Recent Progress on Euclidean (and Related) Spanners

Hung Le

UMassAmherst

Manning College of Information & Computer Sciences

Simons Institute, 01 Aug 2024

Metric t-Spanners

A t-spanner of a point set P in a metric space (X, d_X) is a graph $S = (P, E_S, d_X)$ such that:

$$d_S(u, v) \leq t \cdot d_X(u, v) \qquad \forall u \neq v \in P \tag{1}$$

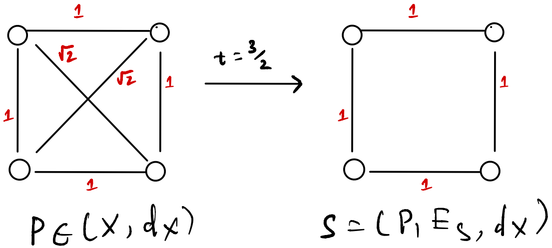

t is called the stretch of the spanner.

- Sparsity: $\frac{|E_S|}{|E(MST)|} = \frac{|E_S|}{n-1}$.
- Lightness $\frac{\omega(E_S)}{\omega(MST)}$

Graph t-Spanners

The t-spanner S is a spanning subgraph of the input graph G, preserving distances up to a factor of t.

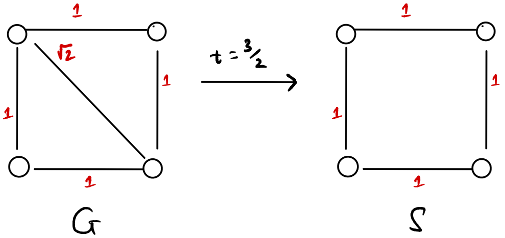

- Sparsity: $\frac{|E_S|}{|E(MST)|} = \frac{|E_S|}{n-1}$.
- Lightness $\frac{\omega(E_S)}{\omega(MST)}$

Doubling Metrics

Two metrics: Euclidean $(\mathbb{R}^d, \|\cdot\|_2)$, and doubling metrics [Assouad;30][Gupta,Krauthgamer,Lee;03].

Doubling Metrics

Two metrics: Euclidean $(\mathbb{R}^d, \|\cdot\|_2)$, and doubling metrics [Assouad;30][Gupta,Krauthgamer,Lee;03].

Doubling Metrics

A metric (X, d_X) has doubling dimension d if any ball of radius R can be covered by 2^d balls of radius $R/2$, $\forall\ R \geq 0$.

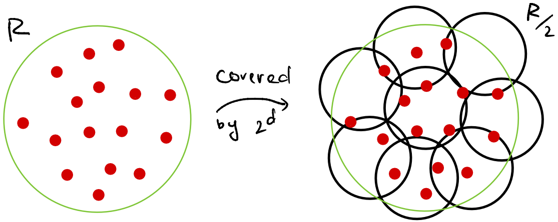

- n-point set in $(\mathbb{R}^d, \|\cdot\|_2)$ has doubling dimension $d + O(1)$.
- Any metric space has doubling dimension $O(\log n)$.

Basic Results: Metric Spanners

[Chew86] constructed $O(1)$-spanner with $O(n)$ edges (or $O(1)$ sparsity) for $P \in \mathbb{R}^2$.

Basic Results: Metric Spanners

[Chew86] constructed $O(1)$-spanner with $O(n)$ edges (or $O(1)$ sparsity) for $P \in \mathbb{R}^2$.

$(1+\epsilon)$-spanner for $P \in \mathbb{R}^d$ with:

- $O_d((1/\epsilon)^{d-1})$ sparsity [Yao;82][Clarkson;87][Keil;88] [Ruppert,Seidel;91][Althofer,Das, Dobkin, Joseph,Soares;93] [Callahan,Kosaraju;93] ...
- $O_d((1/\epsilon)^d)$ lightness [Das,Heffernan,Narasimhan;93] [Das,Narasimhan,Salowe;95] [L,Solomon;19]

Book "Geometric Spanner Networks" by Narasimhan and Smid.

Basic Results: Metric Spanners

[Chew86] constructed $O(1)$-spanner with $O(n)$ edges (or $O(1)$ sparsity) for $P \in \mathbb{R}^2$.

$(1+\epsilon)$-spanner for $P \in \mathbb{R}^d$ with:

- $O_d((1/\epsilon)^{d-1})$ sparsity [Yao;82][Clarkson;87][Keil;88] [Ruppert,Seidel;91][Althofer,Das, Dobkin, Joseph,Soares;93] [Callahan,Kosaraju;93] ...
- $O_d((1/\epsilon)^d)$ lightness [Das,Heffernan,Narasimhan;93] [Das,Narasimhan,Salowe;95] [L,Solomon;19]

Book "Geometric Spanner Networks" by Narasimhan and Smid.

$(1+\epsilon)$-spanner for P in doubling metric of dimension d:

- $O_d((1/\epsilon)^d)$ sparsity [Gao, Guibas,Nguyen;04][Har-Peled,Mendel;04][Smid,09][Chan,Gupta,Maggs,Zhou;16]
- $O_d((1/\epsilon)^{d+1})$ lightness [Gottlieb;15][Filtser,Solomon;16] [Borradaile,L,Wulff-Nilsen;19][L,Solomon;23]

Basic Results: Metric Spanners (cont.)

$(1+\epsilon)$-spanner for $P \in \mathbb{R}^d$ with:

- $O_d((1/\epsilon)^{d-1})$ sparsity
- $O_d((1/\epsilon)^d)$ lightness

$(1+\epsilon)$-spanner for P in doubling metric of dimension d:

- $O_d((1/\epsilon)^d)$ sparsity
- $O_d((1/\epsilon)^{d+1})$ lightness

Observations:

1. Can get stretch $1+\epsilon$ in metric setting.
2. Sparsity and lightness are off by a factor of $1/\epsilon$.
3. Doubling is off by a factor of $1/\epsilon$ compared to Euclidean, due to Euclidean geometry.

Basic Results: Graph Spanners

$(2k-1)$-spanner with sparsity $O(n^{1/k})$ for any $k \geq 1$.
[Peleg,Schäffer;89][Althöfer,Das, Dobkin, Joseph,Soares;93]
[Halperin,Zwick;96]

- Tight under Erdős' Girth Conjecture.

Basic Results: Graph Spanners

$(2k-1)$-spanner with sparsity $O(n^{1/k})$ for any $k \geq 1$.
[Peleg,Schäffer;89][Althöfer,Das, Dobkin, Joseph,Soares;93]
[Halperin,Zwick;96]

- Tight under Erdős' Girth Conjecture.

$(2k-1)(1+\epsilon)$-spanner with lightness $O(n^{1/k}/\epsilon)$.
[Althöfer,Das,Dobkin,Joseph,Soares;93]
[Chandra,Das,Narasimhan,Soares;92]
[Elkin,Neiman,Solomon;14][Chechik,Wulff-Nilsen;16]
[L,Solomon;13][Bodwin;24]

Basic Results: Graph Spanners

$(2k-1)$-spanner with sparsity $O(n^{1/k})$ for any $k \geq 1$.
[Peleg,Schäffer;89][Althöfer,Das, Dobkin, Joseph,Soares;93]
[Halperin,Zwick;96]

- Tight under Erdős' Girth Conjecture.

$(2k-1)(1+\epsilon)$-spanner with lightness $O(n^{1/k}/\epsilon)$.
[Althöfer,Das,Dobkin,Joseph,Soares;93]
[Chandra,Das,Narasimhan,Soares;92]
[Elkin,Neiman,Solomon;14][Chechik,Wulff-Nilsen;16]
[L,Solomon;13][Bodwin;24]

- The dependency on ϵ seems necessary, at least for $\epsilon = kn^{-1/(2k-2)}$ [Bodwin,Flics;24].

Erdős' Girth Conjecture

$G = (V, E)$ with $|E| = \Omega(n^{1+1/k})$ and girth $\geq 2k+1$ exists.

One Algorithm Fits All

<u>Greedy</u>$(G(V,E,w),t)$

sort edges in E in the increasing weight order

$S \leftarrow (V, \emptyset, w)$

for each edge $e=(u,v)$ in the sorted order

  if $d_S(u,v) > t \cdot w(e)$

    $E_S \leftarrow E_S \cup \{e\}$

return S

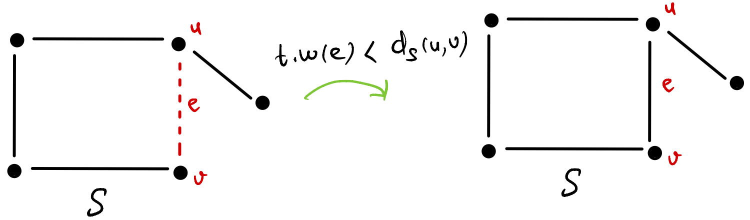

- Achieve all aformentioned bounds in all settings.
- Existentially optimal [Filtser,Solomon;16].

Outline

1. ~~Introduction.~~
2. ~~Basic Results.~~
3. Fault Tolerance.
4. Approximation/Instance Optimality.
5. Structural Alternatives.

Fault-Tolerant Spanners

[Levcopoulos,Narasimhan,Smid;98]: removing at most f points still leaves a t-spanner.

Metric Fault Tolerance

A t-spanner S is f-fault-vertex-tolerant if for any set $F \subseteq P$ of at most f points, $S \setminus F$ is a t-spanner of $P \setminus F$.

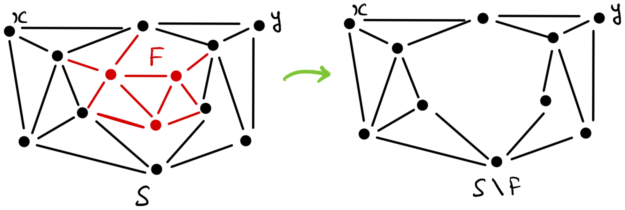

Graph Fault Tolerance

For graph $G = (V, E, \omega)$, S is f-fault-vertex-tolerant if $S \setminus F$ is a t-spanner of $G \setminus F$ for any set F of at most f vertices.

Fault-Tolerant Metric Spanners

Lower bounds for any constant ϵ, d:

- $\Omega(f)$ for sparsity: every vertex must have degree at least f.
- $\Omega(f^2)$ for ligntness [Czumaj,Zhao;03].

Fault-Tolerant Metric Spanners

Lower bounds for any constant ϵ, d:

- $\Omega(f)$ for sparsity: every vertex must have degree at least f.
- $\Omega(f^2)$ for ligntness [Czumaj,Zhao;03].

Tolerating f vertex (edge) faults in $\mathbb{R}^d$:

- $O_{\epsilon,d}(f)$ sparsity.[Levcopoulos,Narasimhan,Smid;98] [Lukovszki;99]
- $O_{\epsilon,d}(f^2)$ lightness. [Czumaj,Zhao;03]

Fault-Tolerant Metric Spanners

Lower bounds for any constant ϵ, d:

- $\Omega(f)$ for sparsity: every vertex must have degree at least f.
- $\Omega(f^2)$ for ligntness [Czumaj,Zhao;03].

Tolerating f vertex (edge) faults in $\mathbb{R}^d$:

- $O_{\epsilon,d}(f)$ sparsity.[Levcopoulos,Narasimhan,Smid;98] [Lukovszki;99]
- $O_{\epsilon,d}(f^2)$ lightness. [Czumaj,Zhao;03]

These results are extensible to doubling metrics. [Chan,Li,Ning;12] [Chan,Li,Ning,Solomon;13][Solomon;14][L,Solomon,Than;23]

Fault-Tolerant Graph Spanners

Recall: Erdős' Girth Conjecture (EGC) implies $\Omega(n^{1/k})$ sparsity lower bound for $(2k-1)$-spanners.

- $\Omega(f^{1-1/k}n^{1/k})$ sparsity/lightness lower bound for f vertex faults, assuming EGC. [Bodwin,Dinitz,Parter,Vassilevska Williams;18]

Fault-Tolerant Graph Spanners

Recall: Erdős' Girth Conjecture (EGC) implies $\Omega(n^{1/k})$ sparsity lower bound for $(2k-1)$-spanners.

- $\Omega(f^{1-1/k}n^{1/k})$ sparsity/lightness lower bound for f vertex faults, assuming EGC. [Bodwin,Dinitz,Parter,Vassilevska Williams;18]
- Sparsity $O(f^{1-1/k}n^{1/k})$. [Chechik,Langberg,Peleg,Roditty;10] Dinitz,Krauthgamer;11][Bodwin,Dinitz,Parter,Vassilevska Williams;18] [Bodwin,Patel;19][Dinitz,Robelle;20] [Bodwin,Dinitz,Robelle;21][Parter;22]

Fault-Tolerant Graph Spanners

Recall: Erdős' Girth Conjecture (EGC) implies $\Omega(n^{1/k})$ sparsity lower bound for $(2k-1)$-spanners.

- $\Omega(f^{1-1/k}n^{1/k})$ sparsity/lightness lower bound for f vertex faults, assuming EGC. [Bodwin,Dinitz,Parter,Vassilevska Williams;18]
- Sparsity $O(f^{1-1/k}n^{1/k})$. [Chechik,Langberg,Peleg,Roditty;10] Dinitz,Krauthgamer;11][Bodwin,Dinitz,Parter,Vassilevska Williams;18] [Bodwin,Patel;19][Dinitz,Robelle;20] [Bodwin,Dinitz,Robelle;21][Parter;22]

Open problem: Light fault-tolerant graph spanners?

Massive failure: $f = \Omega(n)$

Massive failure: $f = \Omega(n)$

Define a reasonable model for massive failure is tricky.

- ν-reliable: faulty set F , there exists $F \subseteq F^+$ such that $|F^+| \leq (1+\nu)|F|$ and distances are preserved for points in $P \setminus F^+$ in $S \setminus F$.[Bose,Dujmovic,Morin,Smid;13][Buchin,Har-Peled,Oláh;20]

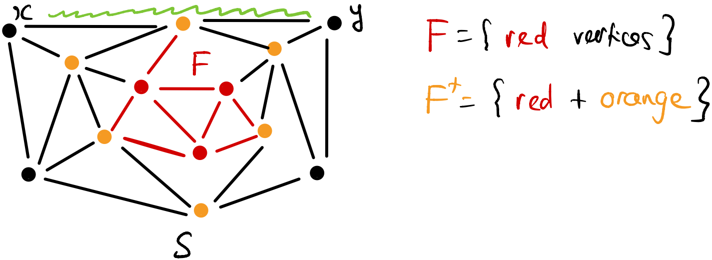

Remark: this model does not work for graphs: (reliably) preserving connectivity needs $\Omega(n^2)$ edges. [Filtser,L;22]

Massive failure: $f = \Omega(n)$

Lower bounds:

- Any reliable $O(1)$-spanner in $\mathbb{R}^d$ must have sparsity $\Omega(\log n)$. [Bose,Dujmovic,Morin,Smid;13].
- (Oblivious) lightness lower bound $\Omega(\log n)$ for any finite stretch. [Filtser,Gitlitz,Neiman;23]

Massive failure: $f = \Omega(n)$

Lower bounds:

- Any reliable $O(1)$-spanner in $\mathbb{R}^d$ must have sparsity $\Omega(\log n)$. [Bose,Dujmovic,Morin,Smid;13].
- (Oblivious) lightness lower bound $\Omega(\log n)$ for any finite stretch. [Filtser,Gitlitz,Neiman;23]

Upper bounds:

- ν-reliable $(1+\epsilon)$-spanner with $O_{\nu,d,\epsilon}(\log n)$ edges in $\mathbb{R}^d$ and doubling dim. d. [Buchin,Har-Peled,Oláh;19][Filtser,L;22].
- (Oblivious) lightness $O_{\epsilon,d,\nu}(\log n)$. [Filtser,Gitlitz,Neiman;23]

Massive failure: $f = \Omega(n)$

Edge counter part of ν-reliable: ψ-dependable where every edge of the spanner may fail independently with probability $1 - \psi$. [Har-Peled,Lusardi;24]

- Sparsity $O_{\epsilon,d}(\log n (1/\psi)^{4/3})$ in $\mathbb{R}^d$, and lower bound $O_{\epsilon,d}(\log n/\psi))$. [Har-Peled,Lusardi;24]

Massive failure: $f = \Omega(n)$

Color fault-tolerant spanners: vertices (or edges) are partitioned into color classes, and at most f color classes might fail. [Petruschka,Sapir,Tzalik;23]

- Sparsity $O(fn^{1/k})$ for edge-colored and $O(f^{1-1/k}n^{1/k})$ for vertex-colored. [Petruschka,Sapir,Tzalik;23]

Massive failure: $f = \Omega(n)$

Bounded degree fault: at most f faulty edges incident to a vertex. [Bodwin,Haeupler,Parter;24]

- Sparsity $O_k(f^{1-1/k}n^{1/k})$ for graph $(2k-1)$-spanners. [Bodwin,Haeupler,Parter;24]
- For points in $\mathbb{R}^d$, bounded degree fault f with $O_{\epsilon,d}(f)$ sparsity. [Biniaz,Carufel,Maheshwari,Smid;24]
- For points in doubling dimension d, $O_{\epsilon,d}(f^2)$ sparsity. [Biniaz,Carufel,Maheshwari,Smid;24]

Massive failure: $f = \Omega(n)$

Bounded degree fault: at most f faulty edges incident to a vertex. [Bodwin,Haeupler,Parter;24]

- Sparsity $O_k(f^{1-1/k}n^{1/k})$ for graph $(2k-1)$-spanners. [Bodwin,Haeupler,Parter;24]
- For points in $\mathbb{R}^d$, bounded degree fault f with $O_{\epsilon,d}(f)$ sparsity. [Biniaz,Carufel,Maheshwari,Smid;24]
- For points in doubling dimension d, $O_{\epsilon,d}(f^2)$ sparsity. [Biniaz,Carufel,Maheshwari,Smid;24]

Open problems:

1. Closing the gap between $\mathbb{R}^d$ and doubling dimension d.
2. Light versions of color or bounded-degree faults.
3. Other models of massive failure

Outline

1. ~~Introduction.~~
2. ~~Basic Results.~~
3. ~~Fault Tolerance.~~
4. Approximation/Instance Optimality.
5. Structural Alternatives.

Approximation/Instance Optimality.

Find a (close to) optimal t-spanner for a given input graph G or point set P. (Optimal per instance.)

Approximation/Instance Optimality.

Find a (close to) optimal t-spanner for a given input graph G or point set P. (Optimal per instance.)

Approximating graph t-spanners: $O(n^{1/k})$-approximation by greedy algorithm (basic results) for $t = 2k - 1$.

Approximation/Instance Optimality.

Find a (close to) optimal t-spanner for a given input graph G or point set P. (Optimal per instance.)

Approximating graph t-spanners: $O(n^{1/k})$-approximation by greedy algorithm (basic results) for $t = 2k - 1$.

- $O(\log n)$-approximation for 2-spanners. [Kortsarz,Peleg;94]
- $\tilde{O}(n^{1/4})$-appproximation for 3-spanners and 4-spanners. [Berman,Raskhodnikova,Ruan;10][Dinitz,Krauthgamer;11] [Berman,Bhattacharyya,Makarychev,Raskhodnikova,Yaroslavtsev;13] [Dinitz,Zhang;16]
- For every constant t and ϵ, $2^{\log^{1-\epsilon} n/t}$-approximation is hard.[Elkin,Peleg;00][Dinitz,Kortsarz,Raz;12]

Approximation/Instance Optimality.

Find a (close to) optimal t-spanner for a given input graph G or point set P. (Optimal per instance.)

Approximating graph t-spanners: $O(n^{1/k})$-approximation by greedy algorithm (basic results) for $t = 2k - 1$.

- $O(\log n)$-approximation for 2-spanners. [Kortsarz,Peleg;94]
- $\tilde{O}(n^{1/4})$-appproximation for 3-spanners and 4-spanners. [Berman,Raskhodnikova,Ruan;10][Dinitz,Krauthgamer;11] [Berman,Bhattacharyya,Makarychev,Raskhodnikova,Yaroslavtsev;13] [Dinitz,Zhang;16]
- For every constant t and ϵ, $2^{\log^{1-\epsilon} n/t}$-approximation is hard.[Elkin,Peleg;00][Dinitz,Kortsarz,Raz;12]

Open problem:

- $O(n^{\epsilon})$-approximation for t-spanners for a constant t?

Metric Spanners

- There have been a few NP-hardness results.
  [Klein,Kutz;06][Giannopoulos, Klein,Knauer,Kutz,Marx;10]
  [Carmi,Chaitman-Yerushalm;13]

- Approximation ratio of $O((1/\epsilon))^{d-1}$ from classical results. No other significant algorithmic result until now.

Metric Spanners

- There have been a few NP-hardness results. [Klein,Kutz;06][Giannopoulos, Klein,Knauer,Kutz,Marx;10] [Carmi,Chaitman-Yerushalm;13]
- Approximation ratio of $O((1/\epsilon))^{d-1}$ from classical results. No other significant algorithmic result until now.
- Bicriteria: $(1+O(\epsilon))$-spanner with $O(\log\log(d/\epsilon)|OPT|)$ where $|OPT|$ is the sparsity/lightness of optimal $(1+\epsilon)$ for the point set $P \in \mathbb{R}^d$ [L,Solomon,Than,Tóth,Zhang; upcoming].

Metric Spanners

- There have been a few NP-hardness results. [Klein,Kutz;06][Giannopoulos, Klein,Knauer,Kutz,Marx;10] [Carmi,Chaitman-Yerushalm;13]
- Approximation ratio of $O((1/\epsilon))^{d-1}$ from classical results. No other significant algorithmic result until now.
- Bicriteria: $(1+O(\epsilon))$-spanner with $O(\log\log(d/\epsilon)|OPT|)$ where $|OPT|$ is the sparsity/lightness of optimal $(1+\epsilon)$ for the point set $P \in \mathbb{R}^d$ [L,Solomon,Than,Tóth,Zhang; upcoming].

Open problems:

1. $(1+\epsilon)$-spanner with $O(1)|OPT|$?
2. t-spanner with $O(\log n)|OPT|$ for high dimensional $\mathbb{R}^d$ for any constant t?
3. Doubling metrics?

Outline

1. ~~Introduction.~~
2. ~~Basic Results.~~
3. ~~Fault Tolerance.~~
4. ~~Approximation/Instance Optimality.~~
5. Structural Alternatives.

Tree Covers

[Awerbuch,Peleg;92][Arya,Das,Mount,Salowe,Smid;95]

A t-tree cover $\mathcal{T}$ for a point set P is collection of trees such that $\forall x, y \in P$, there exists $T_{xy} \in \mathcal{T}$ where:

$$d_{T_{xy}}(x,y) \leq d_X(x,y) \leq t \cdot d_{T_{xy}}(x,y) \tag{2}$$

Want to construct a tree cover with a few trees.

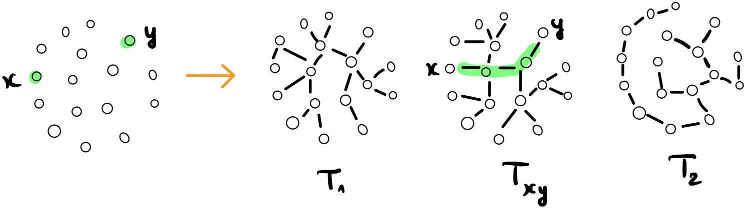

The union $S = \cup_{T \in \mathcal{T}} T$ is a t-spanner of P.

- More: $\mathcal{T}$ allows approximate distances to be queried quickly: just query distances on each tree.

Metric Tree Covers

- In $\mathbb{R}^d$, $(1+\epsilon)$-tree cover with $O_d((1/\epsilon)^{d-1})$ trees.
  [Arya,Das,Mount,Salowe,Smid;95]
  [Chang,Conroy,L,Milenkovic,Solomon,Than;24]

Metric Tree Covers

- In $\mathbb{R}^d$, $(1+\epsilon)$-tree cover with $O_d((1/\epsilon)^{d-1})$ trees.
  [Arya,Das,Mount,Salowe,Smid;95]
  [Chang,Conroy,L,Milenkovic,Solomon,Than;24]

- In doubling metrics of dimension d, the number of trees is $O_d((1/\epsilon)^d)$.[Bartal,Fandina,Neiman,19]

Metric Tree Covers

- In $\mathbb{R}^d$, $(1+\epsilon)$-tree cover with $O_d((1/\epsilon)^{d-1})$ trees.
  [Arya,Das,Mount,Salowe,Smid;95]
  [Chang,Conroy,L,Milenkovic,Solomon,Than;24]
- In doubling metrics of dimension d, the number of trees is $O_d((1/\epsilon)^d)$.[Bartal,Fandina,Neiman,19]

Open problem: Do light tree covers exist? (Every tree in the tree cover has total edge weight $O_{\epsilon,d}(1)\omega(MST)$.)

- [Filtser,Gitlitz,Neiman;23] showed that each tree could be $O_{\epsilon,d}(\log(n))\omega(MST)$.

Tree cover for general graphs

A stretch $2k-1$ tree cover with $n^{1/k}\log^{1-1/k}(n)$ trees.
[Awerbuch,Peleg,92][Awerbuch,Kutten,Peleg,94][Thorup-Zwick-05]

Tree cover for general graphs

A stretch $2k-1$ tree cover with $n^{1/k}\log^{1-1/k}(n)$ trees.
[Awerbuch,Peleg,92][Awerbuch,Kutten,Peleg,94][Thorup-Zwick-05]

[Bartal,Fandina,Neiman,19] asked: what stretch can we get with only 2 trees, 3 trees and so on.

- They showed 2 trees $\tilde{O}(\sqrt{n})$ stretch, k trees $\tilde{O}(n^{1/k})$ stretch.

No lower bounds known.

Tree cover for general graphs

A stretch $2k-1$ tree cover with $n^{1/k}\log^{1-1/k}(n)$ trees.
[Awerbuch,Peleg,92][Awerbuch,Kutten,Peleg,94][Thorup-Zwick-05]

[Bartal,Fandina,Neiman,19] asked: what stretch can we get with only 2 trees, 3 trees and so on.

- They showed 2 trees $\tilde{O}(\sqrt{n})$ stretch, k trees $\tilde{O}(n^{1/k})$ stretch.

No lower bounds known.

Open problem: what is the smallest stretch achieved with 2 trees?

A Path Cover?

A t-tree cover $\mathcal{T}$ where every tree in $\mathcal{T}$ is a path?

A Path Cover?

A t-tree cover $\mathcal{T}$ where every tree in $\mathcal{T}$ is a path?

- Impossible.

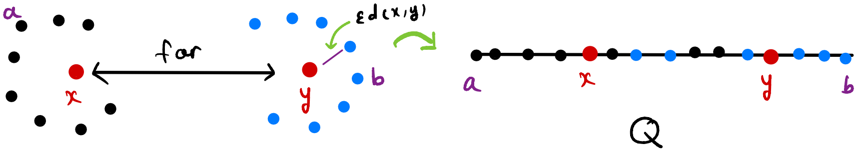

[Chan,Har-Peled,Jones;19]

Locality Sensitive Ordering (LSO)

A set of paths $\mathcal{P}$ such that for every $x \neq y$, there is a path $Q \in \mathcal{P}$ such that $Q[x,y] \subseteq B(x, \epsilon d_X(x,y)) \cup B(y, \epsilon d_X(x,y))$

LSO

Any $P \in \mathbb{R}^d$ or doubling metric of dim. d has an LSO with $O_{\epsilon,d}(1)$ paths, a.k.a., orderings. [Chan,Har-Peled,Jones;19][Filtser,L;22]

LSO

Any $P \in \mathbb{R}^d$ or doubling metric of dim. d has an LSO with $O_{\epsilon,d}(1)$ paths, a.k.a., orderings. [Chan,Har-Peled,Jones;19][Filtser,L;22]

$S = \cup_{Q \in \mathcal{P}} Q$ is a $(1+\epsilon)$-spanner for P.

LSO

Any $P \in \mathbb{R}^d$ or doubling metric of dim. d has an LSO with $O_{\epsilon,d}(1)$ paths, a.k.a., orderings. [Chan,Har-Peled,Jones;19][Filtser,L;22]

$S = \cup_{Q \in \mathcal{P}} Q$ is a $(1+\epsilon)$-spanner for P.

Many other applications, for example: Given $x \in P$, find a $(1+\epsilon)$-appproximate nearest neighbor of P:

- Return the closest point to x among all neighbors of x in paths in $\mathcal{P}$ (two neighbors per path $Q \in \mathcal{P}$.)
- The total running time $O(|\mathcal{P}|) = O_{\epsilon,d}(1)$.

LSO

Any $P \in \mathbb{R}^d$ or doubling metric of dim. d has an LSO with $O_{\epsilon,d}(1)$ paths, a.k.a., orderings. [Chan,Har-Peled,Jones;19][Filtser,L;22]

$S = \cup_{Q \in \mathcal{P}} Q$ is a $(1+\epsilon)$-spanner for P.

Many other applications, for example: Given $x \in P$, find a $(1+\epsilon)$-appproximate nearest neighbor of P:

- Return the closest point to x among all neighbors of x in paths in $\mathcal{P}$ (two neighbors per path $Q \in \mathcal{P}$.)
- The total running time $O(|\mathcal{P}|) = O_{\epsilon,d}(1)$.

Open problem: Other structural alternatives to spanners?

Other Things Not Mentioned Here

- Spanners in high dimensional spaces.
- Spanners in different models of computations.
- Ramsey tree covers.
- Etc.

Thanks!