for all but ϵn^2 pairs $\{u, v\}$

Slack definitions

- Original work on slack was on slack embeddings into ℓ_p spaces
- KSW '04, ABCDGKNS '05, ABN '06
- Basic definition:

Definition (Slack Spanner)

Given a metric (V, d), a $t-$spanner $H = (V, E)$ has ϵ-slack if $d(u, v) \leq d_H(u, v) \leq t \cdot d(u, v)$ for all but ϵn^2 pairs $\{u, v\}$

- More restrictive definition:

Definition (Uniform Slack)

Given a metric (V, d), a t-spanner $H = (V, E)$ has ϵ-uniform slack if for all $u, v \in V$ such that v is ϵ-far from u, $d(u, v) \leq d_H(u, v) \leq t \cdot d(u, v)$

Conversion Theorem

Theorem

Suppose there exists an algorithm to construct a $t(n)$-stretch spanner with $h(n)$ edges for any metric. Then we can find an ϵ-slack spanner with $5 + 6t(\frac{1}{\epsilon})$ stretch and $n + h(\frac{1}{\epsilon})$ edges.

Conversion Theorem

Theorem

Suppose there exists an algorithm to construct a $t(n)$-stretch spanner with $h(n)$ edges for any metric. Then we can find an ϵ-slack spanner with $5 + 6t(\frac{1}{\epsilon})$ stretch and $n + h(\frac{1}{\epsilon})$ edges.

We can apply this to the Althofer spanner:

Corollary

For any metric, for any $0 < \epsilon < 1$, for any integer $k > 0$, there exists a $(12k - 1)$-spanner with ϵ-slack of size $n + O((\frac{1}{\epsilon})^{1+1/k})$

Density Net

- Intuition: Small set of points that approximates the metric
- Recall that $R(u, \epsilon) = \min\{r : |B(u, r)| \geq \epsilon n\}$

Definition

An ϵ-*density net* is a subset N of V such that

1. For all $x \in V$, there is some $y \in N$ s.t. $d(x, y) \leq 2R(x, \epsilon)$
2. $|N| \leq \frac{1}{\epsilon}$

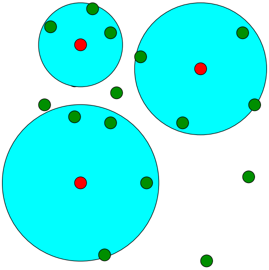

Constructing a Density Net

1. Put points in a list L by non-decreasing value of $R(\cdot, \epsilon)$
2. Initialize $N := 0$.
3. While L is non-empty:
    1. Remove first point v from L
    2. If there exists $u \in N$ s.t. $N_\epsilon(v)$ and $N_\epsilon(u)$ intersect, then discard v; otherwise add v to N

Density Net Example

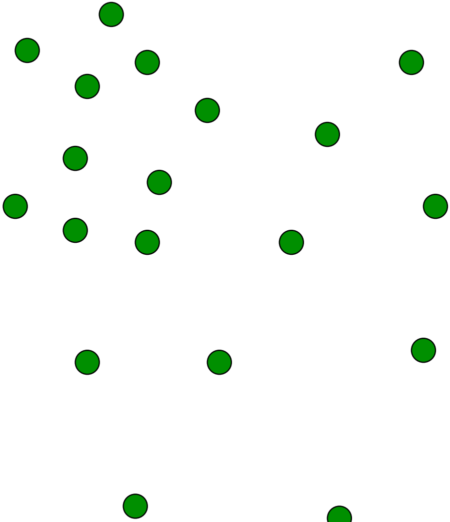

Density Net Example

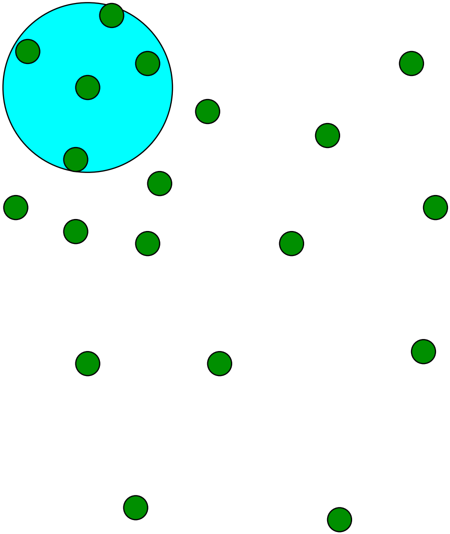

Density Net Example

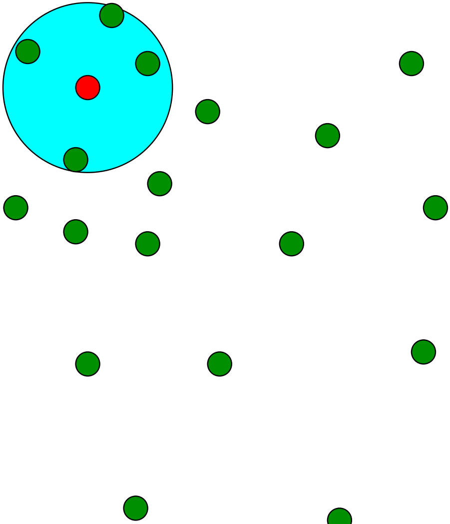

Density Net Example

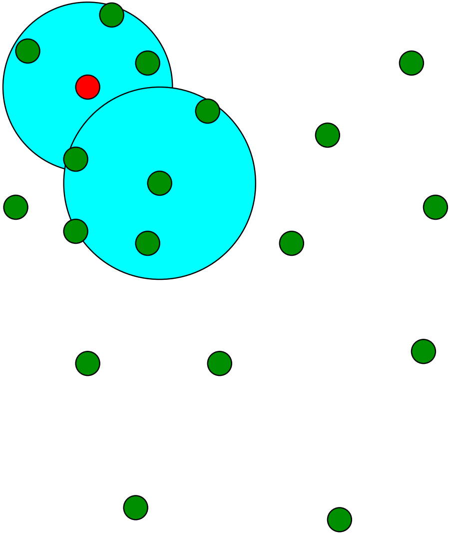

Density Net Example

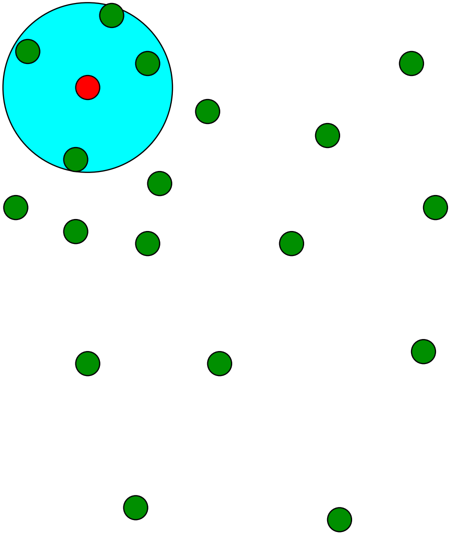

Density Net Example

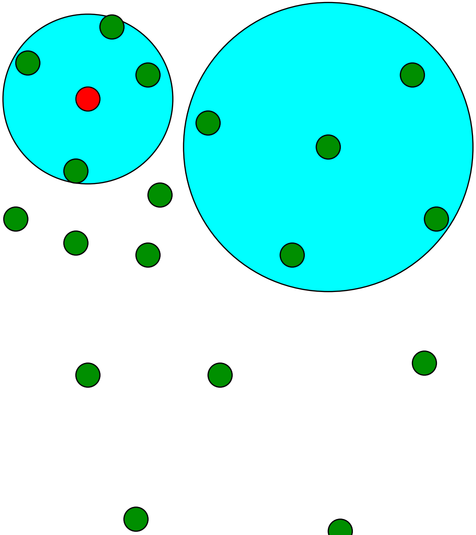

Density Net Example

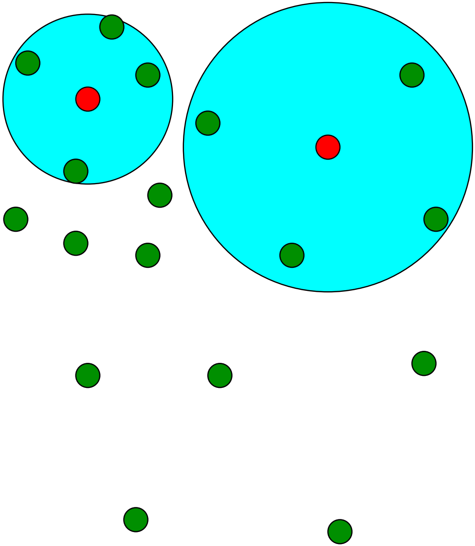

Density Net Example

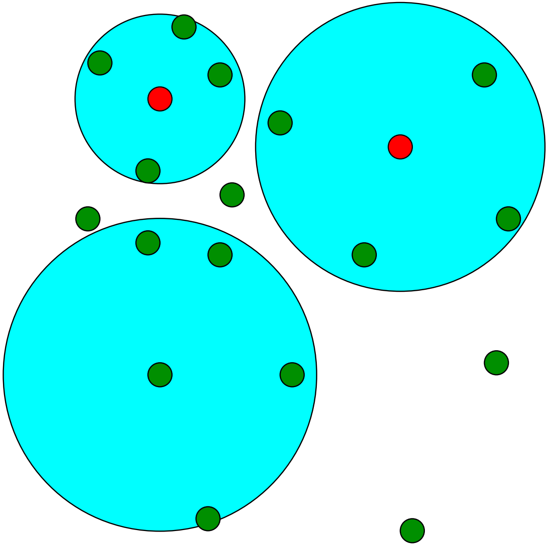

Density Net Example

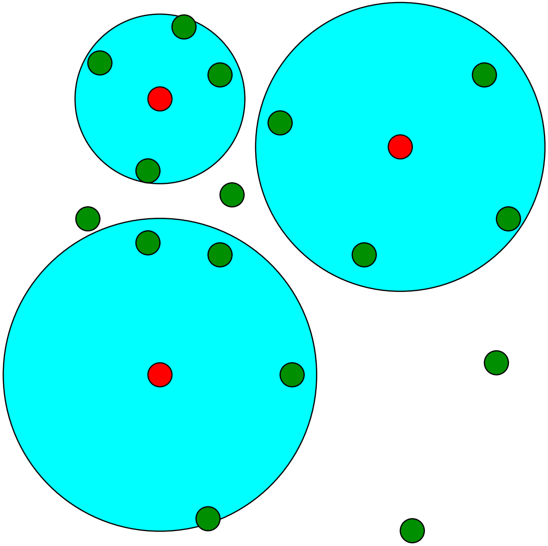

Density Net Proof

- Need to prove:
  1. For all $x \in V$ there is some $y \in N$ such that $d(x, y) \leq 2R(x, \epsilon)$
  2. $|N| \leq \frac{1}{\epsilon}$

Density Net Proof

- Need to prove:
  1. For all $x \in V$ there is some $y \in N$ such that $d(x, y) \leq 2R(x, \epsilon)$
  2. $|N| \leq \frac{1}{\epsilon}$
- Net property:
  - If $x \in N$ then we're good.
  - Else there is $y \in N$ before x s.t. $N_\epsilon(x)$ and $N_\epsilon(y)$ intersect. So $d(x, y) \leq R(x, \epsilon) + R(y, \epsilon) \leq 2R(x, \epsilon)$

Density Net Proof

- Need to prove:
  1. For all $x \in V$ there is some $y \in N$ such that $d(x, y) \leq 2R(x, \epsilon)$
  2. $|N| \leq \frac{1}{\epsilon}$
- Net property:
  - If $x \in N$ then we're good.
  - Else there is $y \in N$ before x s.t. $N_\epsilon(x)$ and $N_\epsilon(y)$ intersect. So $d(x, y) \leq R(x, \epsilon) + R(y, \epsilon) \leq 2R(x, \epsilon)$
- Size property:
  - For different $u, v \in N$, $N_\epsilon(u)$ and $N_\epsilon(v)$ are disjoint
  - Each $|N_\epsilon(u)| \geq \epsilon n$, so $|N| \leq \frac{1}{\epsilon}$

Conversion Algorithm

1. Given metric (V, d),construct ϵ-density net N. Note that $|N| \leq \frac{1}{\epsilon}$
2. Construct $t(\frac{1}{\epsilon})$-spanner with $h(\frac{1}{\epsilon})$ edges on N
3. For all $u \in V \setminus N$, add an edge to the nearest point in N

Conversion Algorithm

1. Given metric (V, d),construct ϵ-density net N. Note that $|N| \leq \frac{1}{\epsilon}$
2. Construct $t(\frac{1}{\epsilon})$-spanner with $h(\frac{1}{\epsilon})$ edges on N
3. For all $u \in V \setminus N$, add an edge to the nearest point in N

Obviously sparse: $O(n + h(\frac{1}{\epsilon}))$ edges

Conversion Example

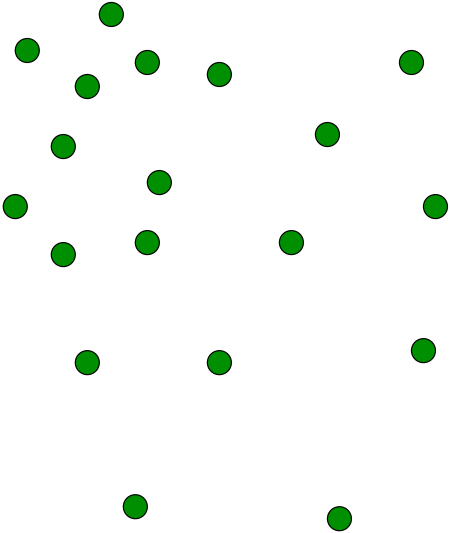

Conversion Example

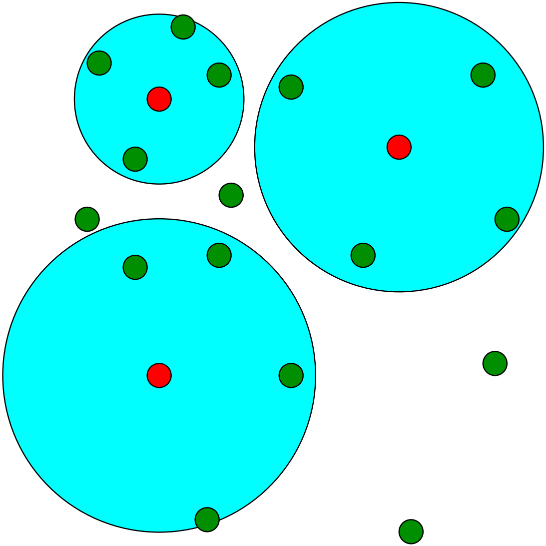

Conversion Example

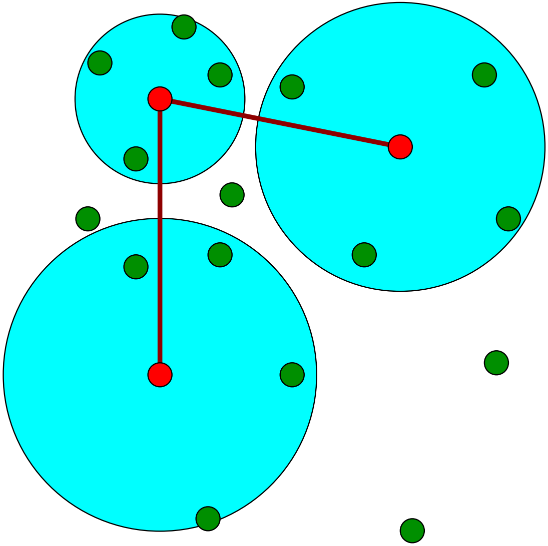

Conversion Example

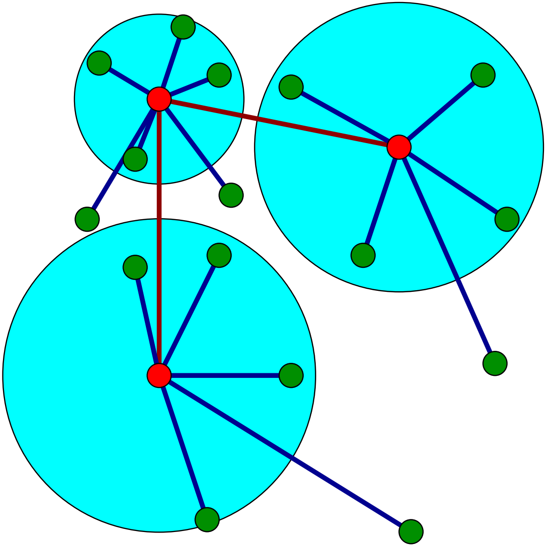

Low Stretch

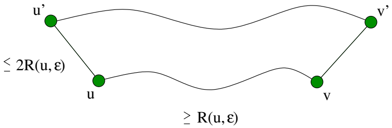

- Let $u, v \in V$ s.t. $v \notin N_\epsilon(u)$
- Let u', v' be the closest points in N to u and v respectively

Low Stretch

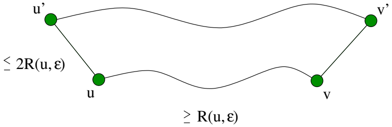

- Let $u, v \in V$ s.t. $v \notin N_\epsilon(u)$
- Let u', v' be the closest points in N to u and v respectively
- $d(u, u') \leq 2R(u, \epsilon) \leq 2d(u, v)$
- $d(v, v') \leq d(v, u') \leq d(v, u) + d(u, u') \leq 3d(u, v)$
- $d(u', v') \leq d(u', u) + d(u, v) + d(v, v') \leq 6d(u, v)$

Low Stretch

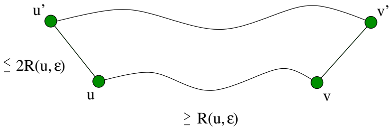

- Let $u, v \in V$ s.t. $v \notin N_\epsilon(u)$
- Let u', v' be the closest points in N to u and v respectively
- $d(u, u') \leq 2R(u, \epsilon) \leq 2d(u, v)$
- $d(v, v') \leq d(v, u') \leq d(v, u) + d(u, u') \leq 3d(u, v)$
- $d(u', v') \leq d(u', u) + d(u, v) + d(v, v') \leq 6d(u, v)$
- By spanner on N, $d_H(u', v') \leq t(\frac{1}{\epsilon})d(u', v') \leq 6t(\frac{1}{\epsilon})d(u, v)$
- So $d_H(u, v) \leq d(u, u') + d_H(u', v') + d(v', v) \leq (5 + 6t(\frac{1}{\epsilon}))d(u, v)$

Subgraph Spanner

- What if our input isn't a metric but a graph?
- Want our spanner to be a subgraph

Subgraph Spanner

- What if our input isn't a metric but a graph?
- Want our spanner to be a subgraph

Theorem

Given a weighted graph $G = (V, E)$, for any integer $k > 0$ and any $0 < \epsilon < 1$, there exists a $(12k - 1)$-spanner with ϵ-slack and $O(n + \sqrt{n}(\frac{1}{\epsilon})^{1+1/k})$ edges.

- Uses pairwise distance preservers of Coppersmith and Elkin to make a subgraph that emulates the spanner on the net

Low Weight

Could also try to minimize the *weight* of the spanner.

Theorem

For any metric, there is an ϵ*-slack spanner with* $O(\log \frac{1}{\epsilon})$ *stretch,* $O(n + \frac{1}{\epsilon})$ *edges, and weight* $O(\log^2(\frac{1}{\epsilon})) \times wt(MST)$

Main idea: use LASTs (Light Approximate Shortest-path Trees)

Gracefully Degrading

- Previous results of the form "You give me an ϵ, I give you an ϵ-slack spanner"
- Could ask for something stronger

Gracefully Degrading

- Previous results of the form "You give me an ϵ, I give you an ϵ-slack spanner"
- Could ask for something stronger
- "I give you a spanner that works *simultaneously* for *all* ϵ"
- Called a *gracefully degrading* spanner

Gracefully Degrading

- Previous results of the form "You give me an ϵ, I give you an ϵ-slack spanner"
- Could ask for something stronger
- "I give you a spanner that works *simultaneously* for *all* ϵ"
- Called a *gracefully degrading* spanner

Theorem

For any metric, there is a spanner H with $O(n)$ edges s.t. for any $0 < \epsilon < 1$, H is a $O(\log \frac{1}{\epsilon})$-spanner with ϵ-slack.

Gracefully Degrading Construction

Intuition: layers of slack spanners for various value of ϵ.

Gracefully Degrading Construction

Intuition: layers of slack spanners for various value of ϵ.
Actually much simpler – only 2 layers needed:

Gracefully Degrading Construction

Intuition: layers of slack spanners for various value of ϵ.
Actually much simpler – only 2 layers needed:

1. Let $\epsilon_0 = n^{-1/2}$, and construct a ϵ_0-density net N of V
2. Connect every vertex to the closest point in N
3. Create a 1-spanner H_0 (e.g. a clique) on N (uses $O(n)$ edges)
4. Use Althofer to make a $\log n$-spanner H' on V
5. Set H to be the union of H_0 and H', together with edges that connect each point in V to its closest point in N

Gracefully Degrading Construction

Intuition: layers of slack spanners for various value of ϵ.
Actually much simpler – only 2 layers needed:

1. Let $\epsilon_0 = n^{-1/2}$, and construct a ϵ_0-density net N of V
2. Connect every vertex to the closest point in N
3. Create a 1-spanner H_0 (e.g. a clique) on N (uses $O(n)$ edges)
4. Use Althofer to make a $\log n$-spanner H' on V
5. Set H to be the union of H_0 and H', together with edges that connect each point in V to its closest point in N

Each step creates $O(n)$ edges, so there are only $O(n)$ edges total

Stretch

Two cases for the stretch:

Stretch

Two cases for the stretch:

1. $\epsilon < \epsilon_0$: Use H' to get stretch $O(\log n) = O(\log n^{1/2}) = O(\log \frac{1}{\epsilon_0}) = O(\log \frac{1}{\epsilon})$ between every pair of points

Stretch

Two cases for the stretch:

1. $\epsilon < \epsilon_0$: Use H' to get stretch $O(\log n) = O(\log n^{1/2}) = O(\log \frac{1}{\epsilon_0}) = O(\log \frac{1}{\epsilon})$ between every pair of points
   - Magic of logs in $O(\cdot)$ notation

Stretch

Two cases for the stretch:

1. $\epsilon < \epsilon_0$: Use H' to get stretch $O(\log n) = O(\log n^{1/2}) = O(\log \frac{1}{\epsilon_0}) = O(\log \frac{1}{\epsilon})$ between every pair of points
    - Magic of logs in $O(\cdot)$ notation
2. $\epsilon \geq \epsilon_0$: Use H_0. Same analysis as for slack spanner, except that stretch in the net is 1, so total stretch is at most 11.

Average Stretch

Gracefully degrading spanner automatically gives us a normal $O(\log n)$-spanner that has $O(1)$ average distortion!

$$\begin{aligned}\frac{1}{\binom{n}{2}} \sum_{\{x,y\} \in \binom{V}{2}} \frac{d_H(x,y)}{d(x,y)} &= \frac{2}{n} \sum_{x \in V} \frac{1}{n-1} \sum_{y \neq x} \frac{d_H(x,y)}{d(x,y)} \\ &\leq \frac{2}{n} \sum_{x \in V} \left(\frac{1}{n^{1/2}} O(\log n) + (1 - \frac{1}{n^{1/2}}) \cdot 11 \right) \\ &= O(1)\end{aligned}$$

Distance Oracles Overview

- Intuition: *all*-pairs shortest path is rarely necessary.
- Distance oracle: data structure/algorithm for computing approximate distances in a metric
- Want to minimize stretch, space, and query time
- First studied by Thorup and Zwick ('01): for any integer $k \geq 1$, oracle with stretch $2k - 1$, space $O(kn^{1+1/k})$, query time $O(k)$
- Implicitly created a spanner, clever way of doing queries based on special structure of spanner

Oracles with Slack

Can create slack oracles using slack embeddings:

Theorem (ABN '06)

For any integer $k \geq 1$, there is an oracle with ϵ-slack, stretch $6k-1$, $O(k)$ query time, and $O(n \log n \log \frac{1}{\epsilon} + k \log n (\frac{1}{\epsilon} \log \frac{1}{\epsilon})^{1+1/k})$ space

Oracles with Slack

Can create slack oracles using slack embeddings:

Theorem (ABN '06)

For any integer $k \geq 1$, there is an oracle with ϵ-slack, stretch $6k-1$, $O(k)$ query time, and $O(n \log n \log \frac{1}{\epsilon} + k \log n (\frac{1}{\epsilon} \log \frac{1}{\epsilon})^{1+1/k})$ space

But slack spanners are better:

Theorem

For every integer $k \geq 1$, there is an oracle with ϵ-slack, stretch $10k-1$, $O(k)$ query times, and $O(n + k(\frac{1}{\epsilon})^{1+1/k})$ space

Same method as used for slack spanners

Gracefully Degrading Oracles

Can do the same thing for gracefully degrading oracles.

Theorem

For any integer k with $1 \leq k \leq O(\log n)$, there is a distance oracle with worst cast stretch of $2k - 1$ and $O(k)$ query time that uses $O(kn^{1+1/k})$ space such that the average distortion and distortion of average are $O(1)$

Improvement over ABN '06 if $k = o(\log n)$

Distance Labeling Overview

- How can we assign each point a short label so that approximate distances can be computed quickly by just comparing labels?
- Used in various networking applications
- Embedding into ℓ_p very natural approach: size of a label is the dimension

Distance Labeling Overview

- How can we assign each point a short label so that approximate distances can be computed quickly by just comparing labels?
- Used in various networking applications
- Embedding into ℓ_p very natural approach: size of a label is the dimension
- One of the original motivations for definition of slack in KSW '04:
    - In general can't have dimension less that $\Omega(\log n)$!
    - Seems to work better in practice

Distance Labeling Overview

- How can we assign each point a short label so that approximate distances can be computed quickly by just comparing labels?
- Used in various networking applications
- Embedding into ℓ_p very natural approach: size of a label is the dimension
- One of the original motivations for definition of slack in KSW '04:
  - In general can't have dimension less that $\Omega(\log n)$!
  - Seems to work better in practice
- Can we do better with spanners than with embeddings?

Slack Labelings

Using embeddings:

Theorem (ABCDGKNS '05)

Any embedding $\varphi : V \to \ell_p$ with ϵ-(uniform) slack must have dimension that depends on $\log n$

Slack Labelings

Using embeddings:

Theorem (ABCDGKNS '05)

Any embedding $\varphi : V \to \ell_p$ with ϵ-(uniform) slack must have dimension that depends on $\log n$

We get rid of all dependence on n by not using an embedding!

Theorem

For any integer k with $1 \le k \le \log \frac{1}{\epsilon}$, we can assign each point a label that uses $O((\frac{1}{\epsilon})^{1/k} \log^{1-1/k} \frac{1}{\epsilon})$ space so that if v is ϵ-far from u, their distance can be computed up to stretch $12k - 1$ in $O(k)$ time

Review

- Ignoring a constant fraction of distances gives us lots of power (e.g. constant stretch, linear size spanners)!
- Using ϵ-density nets to represent metrics gives us good slack and gracefully degrading spanners, distance oracles, and distance labelings

Future Research

- Slack version of (your favorite problem here)
- Additive spanners????

Thank You!

(and please sign up to give your very own theory lunch talk)