

Reliable Spanners: Locality-Sensitive Orderings Strike Back

Hung Le
UMass Amherst



Arnold Filtser
Columbia University



Disclaimer

- ▶ (Lots of?) Buzzwords in this talk:
 - ▶ **Reliable Spanners, Oblivious Reliable Spanners,**
Hop-bounded Reliable Left Spanner
 - ▶ **Locality Sensitive Orderings,** Left-Sided Locality-Sensitive Orderings, Triangle Localitive Sensitive Orderings.
 - ▶ Sparse Covers, Ultrametric Covers.
 - ▶ Doubling Metrics, Planar Metrics, Minor-free Metrics.


- ▶ I try to give the context of these buzzwords and how they are related, glossing over technical details.

Spanners

A t -spanner of a **finite metric space** (X, d_X) , is a weighted graph $H(X, E, w)$ over X where for every pair of points $x, y \in X$:

$$d_X(x, y) \leq d_H(x, y) \leq t \cdot d_X(x, y) \quad (1)$$

d_H is the shortest path metric in H .




- ▶ Parameter t is called the stretch of the spanner.

Locality Sensitive Ordering (LSO) \rightarrow

Chen, Har-Peled, Jones
ITCS 19

(τ, ϵ) -LSO: Given a metric space (X, d_X) , a collection Σ of **linear orderings** is a (τ, ϵ) -LSO if:

- ▶ $|\Sigma| \leq \tau$.
- ▶ $\forall x \neq y \in X$, there is an ordering $\sigma \in \Sigma$ such that (w.l.o.g.) $x \prec_{\sigma} y$ and the points between x and y in σ could be partitioned into two consecutive intervals I_x, I_y where $I_x \subseteq B_X(x, \epsilon \cdot d_X(x, y))$ and $I_y \subseteq B_X(y, \epsilon \cdot d_X(x, y))$.



LSO for \mathbb{R}^d

think ϵ as a constant
 < 1

Theorem (CHJ09): Any n -point set in \mathbb{R}^d has a $O(\epsilon^{-d}, \epsilon)$ -LSO.


- ▶ Chan, Har-Peled, and Jones (CHJ09) provided many applications of LSO in solving: dynamic bichromatic closest pairs, dynamic spanners, dynamic approximate MST, static and dynamic fault-tolerant spanners.
- ▶ For doubling metrics, Chan, Har-Peled, and Jones (CHJ09) construct a $((\epsilon^{-1} \log n)^{O(d)}, \epsilon)$ -LSO.

Spanners from LSO

→ [CH519]

Lemma: If (X, d_X) admits a (τ, ϵ) -LSO, it has a $(1 + O(\epsilon))$ -spanner with $O(\tau n)$ edges.

↙ H



$$w_H(x, y) = d_X(x, y)$$

Spanners from LSO: The stretch

$$\begin{aligned}d_H(x, y) &\leq (1 + 2\epsilon)d_X(x, y) + \epsilon(d_H(x, p) + d_H(y, q)) \\ &= (1 + O(\epsilon))d_X(x, y)\end{aligned}$$

- ▶ Stretch guarantee crucially relies on that $\epsilon < 1$.

Deterministic Reliable Spanner

A weighted graph $H(X, E, w)$ is a (deterministic) ϑ -reliable t -spanner of a metric space (X, d_X) if for every set $B \subseteq X$ of points, there is a set $B^+ \supseteq B$, called an *extension* of B , such that:

1. $|B^+| \leq (1 + \vartheta)|B|$.
2. For every $x, y \notin B^+$, $d_X(x, y) \leq d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$.

Deterministic Reliable Spanner

A weighted graph $H(X, E, w)$ is a (deterministic) ϑ -reliable t -spanner of a metric space (X, d_X) if for every set $B \subseteq X$ of points, there is a set $B^+ \supseteq B$, called an *extension* of B , such that:

1. $|B^+| \leq (1 + \vartheta)|B|$.
 2. For every $x, y \notin B^+$, $d_X(x, y) \leq d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$.
- ▶ ϑ is a (small) constant. The size of B could be $\Omega(n)$, and we don't need to know $|B|$ in advance.
- ▶ f -fault tolerant spanners: need to know f in advance, and have $\Omega(n^2)$ edges when $f = \Omega(n)$.

Oblivious Reliable Spanner

An oblivious ϑ -reliable t -spanner is a distribution \mathcal{D} over t -spanners $H(X, E, w)$ of (X, d_X) , such that for every set $B \subseteq X$, there exists B^+ of B such that:

1. $\mathbb{E}_{H \sim \mathcal{D}} [|B^+|] \leq (1 + \vartheta) |B|$
2. For every $x, y \notin B^+$, $d_X(x, y) \leq d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$.

Oblivious Reliable Spanner

An oblivious ϑ -reliable t -spanner is a distribution \mathcal{D} over t -spanners $H(X, E, w)$ of (X, d_X) , such that for every set $B \subseteq X$, there exists B^+ of B such that:

1. $\mathbb{E}_{H \sim \mathcal{D}} [|B^+|] \leq (1 + \vartheta)|B|$
 2. For every $x, y \notin B^+$, $d_X(x, y) \leq d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$.
- There exists $P \in \mathbb{R}^1$ such that any **deterministic** reliable t -spanner must have $\Omega(n \log n)$ edges, for any constant t , while it is possible to construct an **oblivious** ϑ -reliable $(1 + \epsilon)$ -spanner with $O(n)$ edges.

Reliable Spanners from LSO

If (X, d_X) admits a (τ, ϵ) -LSO, it has a:

- ▶ **deterministic** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $\tilde{O}(n\tau^7 \log n)$ edges (for constant ϑ, ϵ) [BHO19].
- ▶ **oblivious** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $O(n(\log \log n)\tau^2)$ edges (for constant ϑ, ϵ) [BHO20].

Other Metrics Spaces?

- ▶ Doubling metrics of dimension d : ϑ -reliable $(1 + \epsilon)$ -spanners with $O(n(\log n)^{O(d)})$ edges [BHO19].
- ▶ Tree metrics: oblivious ϑ -reliable $(3 + \epsilon)$ -spanners with $O(n \log n (\log \Delta)^2)$ edges [HO20].
- ▶ Planar metrics: oblivious ϑ -reliable $(3 + \epsilon)$ -spanners with $O(n \log^2 n (\log \Delta)^2)$ edges [HO20].
- ▶ General metrics: oblivious ϑ -reliable $O(k)$ -spanners with $O(n^{1+1/k} \log n (\log \Delta)^2)$ edges [HO20].

Δ is the spread of the metric: $\Delta = \frac{\max_{x,y} d_X(x,y)}{\min_{x \neq y} d_X(x,y)}$.

- ▶ Δ could be $2^{\Omega(n)}$, or worst, and the dependency on $\log(\Delta)$ is inherent in the technique of Har-Peled and Oláh [HO20].

Sparse Cover

A (τ, ρ) -cover is a collection \mathcal{C} of clusters such that:

- ▶ every point belongs to at most τ clusters.
- ▶ for every pair $x, y \in X$ there is a cluster $C \in \mathcal{C}$ such that $x, y \in C$ and $\frac{1}{\rho} \cdot \text{diam}(C) \leq d_X(x, y) \leq \text{diam}(C)$.


Theorem: (τ, ρ) -cover implies oblivious ϑ -reliable 2ρ -spanner with $O(n\tau^2 \log \tau)$ edges.

- ▶ Tree metrics and planar metrics have $(\log(\Delta)\epsilon^{-1}, 2 + \epsilon)$ -cover.
- ▶ General metrics have $(\log(\Delta)n^{1/k}, O(k))$ -cover.

Our Results

Family	stretch	guarantee	size	ref
Doubling dim. d	$1 + \epsilon$	Determi.	$n\tilde{O}(\log n)$	This paper
	$1 + \epsilon$	Oblivious	$n\tilde{O}(\log \log n)^2$	This paper
General metric	$512 \cdot k$	Oblivious	$O(n^{1+\frac{1}{k}} k \log(n) \log^2 \Delta)$	[HO20]
	$8k - 2$	Oblivious	$O(n^{1+\frac{1}{k}} k^2 \tilde{O}(\log n))$	This paper
Tree	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\log^2 n \log^2 \Delta)$	[HO20]
	2	Oblivious	$n \cdot O(\log^3 n)$	This paper
	$t < 2$	Oblivious	$\Omega(n^2)$	This paper
Planar	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\log^2 n \log^2 \Delta)$	[HO20]
	$2 + \epsilon$	Oblivious	$n \cdot O(\log^5 n)$	This paper
Minor-free	$2 + \epsilon$	Oblivious	$n \cdot O(\log^5 n)$	This paper

Our Techniques: LSO (strikes back)



The Rest of the Talk

- ▶ Reliable Spanners for Doubling Metrics.
- ▶ Reliable Spanners for Tree Metrics.

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 $R/2$.

- ▶ n -point set in $(\mathbb{R}^d, \|\cdot\|_2)$ has doubling dimension $\Theta(d)$.

Recall: Reliable Spanners from LSO

If (X, d_X) admits a (τ, ϵ) -LSO, it has a:

- ▶ **deterministic** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $\tilde{O}(n\tau^7 \log n)$ edges (for constant ϑ, ϵ) [BHO19].
- ▶ **oblivious** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $O(n(\log \log n)\tau^2)$ edges (for constant ϑ, ϵ) [BHO20].

Doubling Metrics Admit LSO

Theorem (FL20): For every $\epsilon \in (0, \frac{1}{6})$, every metric space (X, d_X) of doubling dimension d admits an $(\epsilon^{-O(d)}, \epsilon)$ -LSO.

- ▶ Previous result [CHJ19]: (X, d_X) of doubling dimension d admits an $((\log n \epsilon^{-1})^d, \epsilon)$ -LSO.

Corollary: If (X, d_X) has doubling dimension d , it has

- ▶ **deterministic** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $\tilde{O}(n \log n)$ edges (for constant ϑ, ϵ)
- ▶ **oblivious** ϑ -reliable $(1 + O(\epsilon))$ -spanner with $O(n(\log \log n)^2)$ edges (for constant ϑ, ϵ)

Ultrametrics

- ▶ An ultrametric is a metric with a strengthened triangle inequality:

$$d(x, z) \leq \max\{d(x, y), d(y, z)\} \forall x, y, z \in X$$

- ▶ An ultrametric can be represented by a k -HST:

k -HST: A metric (X, d_X) is a k -hierarchical well-separated tree if there exists a bijection φ from X to leaves of a rooted tree T :

1. Each node $v \in T$ has a label Γ_v s.t $\Gamma_v = 0$ if v is a leaf and $\Gamma_v \geq k\Gamma_u$ if u is any child of v .
2. $d_X(x, y) = \Gamma_{\text{lca}(\varphi(x), \varphi(y))}$.

LSO for Ultrametrics

- ▶ A (δ, k) -ultrametric is a k -HST where every vertex of the tree T has degree at most δ .

Lemma (FL20): Every (δ, k) -ultrametric (U, d_U) of degree δ admits a $(\lceil \frac{\delta}{2} \rceil, \frac{1}{k})$ -LSO.

Ultrametric Cover

Theorem (FL20): For any metric (X, d_X) of doubling dimension d , we can construct a set $\mathcal{U} = \{U_1, U_2, \dots\}$ such that:

- ▶ $|\mathcal{U}| = \epsilon^{-O(d)}$.
- ▶ Each U_i is an $(\epsilon^{-O(d)}, \frac{1}{\epsilon})$ -ultrametrics.
- ▶ For every $x \neq y \in X$, there exists $U_i \in \mathcal{U}$ such that:

$$d_X(x, y) \leq d_{U_i}(x, y) \leq (1 + \epsilon)d_X(x, y)$$

Putting Things Together

Theorem (FL20): For every $\epsilon \in (0, \frac{1}{6})$, every metric space (X, d_X) of doubling dimension d admits an $(\epsilon^{-O(d)}, \epsilon)$ -LSO.

The Rest of the Talk

- ▶ Reliable Spanners for Doubling Metrics.
- ▶ Reliable Spanners for Tree Metrics.

Tree Metrics

- ▶ A tree metric (X, d_X) is the shortest path metric of a tree.

Lemma (FL20): Any (oblivious/deterministic) ϑ -reliable t -spanner for $t < 2$ of a star must have $\Omega(n^2)$ edges.

Tree Metrics

- ▶ A tree metric (X, d_X) is the shortest path metric of a tree.

Lemma (FL20): Any (oblivious/deterministic) ϑ -reliable t -spanner for $t < 2$ of a star must have $\Omega(n^2)$ edges. * Deterministic ϑ -reliable


$(2t)$ -spanner for $t \geq 1$ of a star must have $\Omega(n^{1+\frac{1}{t}})$ edges [HO20].

- ▶ There exists an oblivious ϑ -reliable $(3 + \epsilon)$ -spanner with $O(n \log^2 n \log^2(\Delta))$ edges.

Question: Can we construct an oblivious ϑ -reliable 2-spanner with $O(n \text{polylog}(n))$ edges?

Tree Metrics: Our Results

Theorem (FL20): Tree metrics have an oblivious ϑ -reliable 2-spanner with $O(n \log^3 n)$ edges.



Lemma (FL20): Any (oblivious/deterministic) ϑ -reliable t -spanner for $t < 2$ of a star must have $\Omega(n^2)$ edges.

New Tools (1): Left-sided Locality-Sensitive Ordering

(τ, ρ) -left-sided LSO A collection Σ of *partial orderings* of (X, d_X) is a (τ, ρ) -left-sided LSO if:

- ▶ every point $x \in X$ belongs to at most τ orderings in Σ .
- ▶ $\forall x, y \in X$, there is an order $\sigma \in \Sigma$ such that for every $x' \preceq_\sigma x$ and $y' \preceq_\sigma y$ it holds that $d_X(x', y') \leq \rho \cdot d_X(x, y)$.

Left-Sided LSO for Tree

Lemma (FL20): n -point tree metric (X, d_X) has a $(\log n, 1)$ -Left-Sided LSO.

Construction:

- ▶ v is a center of a tree T if every connected component of $T \setminus \{v\}$ has at most $n/2$ vertices.

- ▶ There always exists a center vertex for a given tree T [Jordan, 1869].

Lemma (FL20): n -point tree metric (X, d_X) has a $(\log n, 1)$ -Left-Sided LSO.

$\forall x, y \in X$, there is an order $\sigma \in \Sigma$ such that for every $x' \preceq_\sigma x$ and $y' \preceq_\sigma y$ it holds that $d_X(x', y') \leq \rho \cdot d_X(x, y)$.

New Tools (2): 2-Hop Reliable Left Spanner

Given a path graph P_n , a 2-hop left spanner H is a graph such that for every $a < b$ there is $c \leq a, b$, such that $\{c, a\}, \{c, b\} \in H$.

New Tools (2): 2-Hop Reliable Left Spanner

Given a path graph P_n , a 2-hop left spanner H is a graph such that for every $a < b$ there is $c \leq a, b$, such that $\{c, a\}, \{c, b\} \in H$.

A distribution \mathcal{D} over 2-hop left spanners H is ϑ -reliable if for every attack $B \subseteq X$, there is a set $B^+ \subseteq X$ such that $\mathbb{E}[|B^+|] \leq (1 + \vartheta)|B|$, and for every $a < b$ such that $a, b \notin B^+$, there is $c \leq a, b$, $c \notin B$ such that $\{c, a\}, \{c, b\} \in E(H)$.

2-Hop Reliable Left Spanners for Path Graphs

Lemma (FL20): For every $\vartheta \in (0, 1)$, the n -vertex path P_n admits an oblivious ϑ -reliable 2-hop left spanner with $n \cdot O(\vartheta^{-1} \log n)$ edges.

- ▶ Sample $a \in [n]$ with probability $\frac{c}{a\vartheta}$.
 - ▶ N be the set of sampled vertices.
 - ▶ If a is sampled, add edge from a to every vertex in $\{a + 1, a + 2, \dots, n\}$.

Putting Things Together

Theorem (FL20): Tree metrics have an oblivious ϑ -reliable 2-spanner with $O(n \log^3 n)$ edges.

Conclusion

- ▶ Reliable Spanners from Locality-Sensitive Orderings.
- ▶ Open Problems:
 - ▶ Can we construct an oblivious reliable 2-spanner for trees with $O(n)$ edges?
 - ▶ Can we construct (deterministic or oblivious) reliable spanners for **graphs** where the spanner must be the subgraph of the input graph?
 - ▶ Several other open problems in our paper, check it out.

