Euclidean Steiner Spanners: Light and Sparse

Sujoy Bhore

Indian Institute of Technology Bombay

RTA, 2023

⋆Joint work with Csaba Tóth

Game of Missing Links! ...

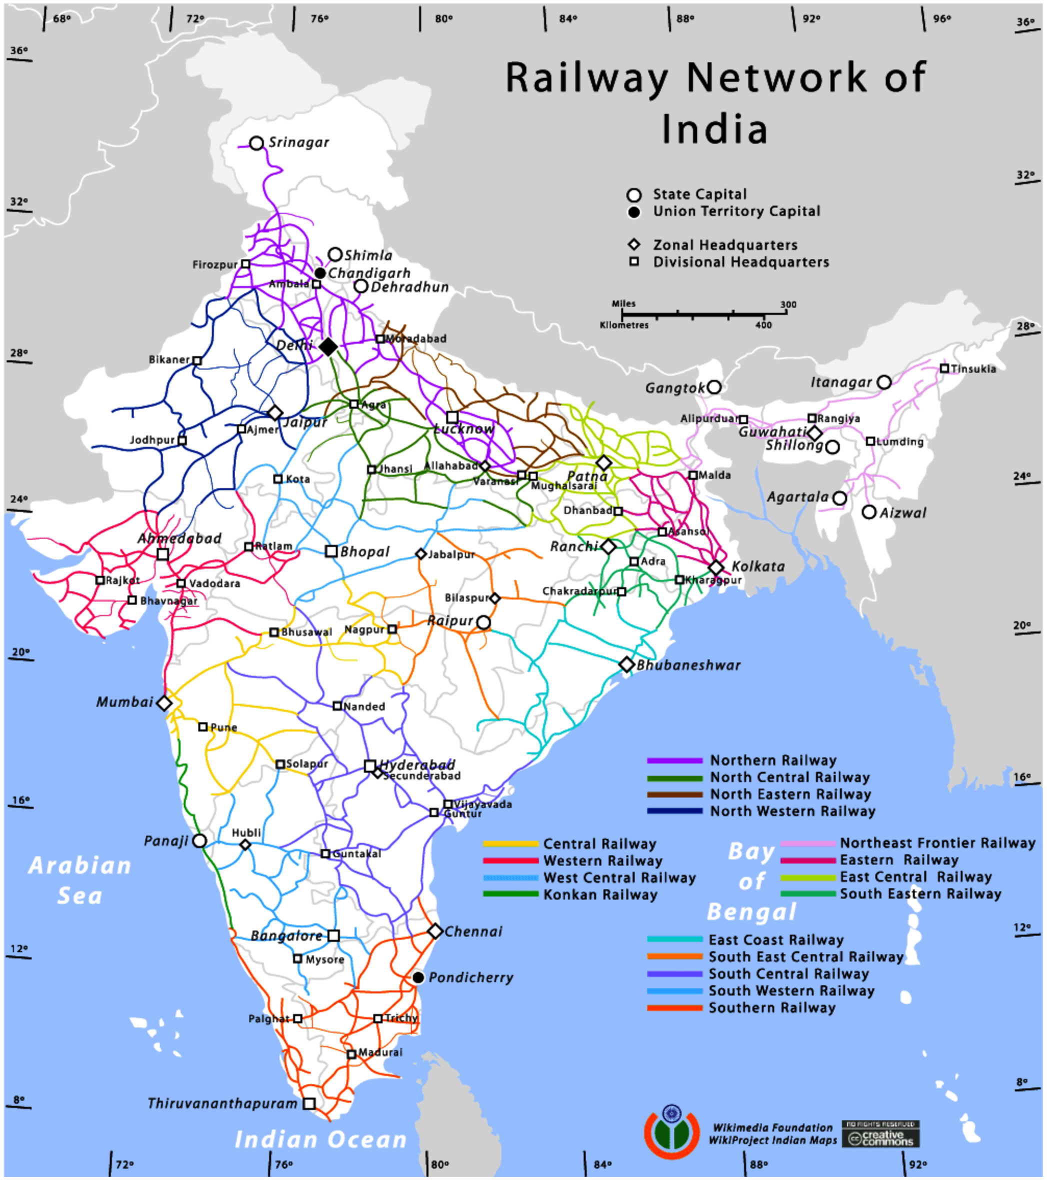

(Source.- Wiki.)

Game of Missing Links! ...

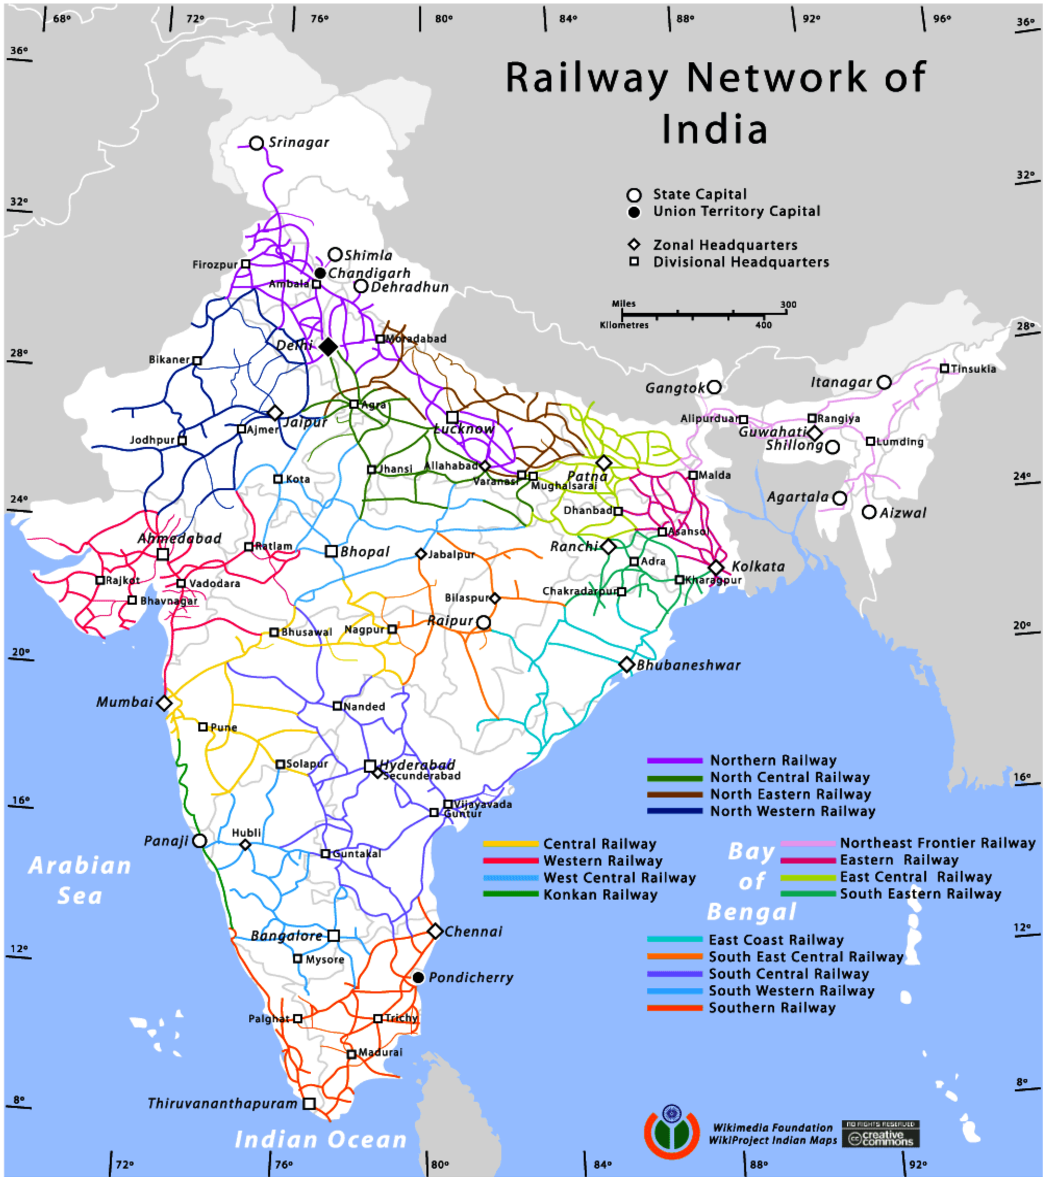

(Source.- Wiki.)

Q. Can I go from Mumbai to Kolkata... ?

Game of Missing Links! ...

(Source.- Wiki.)

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

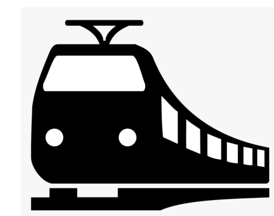

Game of Missing Links! ...

(Source.- Wiki.)

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

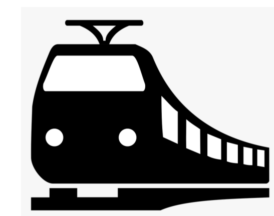

There is a problem in the Nagpur - Raipur track.

Game of Missing Links! ...

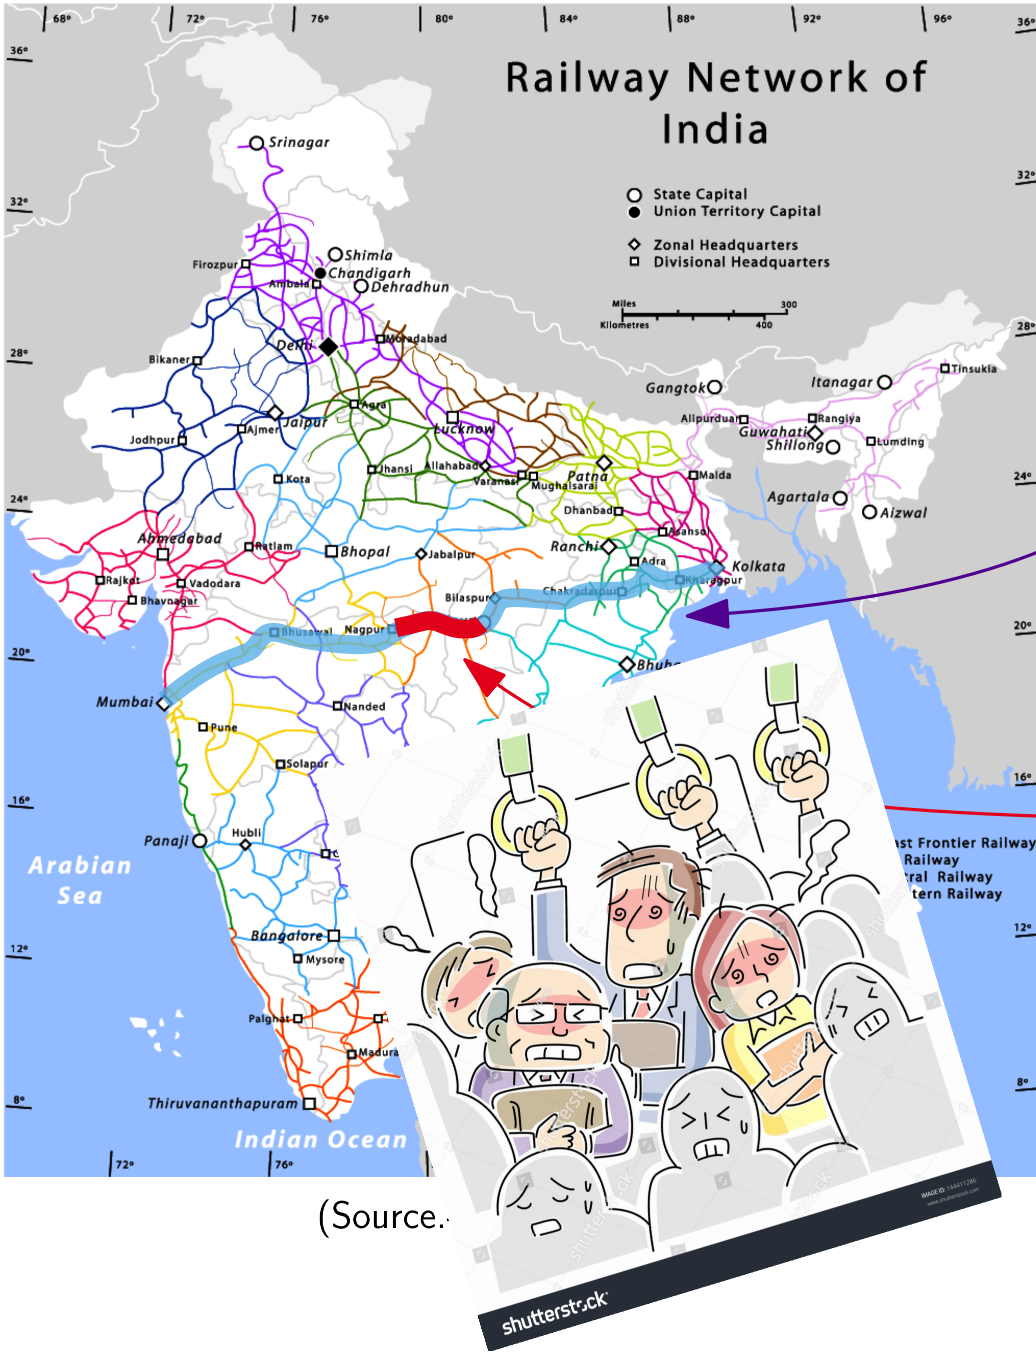

(Source.

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

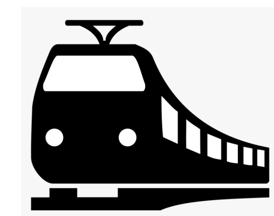

There is a problem in the Nagpur - Raipur track.

Game of Missing Links! ...

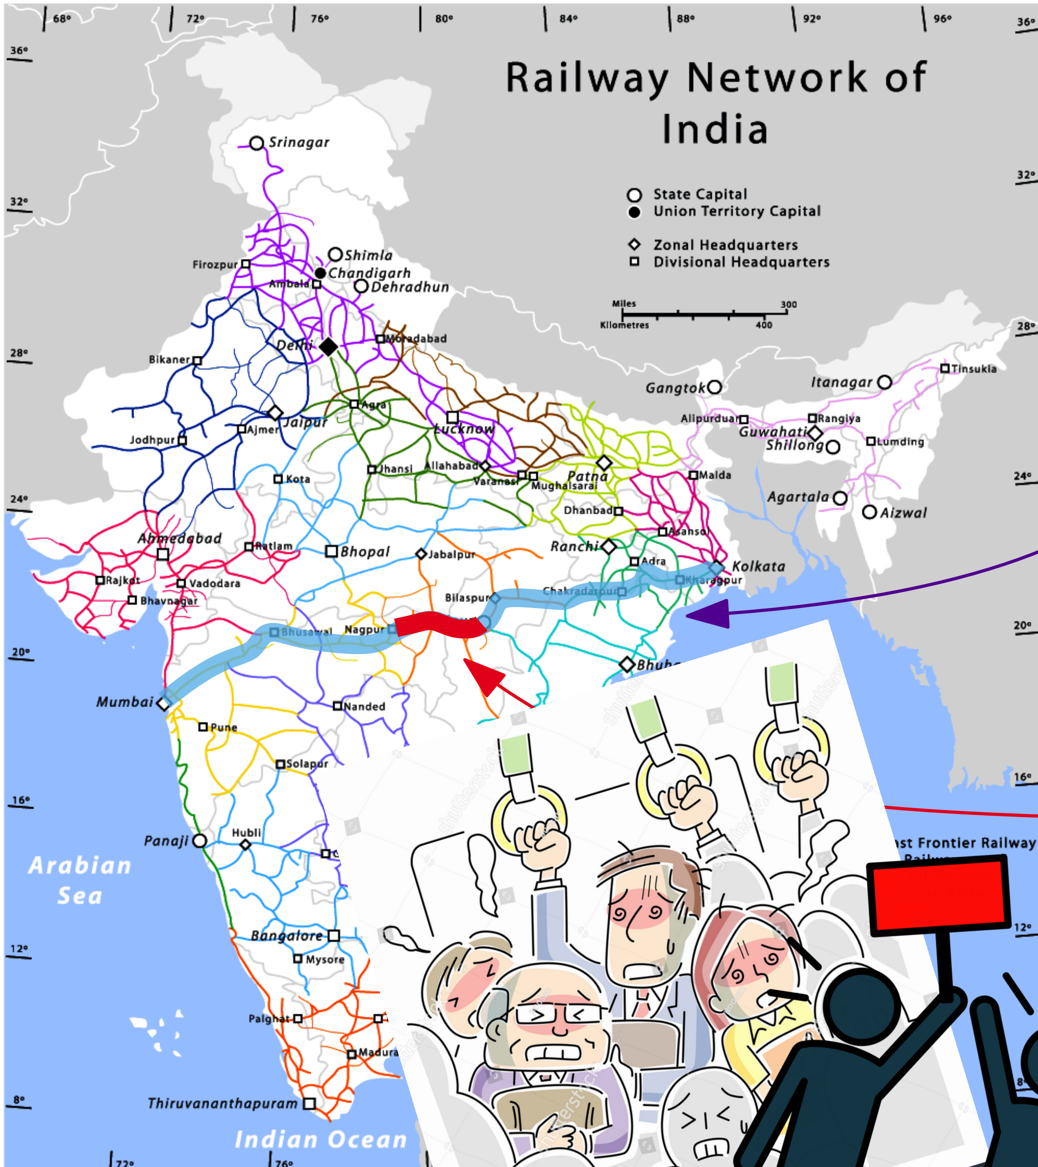

(Source.

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

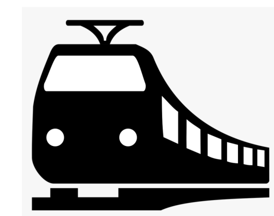

There is a problem in the Nagpur - Raipur track.

Game of Missing Links! ...

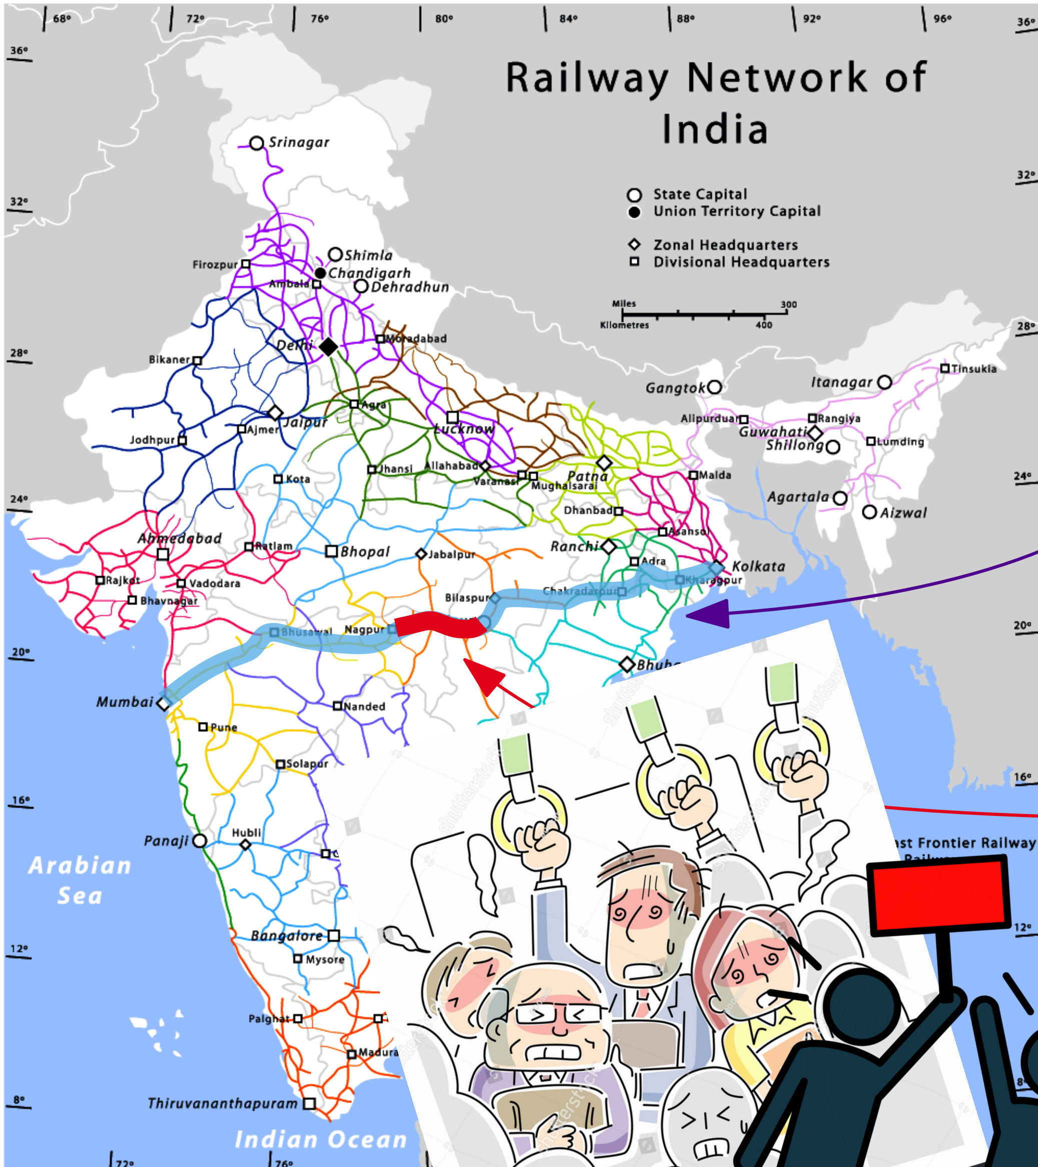

(Source.

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

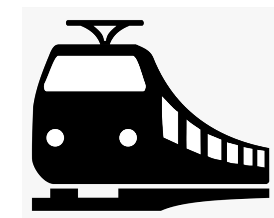

There is a problem in the Nagpur - Raipur track.

Game of Missing Links! ...

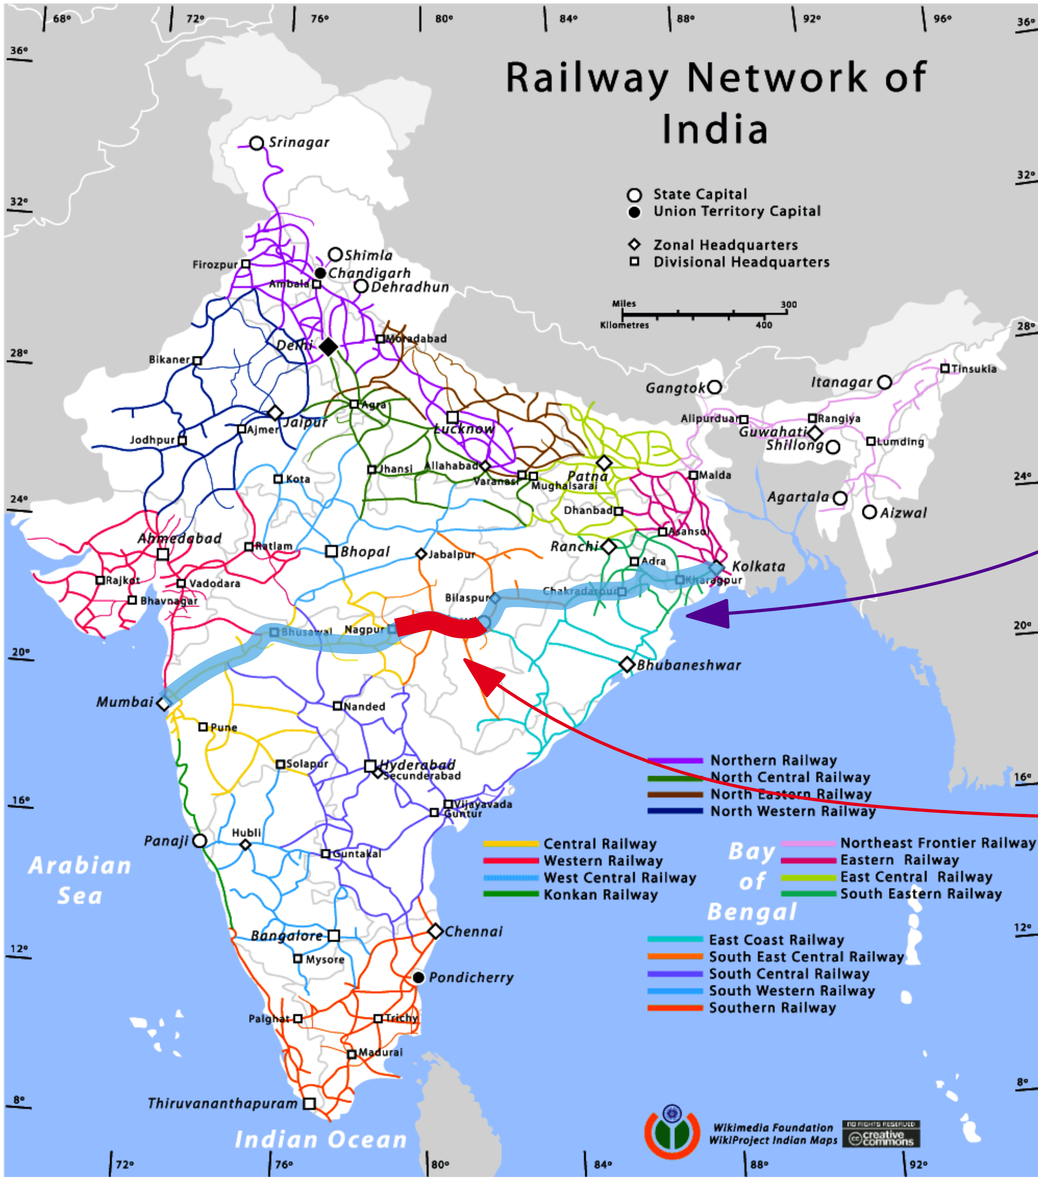

(Source.- Wiki.)

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

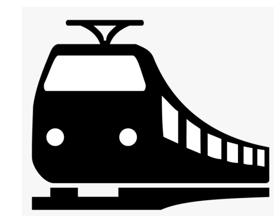

There is a problem in the Nagpur - Raipur track.

Soln.: Build robust networks that don't fail often!

Game of Missing Links! ...

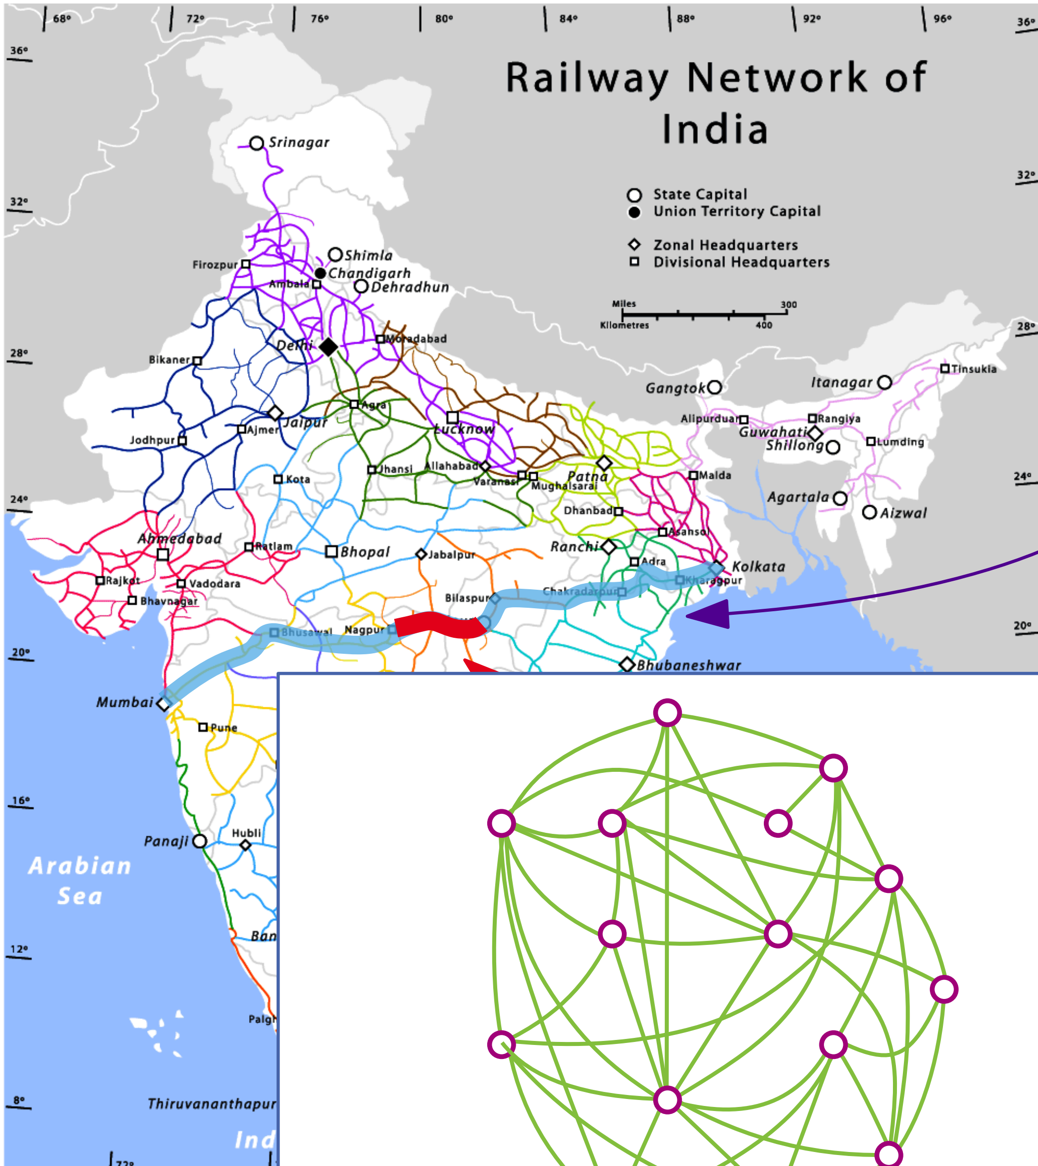

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

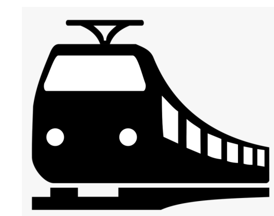

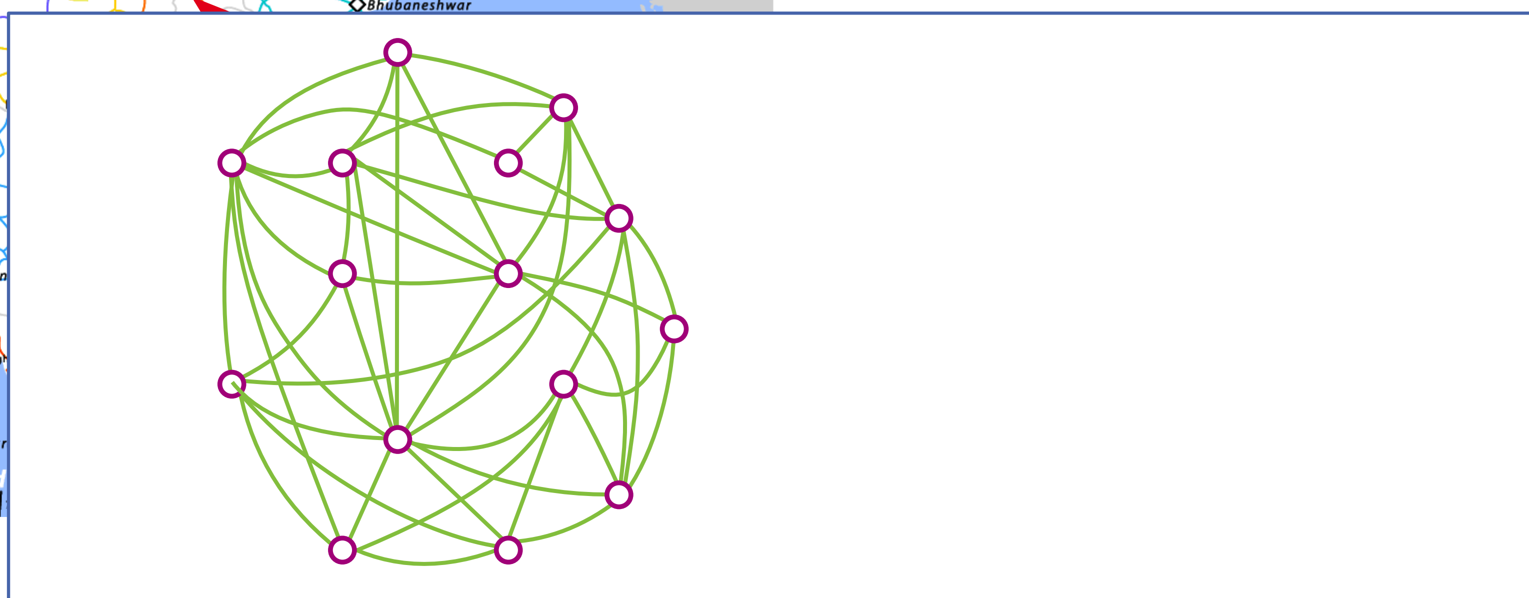

in the

track.

s that

Game of Missing Links! ...

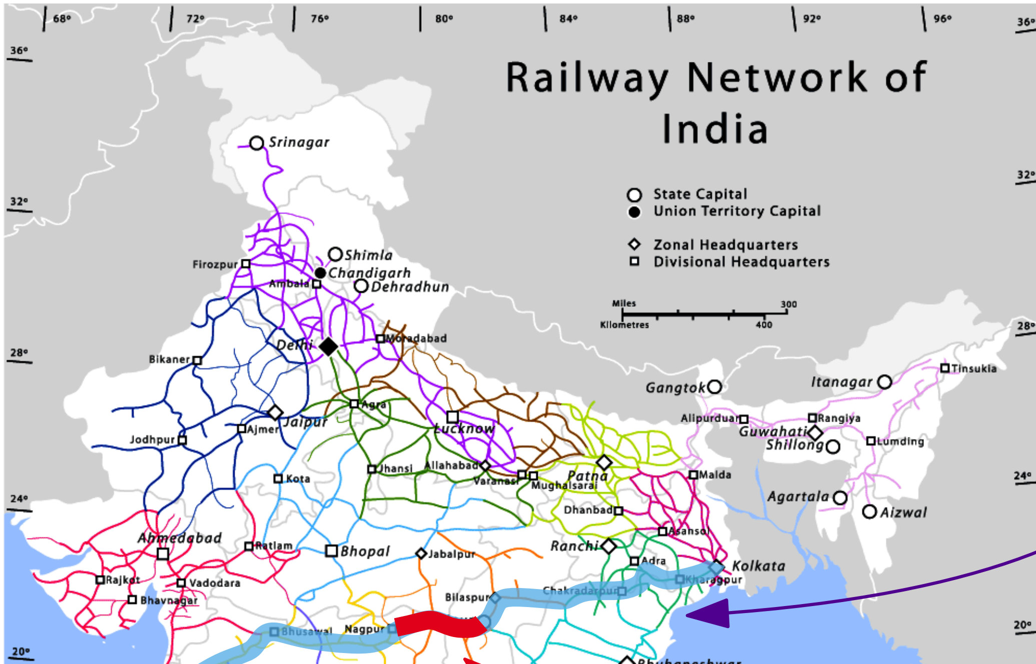

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

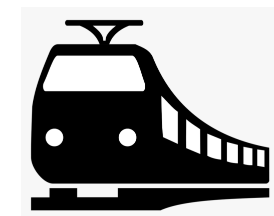

in the

track.

s that

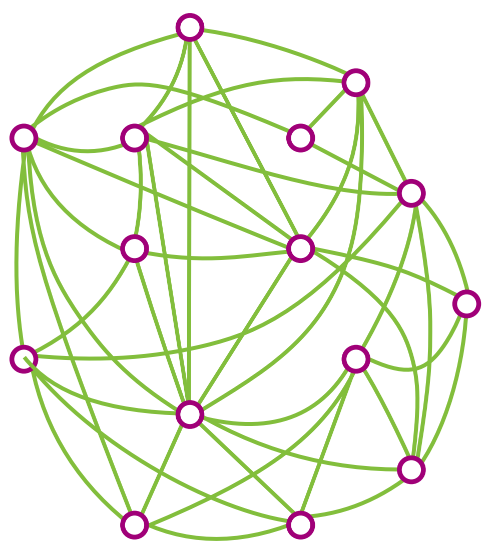

We need sparsification -

- Less connections.
- Less costly.
- Less congestion.
- Reduced travel time.

Game of Missing Links! ...

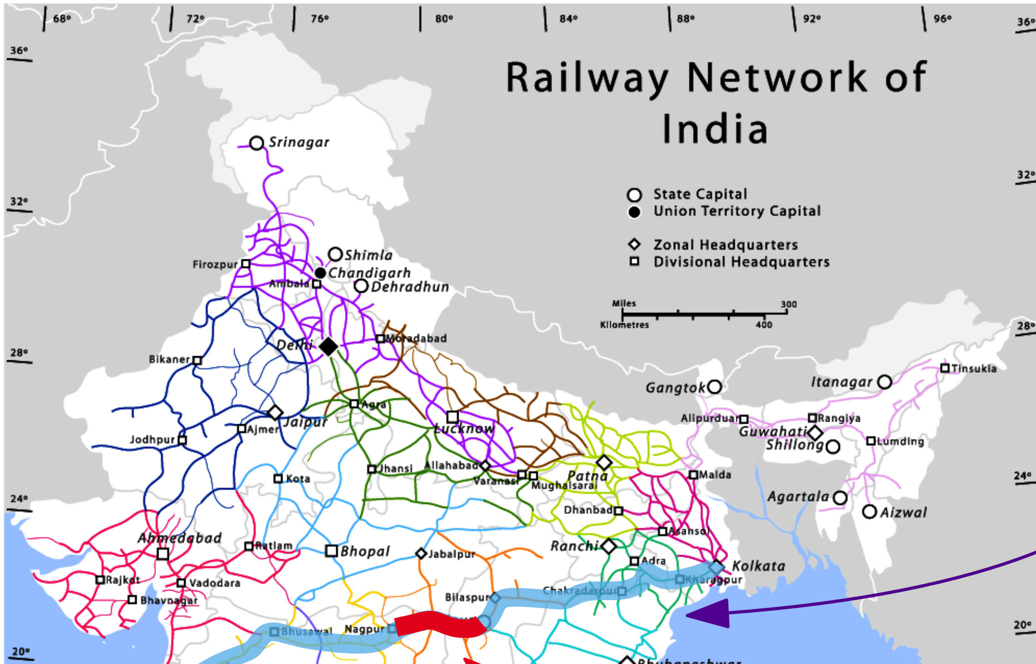

Q. Can I go from Mumbai to Kolkata... ?

Of course! there you go ...

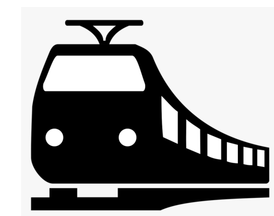

in the

track.

s that

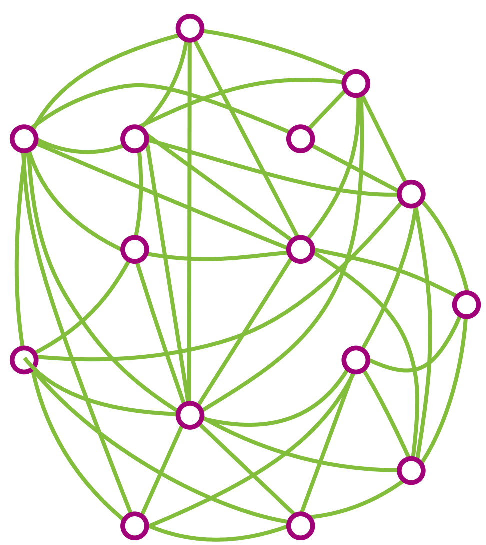

- Fault-tolerant.
- Robust.
- Steiner Spanner!

Graph Spanners

- Given an edge-weighted graph G, a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion.

Graph Spanners

- Given an edge-weighted graph G, a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion.

- Formally, a subgraph H of a given edge-weighted graph G is a t-*spanner*, for some $t \geq 1$, if for every $pq \in \binom{V(G)}{2}$ we have $d_H(p,q) \leq t \cdot d_G(p,q)$, where $d_G(p,q)$ denotes the length of the shortest path in G.

Graph Spanners

- Given an edge-weighted graph G, a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion.

- Formally, a subgraph H of a given edge-weighted graph G is a t-*spanner*, for some $t \geq 1$, if for every $pq \in \binom{V(G)}{2}$ we have $d_H(p,q) \leq t \cdot d_G(p,q)$, where $d_G(p,q)$ denotes the length of the shortest path in G.

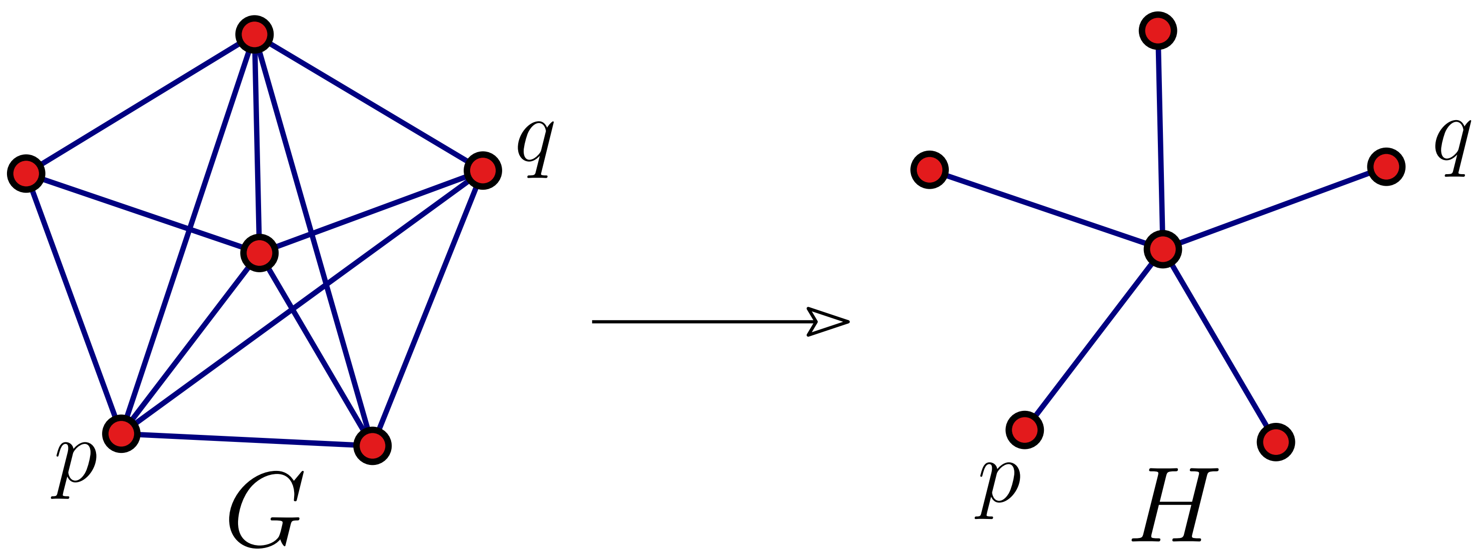

Graph Spanners

- Given an edge-weighted graph G, a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion.

- Formally, a subgraph H of a given edge-weighted graph G is a t-*spanner*, for some $t \geq 1$, if for every $pq \in \binom{V(G)}{2}$ we have $d_H(p,q) \leq t \cdot d_G(p,q)$, where $d_G(p,q)$ denotes the length of the shortest path in G.

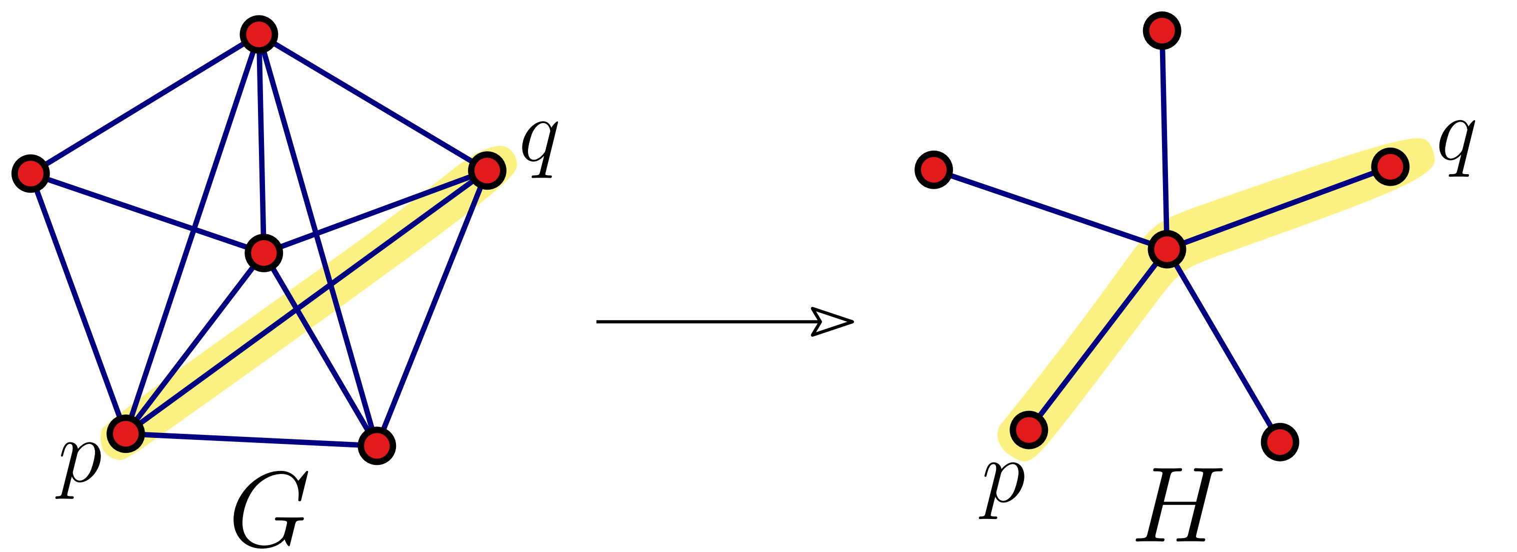

Graph Spanners

- Given an edge-weighted graph G, a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion.

- Formally, a subgraph H of a given edge-weighted graph G is a t-*spanner*, for some $t \geq 1$, if for every $pq \in \binom{V(G)}{2}$ we have $d_H(p,q) \leq t \cdot d_G(p,q)$, where $d_G(p,q)$ denotes the length of the shortest path in G.

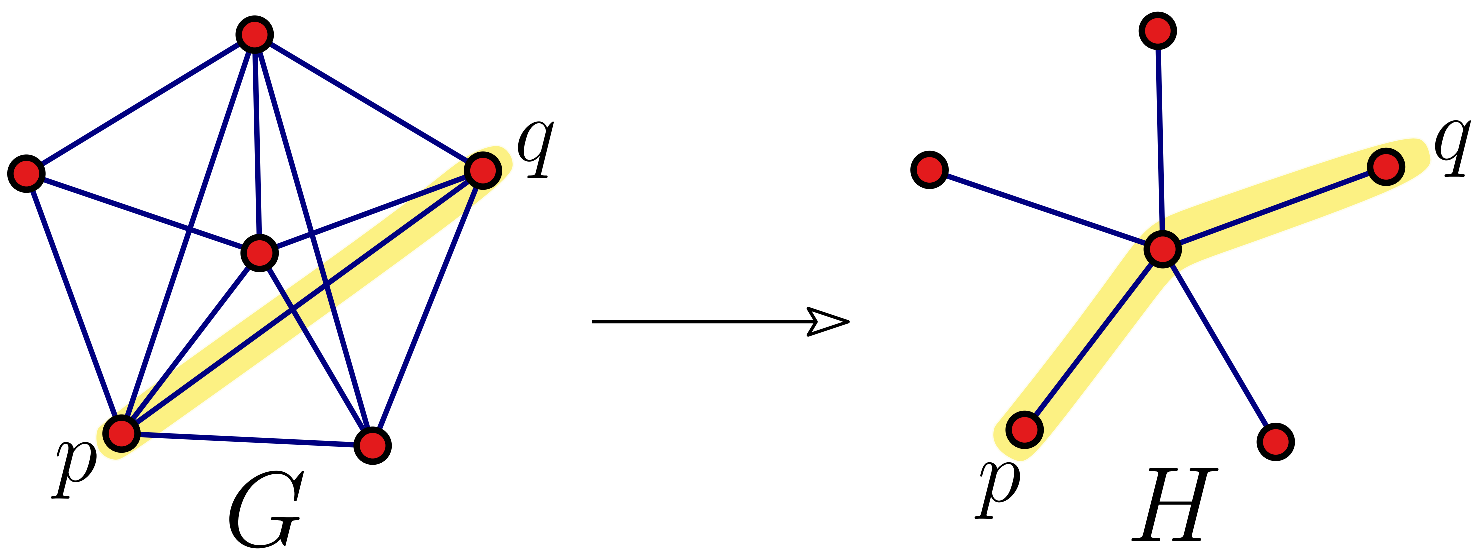

- The parameter t is called the stretch factor/dialation factor of the spanner.

Graph Spanners - Applications!

- Graph spanners were introduced by Peleg and Schäffer [Journal of Graph Theory'89].
- Spanners are fundamental graph structures with numerous applications –
  - Distributed queuing protocol [Demmer & Herlihy, DISC'98].
  - Compact routing scheme [Thorup & Zwick, SPAA'01].
  - Online load balancing [Awerbuch et al., STOC'92].
  - Wireless sensor networks [Shpungin & Segal, INFOCOM'09].
  - Motion planning in robotics control optimization [Cai & Keil, IJCGA'97].
  - Distributed systems and communications [Peleg, SIAM M. DMA'00; Demmer & Herlihy, DISC'98].
  - and many others ...

Geometric Spanners

- In geometric settings, a t-spanner for a finite point set P of points in $\mathbb{R}^d$, is a subgraph underlying of the complete graph $G_P = (P, \binom{P}{2})$ that preserves the pairwise Euclidean distances between points in S to within a factor of t, that is the *stretch factor*.
- The edge weights of G_P are the Euclidean distances between the vertices.

Geometric Spanners

- In geometric settings, a t-spanner for a finite point set P of points in $\mathbb{R}^d$, is a subgraph underlying of the complete graph $G_P = (P, \binom{P}{2})$ that preserves the pairwise Euclidean distances between points in S to within a factor of t, that is the *stretch factor*.
- The edge weights of G_P are the Euclidean distances between the vertices.

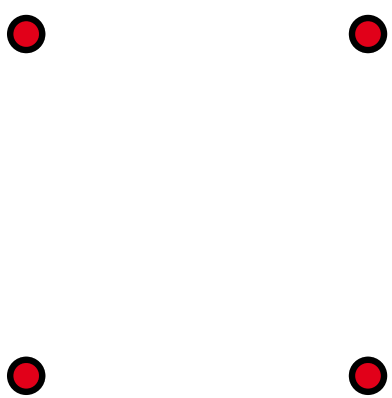

P

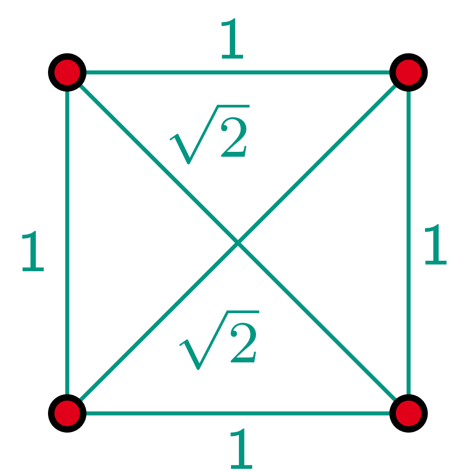

G_P

Geometric Spanners

- In geometric settings, a t-spanner for a finite point set P of points in $\mathbb{R}^d$, is a subgraph underlying of the complete graph $G_P = (P, \binom{P}{2})$ that preserves the pairwise Euclidean distances between points in S to within a factor of t, that is the *stretch factor*.
- The edge weights of G_P are the Euclidean distances between the vertices.

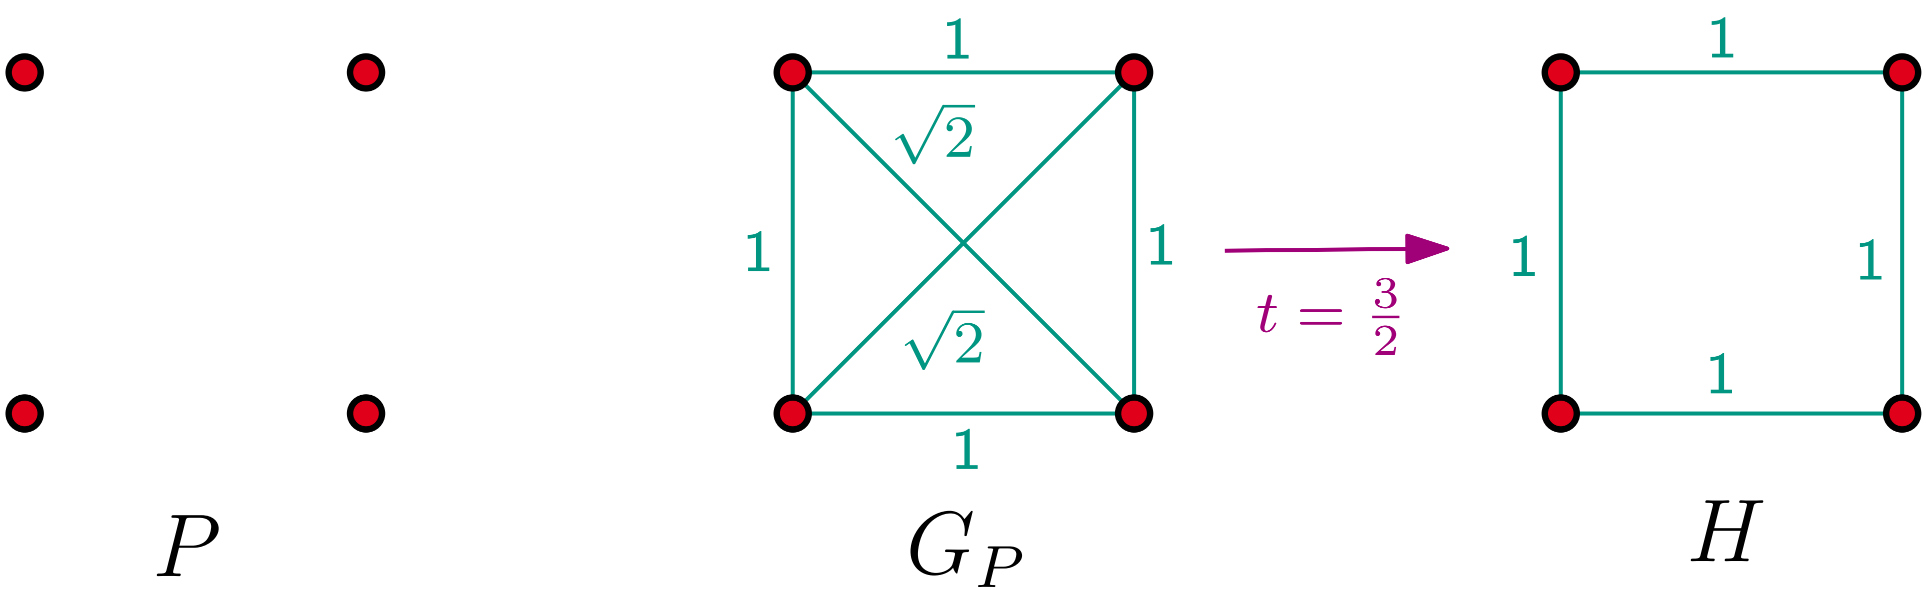

Geometric Spanners - Applications!

- Chew [SoCG'1986] initiated the study of Euclidean spanners, and showed that for a set of n points in $\mathbb{R}^2$, there exists a spanner with $O(n)$ edges and constant stretch factor.
- Geometric spanners have applications accross domains –
  - Topology control in wireless networks [Schindelhauer et al., Comp. Geom.'07].
  - Efficient regression in metric spaces [Gottlieb et al., IEEE T. Inf. Th.'17].
  - Approximate distance oracles [Gudmundsson et al. TALG'08].
  - Euclidean spanners are relevant in the context of other fundamental NP-hard problems, such as Euclidean TSP, Euclidean minimum Steiner tree [Rao and Smith, STOC'1998].

Various Types of Geometric Spanner Constructions

- Bounded-degree spanners [Bose et al., Algorithmica'05]
- α-diamond spanners [Das & Joseph, ISOA'98].
- Well-separated pair decomposition (WSPD) [Callahan, FOCS'93; Gudmundsson et al., SIAM J. Comp.'02]
- Skip-lists [Arya et al., FOCS'94].
- Path-greedy [Althöfer et al., DCG'93].
- Gap-greedy [Arya & Smid, Algorithmica'97].
- Locality sensitive orderings [Chan et al., SIAM J. Comp.'20].
- See the book of Narasimhan and Smid on geometric spanners, and the survey of Bose et al.

Sparsity

- The sparsity of a spanner H is the ratio

$$\frac{|E(H)|}{|E(MST)|} \approx \frac{|E(H)|}{|V(G)|}$$

between the number of edges of H and an MST.

- Since H is connected, $\liminf_{|V(G)|\to\infty}$ sparsity(H) ≥ 1.

Sparsity

- The sparsity of a spanner H is the ratio
$$\frac{|E(H)|}{|E(MST)|} \approx \frac{|E(H)|}{|V(G)|}$$
between the number of edges of H and an MST.
- Since H is connected, $\liminf_{|V(G)|\to\infty}$ sparsity(H)≥ 1.

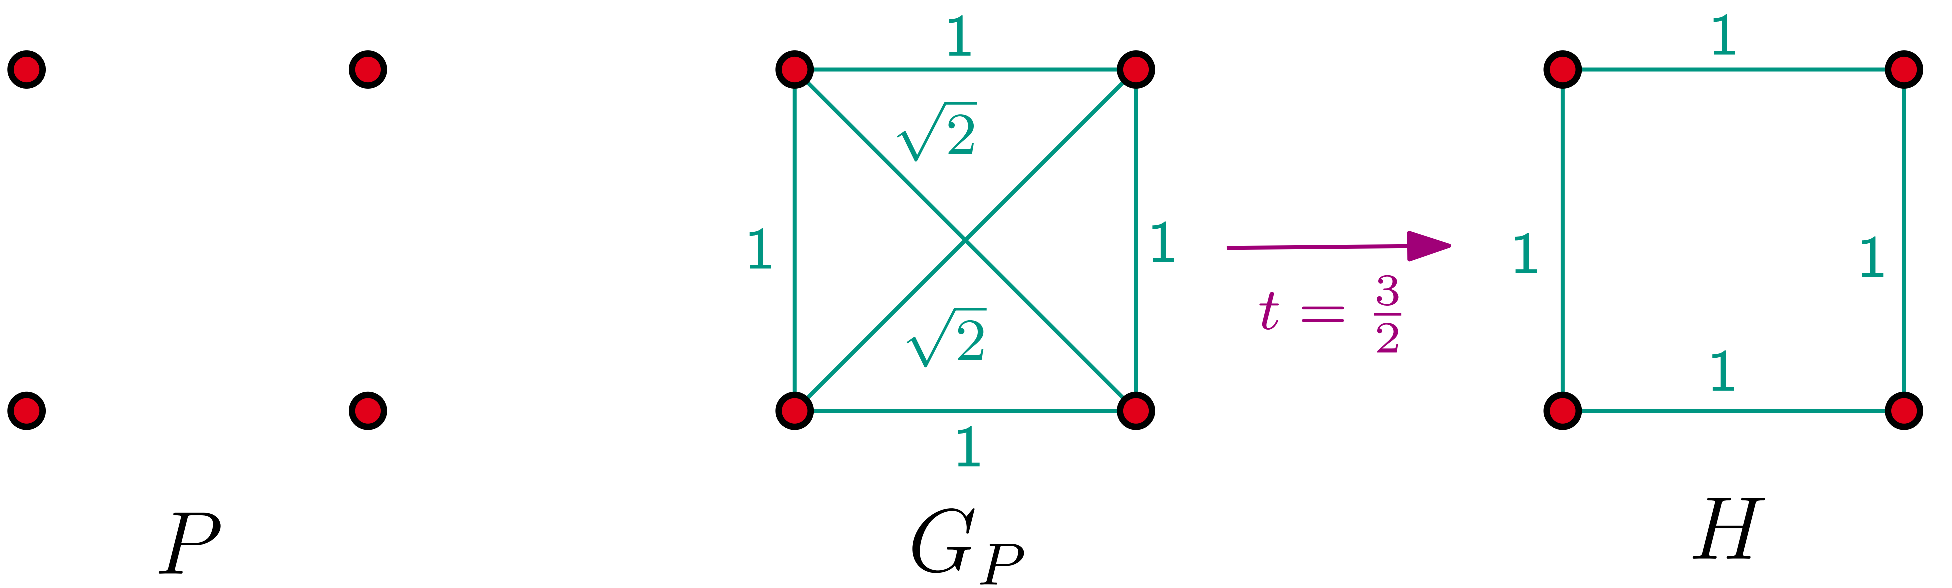

$Sparsity = \frac{4}{3}$

A brief History on Sparse Euclidean Spanners

- Q: How sparse a spanner should be ...?

A brief History on Sparse Euclidean Spanners

- Q: How sparse a spanner should be ...?

- Preferably, $O(|S|)$ with stretch factor $(1+\varepsilon)$, for a set S of n points.

A brief History on Sparse Euclidean Spanners

- Q: How sparse a spanner should be ...?

  - Preferably, $O(|S|)$ with stretch factor $(1+\varepsilon)$, for a set S of n points.

- Chew [SoCG'86] showed an existence of spanners with linear number of edges with stretch factor $\sqrt{10}$.
- Clarckson [STOC'87] designed first $(1+\varepsilon)$-spanner; Keil [SWAT'88] gave an alternative algorithm.
- Delanauy triangulation of the point set S is a 2.42-spanner [DCG'92].
- Θ-graphs help designing spanners in $\mathbb{R}^2$.
  This was generalized to $\mathbb{R}^d$ by Ruppert and Seidel [CCCG'91].

A brief History on Sparse Euclidean Spanners

- Q: How sparse a spanner should be ...?

  - Preferably, $O(|S|)$ with stretch factor $(1+\varepsilon)$, for a set S of n points.

- Chew [SoCG'86] showed an existence of spanners with linear number of edges with stretch factor $\sqrt{10}$.
- Clarckson [STOC'87] designed first $(1+\varepsilon)$-spanner; Keil [SWAT'88] gave an alternative algorithm.
- Delanauy triangulation of the point set S is a 2.42-spanner [DCG'92].
- Θ-graphs help designing spanners in $\mathbb{R}^2$.
  This was generalized to $\mathbb{R}^d$ by Ruppert and Seidel [CCCG'91].

Question: Is the trade-off between the stretch factor $1+\varepsilon$ and the sparsity $O(\varepsilon^{-d+1})$ tight?

Lightness

- For a finite set S in a metric space, the lightness of a spanner H is

$$\frac{\|H\|}{\|MST(S)\|} = \frac{\sum_{e \in E(H)} \|e\|}{\|MST(S)\|},$$

the ratio of the weight of H to the weight of a Euclidean MST of S.

Lightness

- For a finite set S in a metric space, the lightness of a spanner H is

$$\frac{\|H\|}{\|MST(S)\|} = \frac{\sum_{e \in E(H)} \|e\|}{\|MST(S)\|},$$

the ratio of the weight of H to the weight of a Euclidean MST of S.

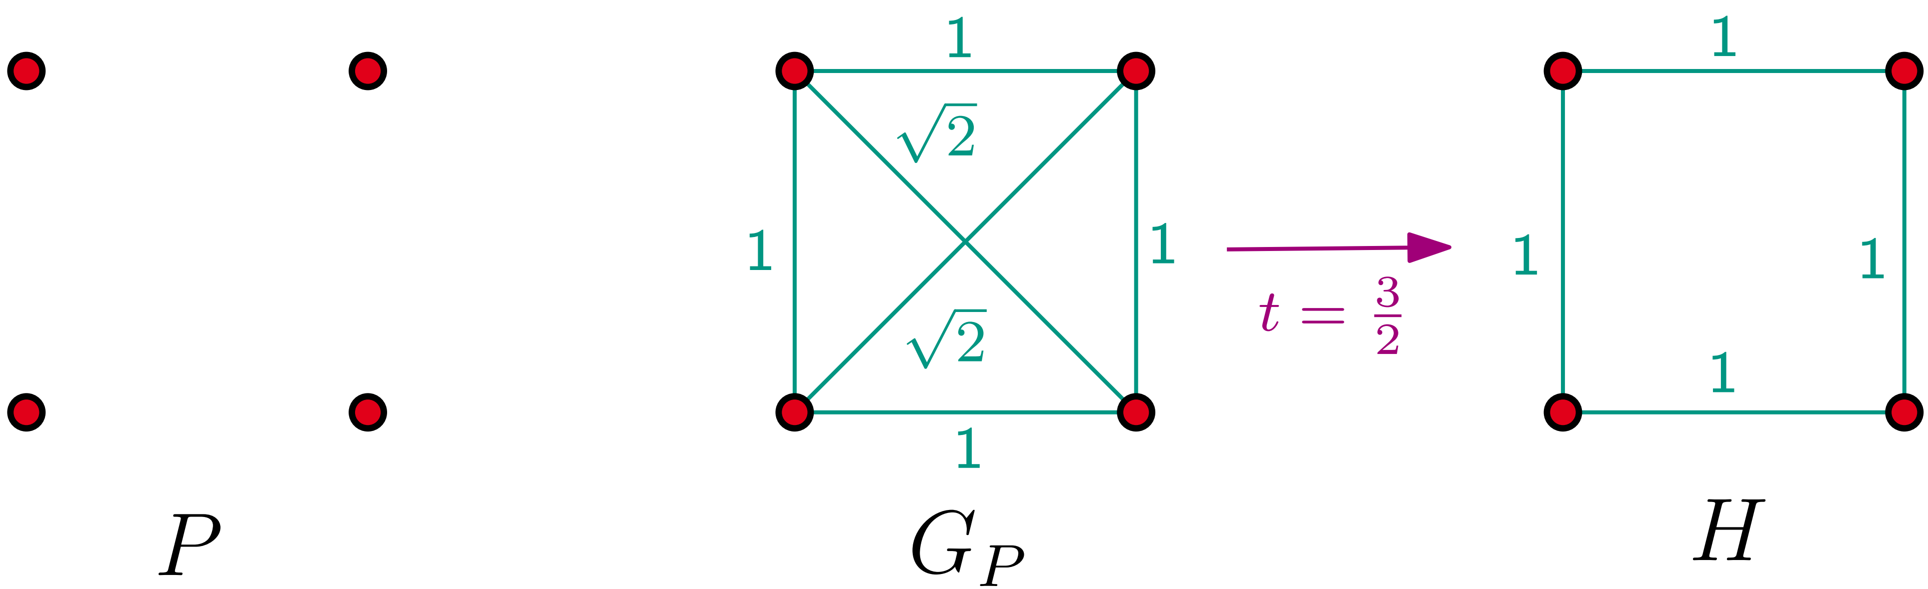

$$Lightness = \frac{4}{3}$$

A brief history on Light Spanners

- Greedy-spanner has constant lightness in $\mathbb{R}^3$ [Das et al.,SoCG'93]; which was later generalized to $\mathbb{R}^d$ [Das et al., SODA'95].
- In fact, greedy spanner has lightness $\varepsilon^{-O(d)}$ in $\mathbb{R}^d$, for every constant d. [Rao & Smith, STOC'98].
- $(1+\varepsilon)$-spanner with lightness ε^{-2d} exists [Narasimhan & Smid. Geometric Spanner Networks].
- A metric of doubling dimension d has a spanner of lightness $(d/\varepsilon)^{O(d)}$ [Gottlieb, FOCS'15]
- Greedy $(1+\varepsilon)$-spanner of a finite metric space of doubling dimension d has lightness $\varepsilon^{-O(d)}$ [Borradaile et al., SODA'19].

A brief history on Light Spanners

- Greedy-spanner has constant lightness in $\mathbb{R}^3$ [Das et al.,SoCG'93]; which was later generalized to $\mathbb{R}^d$ [Das et al., SODA'95].
- In fact, greedy spanner has lightness $\varepsilon^{-O(d)}$ in $\mathbb{R}^d$, for every constant d. [Rao & Smith, STOC'98].

Question: What is the best possible constant in the exponent?

- $(1+\varepsilon)$-spanner with lightness ε^{-2d} exists [Narasimhan & Smid. Geometric Spanner Networks].
- A metric of doubling dimension d has a spanner of lightness $(d/\varepsilon)^{O(d)}$ [Gottlieb, FOCS'15]
- Greedy $(1+\varepsilon)$-spanner of a finite metric space of doubling dimension d has lightness $\varepsilon^{-O(d)}$ [Borradaile et al., SODA'19].

Spanners in Metric Spaces.

	Stretch	Sparsity	Lightness
General	$(2k-1) \cdot (1+\varepsilon)$	$O(n^{1/k})$ [ADDJS93]	$O(n^{1/k}/\varepsilon^3)$ [CW16]
Euclidean	$(1+\varepsilon)$	$O(\varepsilon^{1-d})$ [Yao82]	$O(\varepsilon^{-d})$ [LS19]
Doubling	$(1+\varepsilon)$	$O(\varepsilon^{-O(d)})$ [HM05]	$O(\varepsilon^{-O(d)})$ [BLW19]
Minor-free	$(1+\varepsilon)$	$O(1)$ better.. ?	$O(\varepsilon^{-3})$ [BLW17]

Precise Dependency on $\varepsilon > 0$ & d.

- Le and Solomon in FOCS'19 established the dependencies of ε in the lightness and sparsity bounds of Euclidean $(1+\varepsilon)$-spanners.

Precise Dependency on $\varepsilon > 0$ & d.

- Le and Solomon in FOCS'19 established the dependencies of ε in the lightness and sparsity bounds of Euclidean $(1+\varepsilon)$-spanners.

For every $\varepsilon > 0$ and constant $d \in \mathbb{N}$, and a set S of n points in $\mathbb{R}^d$,

- every $(1+\varepsilon)$-spanner must have lightness $\Omega(\varepsilon^{-d})$ and sparsity $\Omega(\varepsilon^{-d+1})$, whenever $\varepsilon = \Omega(n^{-1/(d-1)})$.
- The greedy $(1+\varepsilon)$-spanner in $\mathbb{R}^d$ has lightness $O(\varepsilon^{-d} \log \varepsilon^{-1})$.

Last few years (highlights) ...

Trade-off Between Degree, Diameter, and Lightness -

- Optimal Euclidean Spanners: Really Short, Thin, and Lanky [Elkin & Solomon, J.ACM 2015]

Trade-off Between Degree, Lightness -

- Unified Framework for Light Spanners [Le & Solomon, STOC 2023]

Trade-off Between Degree & Sparsity -

- Sparse Euclidean Spanners with Optimal Diameter: A General and Robust Lower Bound via a Concave Inverse-Ackermann Function [Le, Milenkovic & Solomon, SoCG 2023]

Steiner Points - The Game Changer!

- Steiner points can substantially improve bounds on the lightness and sparsity of Euclidean $(1+\varepsilon)$-spanners [Le and Solomon, FOCS'19].

Steiner Points - The Game Changer!

- Steiner points can substantially improve bounds on the lightness and sparsity of Euclidean $(1+\varepsilon)$-spanners [Le and Solomon, FOCS'19].

	Sparsity	Lightness
Lower Bound	$\Omega(\varepsilon^{-\frac{1}{2}} / \log \varepsilon^{-1})$, for $d = 2$ [Le & Solomon, FOCS'19]	$\Omega(\varepsilon^{-1} / \log \varepsilon^{-1})$, for $d = 2$ [Le & Solomon, FOCS'19]
Upper Bound	$O(\varepsilon^{(1-d)/2})$ for d-space [Le & Solomon, FOCS'19]	$O(\varepsilon^{-1} \log \Delta)$, for $d = 2$ [Le & Solomon, ESA'20] $\tilde{O}(\varepsilon^{-(d+1)/2})$, for $d \geq 3$ [Le & Solomon, ArXiv'20]

Effectiveness of Steiner Points - Steiner Ratios, etc.

- Steiner points can improve the weight of the network in the single-source setting.

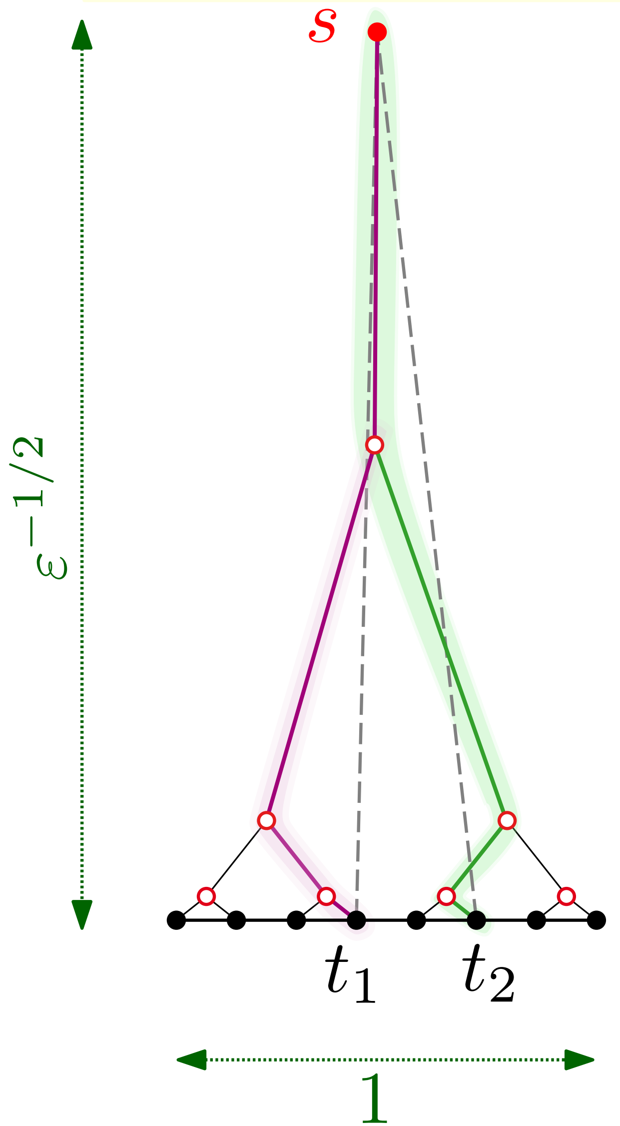

Exponential improvement on the lightness in a metric space [Elkin & Solomon, SICOMP'15]
Quadratic improvement on the lightness in Euclidean spaces [Solomon, JoCG'15].

Effectiveness of Steiner Points - Steiner Ratios, etc.

- Steiner points can improve the weight of the network in the single-source setting.

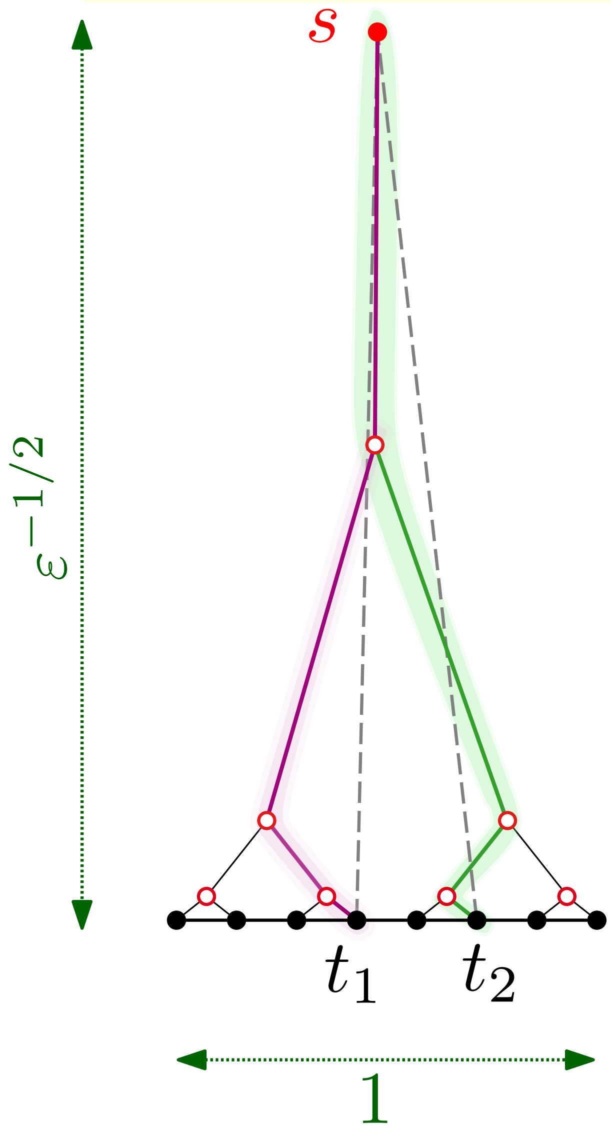

Exponential improvement on the lightness in a metric space [Elkin & Solomon, SICOMP'15]
Quadratic improvement on the lightness in Euclidean spaces [Solomon, JoCG'15].

Shallow-light tree of weight $O(\varepsilon^{-1/2})$.
Without Steiner point, we would need a star centered at s, of weight $\Theta(\varepsilon^{-1})$ to guarantee a stretch factor $\leq 1 + \varepsilon$.

Improved Bounds

Lower Bound [Bhore & Tóth, SIDMA'22]

Theorem 1 *Let a positive integers* d *and real* $\varepsilon > 0$ *be given such that* $\varepsilon \leq 1/d$. *Then there exists a set* S *of* n *points in* $\mathbb{R}^d$ *such that any Euclidean Steiner* $(1+\varepsilon)$*-spanner for* S *has lightness* $\Omega(\varepsilon^{-d/2})$ *and sparsity* $\Omega(\varepsilon^{(1-d)/2})$.

Improved Bounds

Lower Bound [Bhore & Tóth, SIDMA'22]

Theorem 1 *Let a positive integers d and real $\varepsilon > 0$ be given such that $\varepsilon \leq 1/d$. Then there exists a set S of n points in $\mathbb{R}^d$ such that any Euclidean Steiner $(1+\varepsilon)$-spanner for S has lightness $\Omega(\varepsilon^{-d/2})$ and sparsity $\Omega(\varepsilon^{(1-d)/2})$.*

Upper Bound [Bhore &Tóth, SoCG'21]

Theorem 2 *For every set S of n points in Euclidean plane, there exists a Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.*

Where do we stand ...

- All bounds are for Euclidean Steiner $(1+\varepsilon)$-spanner

	Sparsity	Lightness
Lower Bound	$\Omega(\varepsilon^{-1/2}/\log\varepsilon^{-1})$ [Le&Solomon, FOCS'19] $\Omega(\varepsilon^{(1-d)/2})$ [Bhore&Tóth, SIDMA'22]	$\Omega(\varepsilon^{-1}/\log\varepsilon^{-1})$, for $d=2$ [Le&Solomon, FOCS'19] $\Omega(\varepsilon^{-d/2})$ [Bhore&Tóth, SIDMA'22]
Upper Bound	$O(\varepsilon^{(1-d)/2})$ for d-space [Le & Solomon, FOCS'19]	$\Omega(\varepsilon^{-1}\log\Delta)$, for $d=2$ [Le & Solomon,ESA'20] $\tilde{O}(\varepsilon^{-(d+1)/2})$, for $d\geq 3$ [Le & Solomon, STOC'23] $O(\varepsilon^{-1})$, for $d=2$ [Bhore&Tóth, SoCG'21]

Lower Bounds

Basic Observation. Every ab-path of length at most $(1+\varepsilon)\|ab\|$ lies in the ellipoid $\mathcal{E}_{ab}$ with foci a and b and great axis $(1+\varepsilon)\|ab\|$.

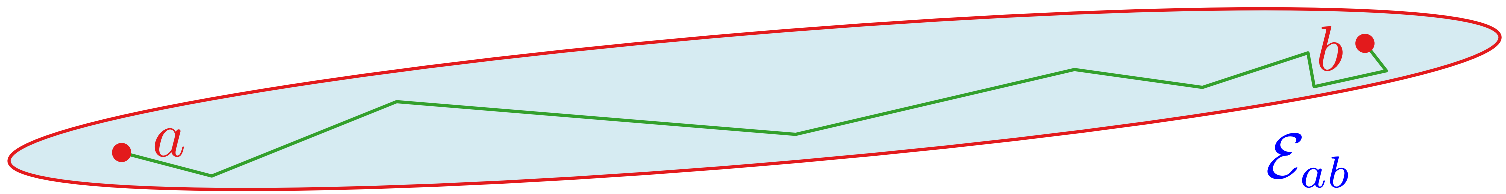

The spacing between the points guarantees that we obtain *disjoint* ellipsoids $\mathcal{E}_{ab}$ for a family of parallel ab pairs.

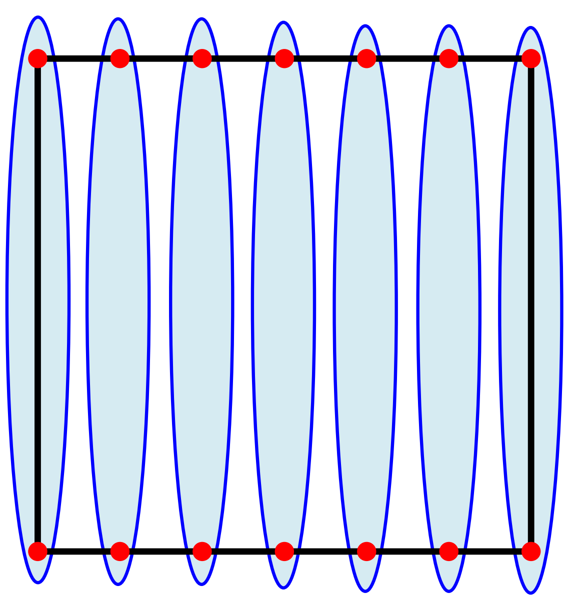

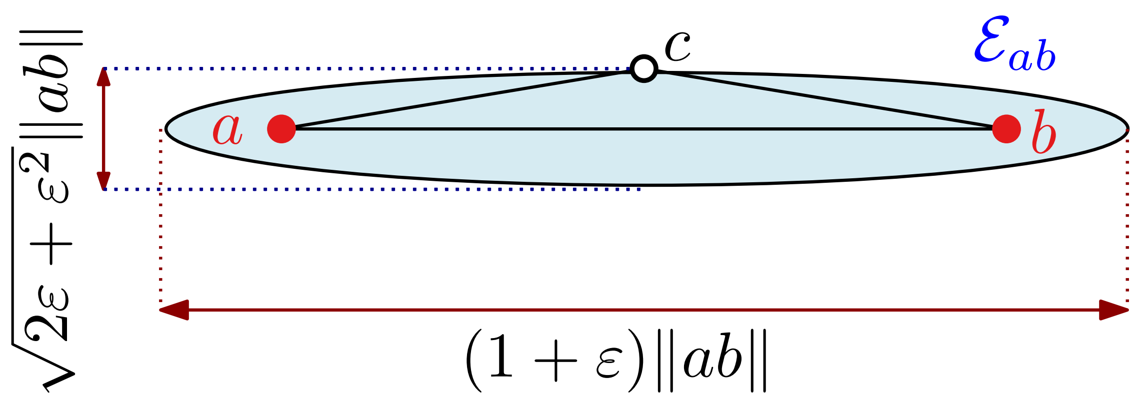

Lower Bounds

Basic Observation. Every ab-path of length at most $(1+\varepsilon)\|ab\|$ lies in the ellipoid $\mathcal{E}_{ab}$ with foci a and b and great axis $(1+\varepsilon)\|ab\|$.

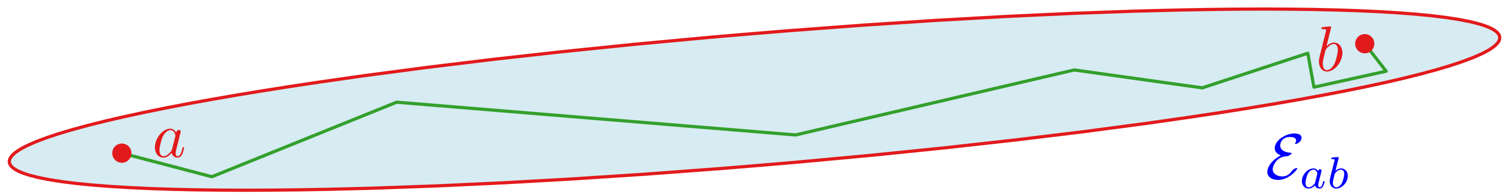

The spacing between the points guarantees that we obtain *disjoint* ellipsoids $\mathcal{E}_{ab}$ for a family of parallel ab pairs.

However, ellipsoids $\mathcal{E}_{ab}$ and $\mathcal{E}_{cd}$ may overlap in general.

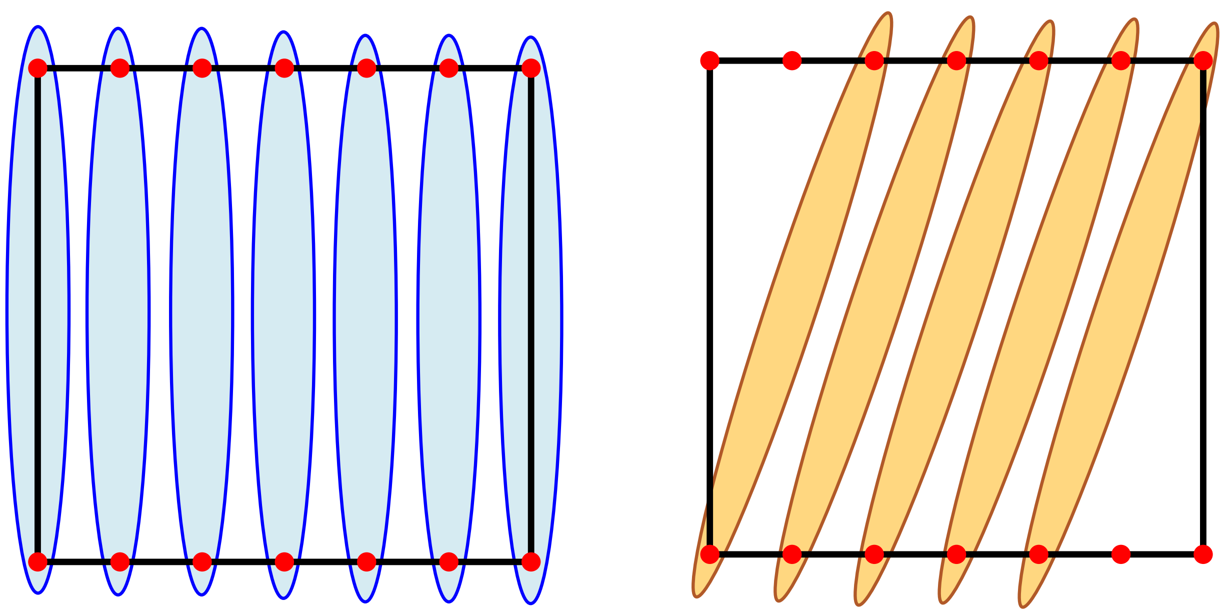

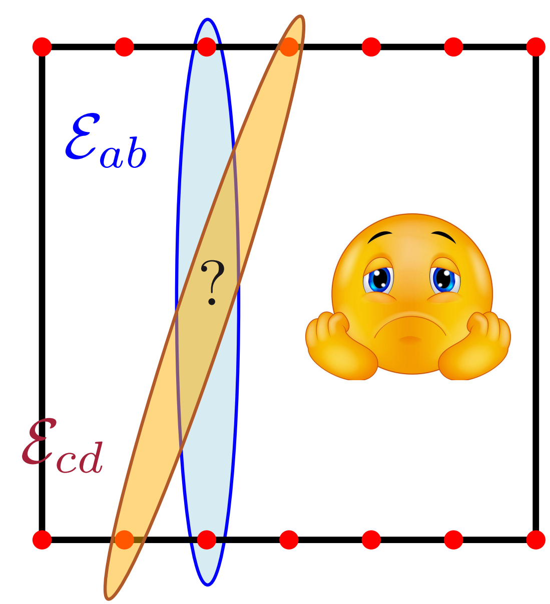

Lower Bounds

An ab-path P_{ab} of weight $\leq (1+\varepsilon)\|ab\|$ is "nearly" parallel to ab.

Let $E(\alpha) = \{\text{edges } e \text{ in } P_{ab} \text{ with } \angle(ab, e) \leq \alpha\}$

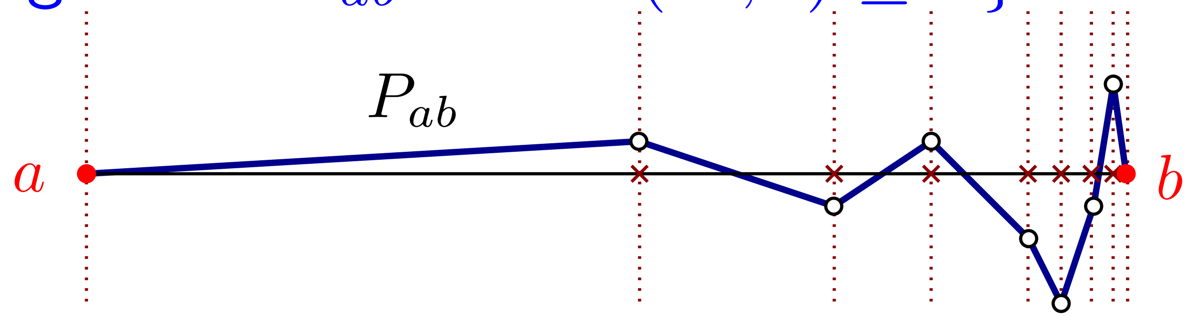

Lemma. For $i = 1, \ldots, \lfloor 1/\sqrt{\varepsilon} \rfloor$, $\|E(i \cdot \sqrt{\varepsilon})\| \geq \left(1 - \frac{2}{i^2}\right) \|ab\|$.

Corollary. Every ab-path of weight $\leq (1+\varepsilon)\|ab\|$ contains edges of direction $\angle(ab, e) \leq 2 \cdot \sqrt{\varepsilon}$ of total weight $\geq \frac{1}{2}\|ab\|$.

In each ellipsoid $\mathcal{E}_{ab}$,
we count only edges of
direction $\angle(ab, e) \leq 2 \cdot \sqrt{\varepsilon}$.
$\Rightarrow$ no edge is counted twice.

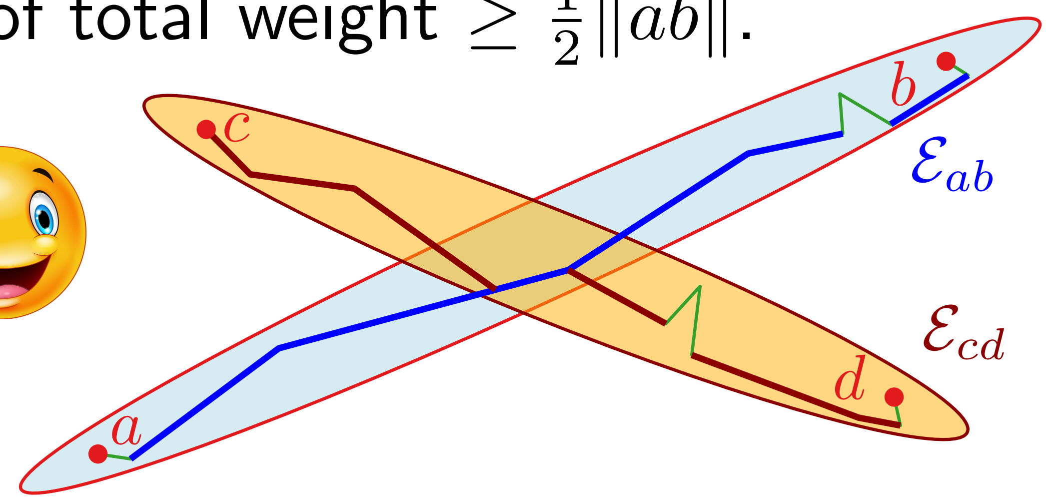

Lower Bounds

Lightness Lower Bounds

Theorem. $\forall d \geq 2$, $\forall \varepsilon > 0$, with $\varepsilon \leq 1/d$, there is a set S of n points in $\mathbb{R}^d$ such that any Euclidean Steiner $(1+\varepsilon)$-spanner N for S has lightness $\Omega_d(\varepsilon^{\frac{-d}{2}})$.

Construction. Let $S = A \cup B$, grids on two opposite faces of a unit cube, with $\frac{d}{\sqrt{\varepsilon}}$ spacing.

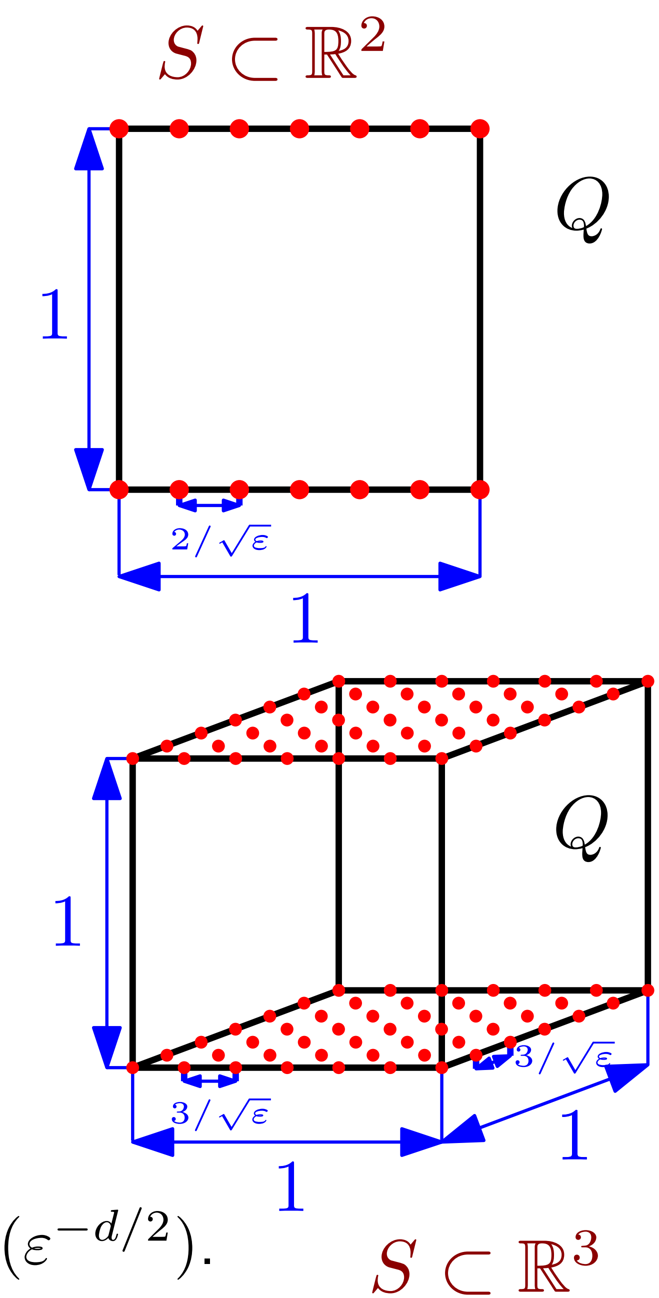

Lightness analysis.

- $|S| = \Theta_d(\varepsilon^{(1-d)/2})$.
- $\|\text{MST}(S)\| = 1 + (|S| - 2)\frac{d}{\sqrt{\varepsilon}} = \Theta_d(\varepsilon^{1-d/2})$.
- $\|N\| \geq \sum_{(a,b)\in A\times B} \|E_{ab}(2\sqrt{\varepsilon})\|$
  $\geq \sum_{(a,b)\in A\times B} \frac{1}{2}\|ab\|$
  $\geq |A \times B| \cdot \frac{1}{2} \geq \Omega(|S|^2) \geq \Omega_d(\varepsilon^{1-d})$.
- $\text{Lightness}(N) = \frac{\|N\|}{|\text{MST}(S)\|} \geq \Omega_d(\frac{\varepsilon^{1-d}}{\varepsilon^{1-d/2}}) \geq \Omega_d(\varepsilon^{-d/2})$.

Lower Bounds

Sparsity Lower Bounds

Theorem. $\forall d \geq 2$, $\forall \varepsilon > 0$, with $\varepsilon \leq 1/d$, there is a set S of n points in $\mathbb{R}^d$ such that any Euclidean Steiner $(1+\varepsilon)$-spanner N for S has sparsity $\Omega_d(\varepsilon^{\frac{1-d}{2}})$.

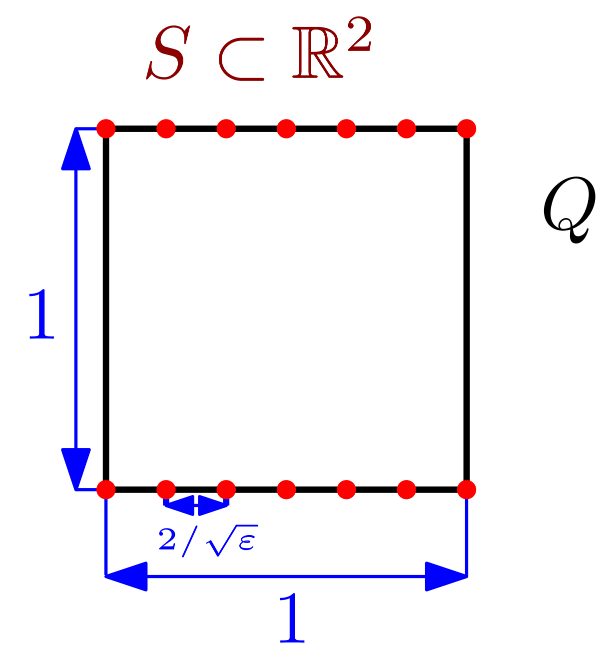

Construction. Let $S = A \cup B$, grids on two opposite faces of a unit cube, with $\frac{d}{\sqrt{\varepsilon}}$ spacing.

Sparsity analysis.

- $|S| = \Theta_d(\varepsilon^{(1-d)/2})$.
- $\|N\| \geq \Omega_d(\varepsilon^{1-d})$ [cf. lightness analysis]
- We may assume $N \subset Q$ [Le & Solomon, 2019], hence $\forall e \in E(N) : \|e| \leq \text{diam}(Q) = \sqrt{d}$.
- $|E(N)| \geq \frac{\|N\|}{\max_{e \in E(N)} \|e\|} \geq \Omega_d(\frac{\varepsilon^{1-d}}{\sqrt{d}}) \geq \Omega_d(\varepsilon^{1-d})$.
- Sparsity$(N) = \frac{|E(N)|}{|S|} \geq \Omega_d(\frac{\varepsilon^{1-d}}{\varepsilon^{(1-d)/2}}) \geq \Omega_d(\varepsilon^{\frac{1-d}{2}})$.

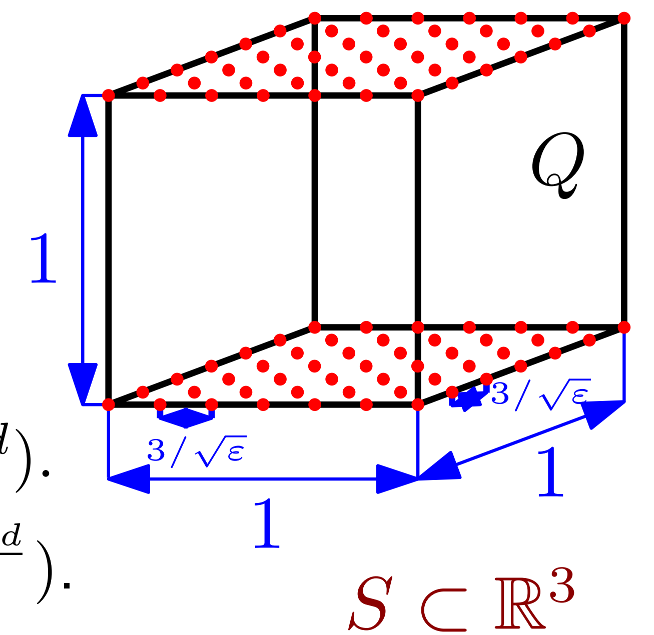

Upper Bound

- We prove and construct, for every $\varepsilon > 0$ and every set of n points in $\mathbb{R}^2$, a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.

Three Major Components

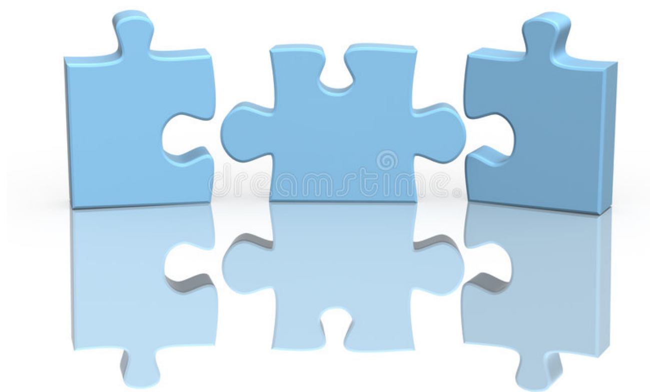

Upper Bound

- We prove and construct, for every $\varepsilon > 0$ and every set of n points in $\mathbb{R}^2$, a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.

Three Major Components

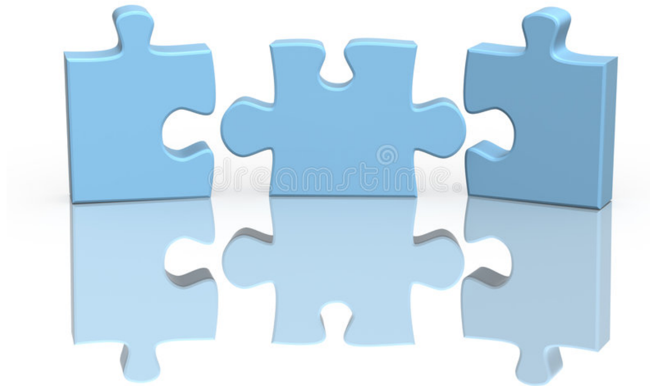

Directional Spanners

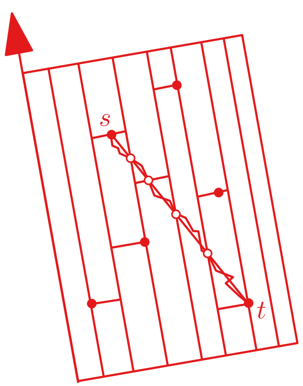

Upper Bound

- We prove and construct, for every $\varepsilon > 0$ and every set of n points in $\mathbb{R}^2$, a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.

Three Major Components

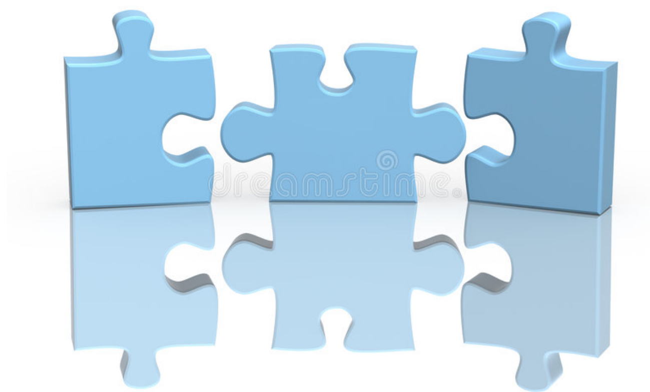

Directional Spanners

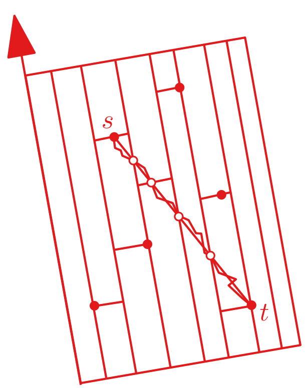

Generalized Shallow-light Trees

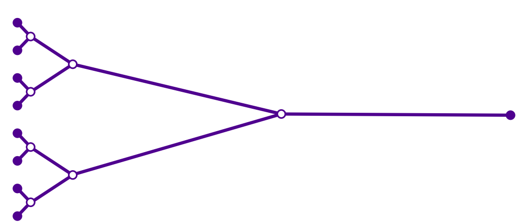

Upper Bound

- We prove and construct, for every $\varepsilon > 0$ and every set of n points in $\mathbb{R}^2$, a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.

Three Major Components

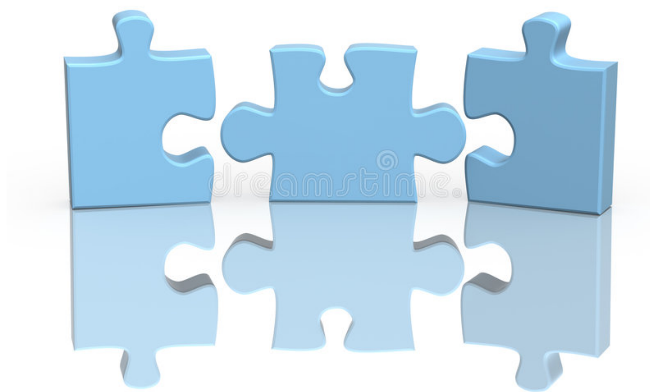

Directional Spanners

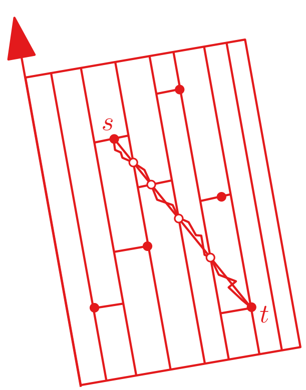

Generalized Shallow-light Trees

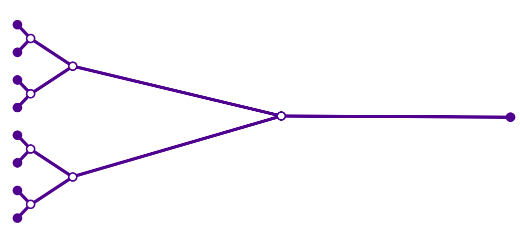

Modified Window-Partitoning Scheme

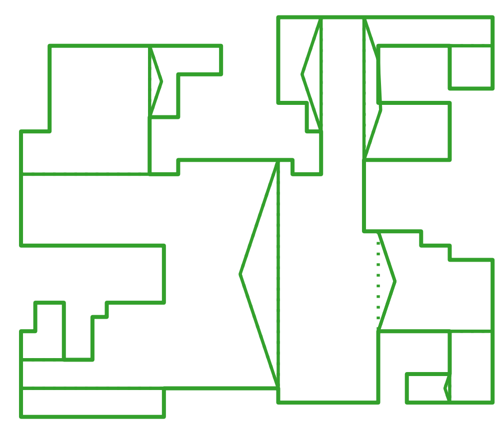

Upper Bound

Direction of a segment -

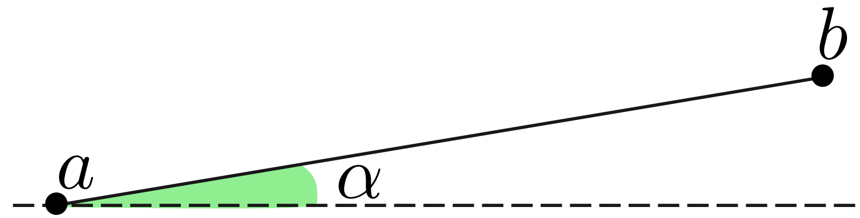

Angle-bounded path -

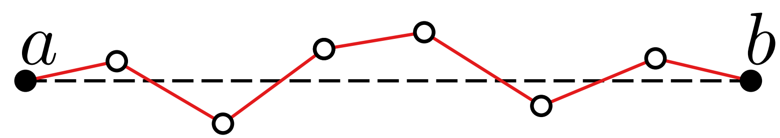

Directional $(1 + \varepsilon)$-Spanner

- For an interval $D \subset [0, \pi)$ of directions, we construct a Euclidean Steiner $(1 + \varepsilon)$-spanner restricted to points pairs whose directions are in D.

- Definition - A geometric graph G is a *directional* $(1 + \varepsilon)$*-spanner* for S and D if for every $a, b \in S$, where the direction of ab is in D, graph G contains an ab-path of weight at most $(1 + \varepsilon)\|ab\|$.

The main Lemma

Lemma 1 For a set S of n points in the plane, and for the interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$ of directions, there exists a directional $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{-1/2}\|MST(S)\|)$.

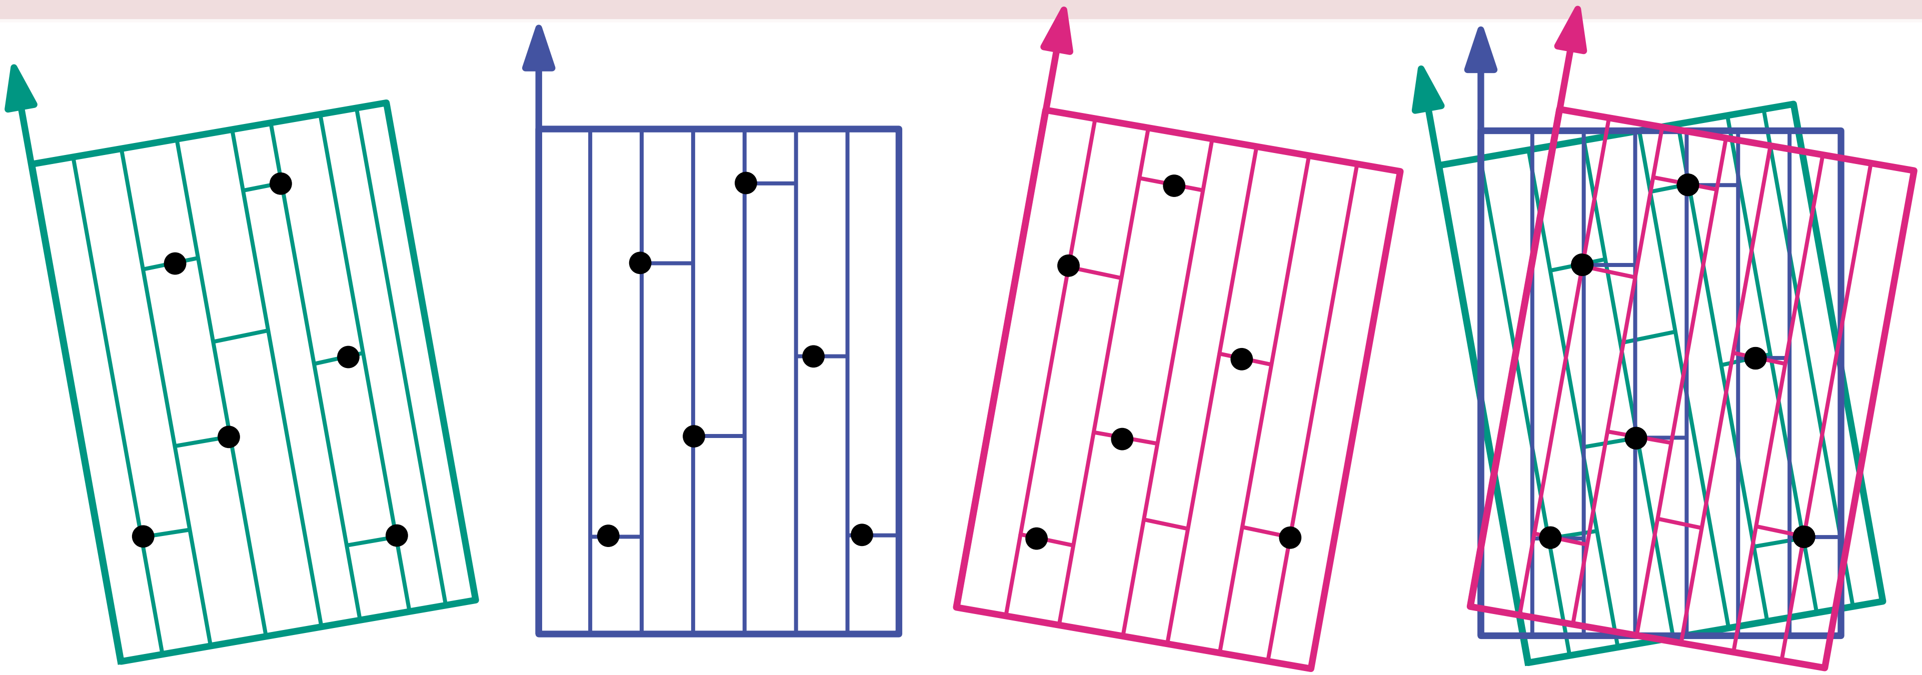

- Let $N = \bigcup_{i=1}^{k} N_i$ be the union of the networks N_i for $i \in \{1, \ldots, k\}$.
- The total weight of N, for $k = O(\varepsilon^{-1/2})$, is
  $\|N\| = \sum_{i=1}^{k} \|N_i\| \leq kO\left(\varepsilon^{-\frac{1}{2}}\|MST(S)\|\right) \leq O\left(\varepsilon^{-1}\|MST(S)\|\right)$.

The main Lemma

Lemma 1 For a set S of n points in the plane, and for the interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$ of directions, there exists a directional $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{-1/2}\|MST(S)\|)$.

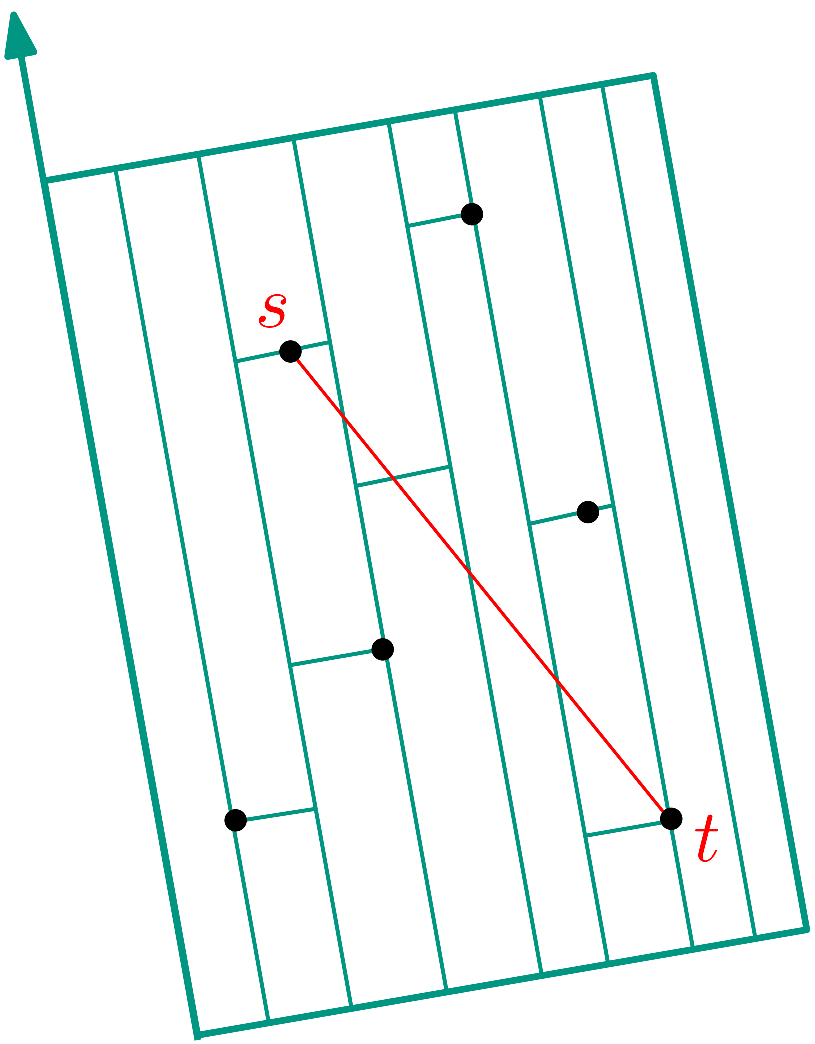

Strategy: for each interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$.

1. Find a tiling of a bounding box of S, of weight $O(\|MST(S)\|)$.
2. For each tile P, construct a directional spanner for a finite point set on the boundary of P, of weight $O(\varepsilon^{-1/2}\,\mathrm{per}(P))$.

The main Lemma

Lemma 1 For a set S of n points in the plane, and for the interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$ of directions, there exists a directional $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{-1/2}\|MST(S)\|)$.

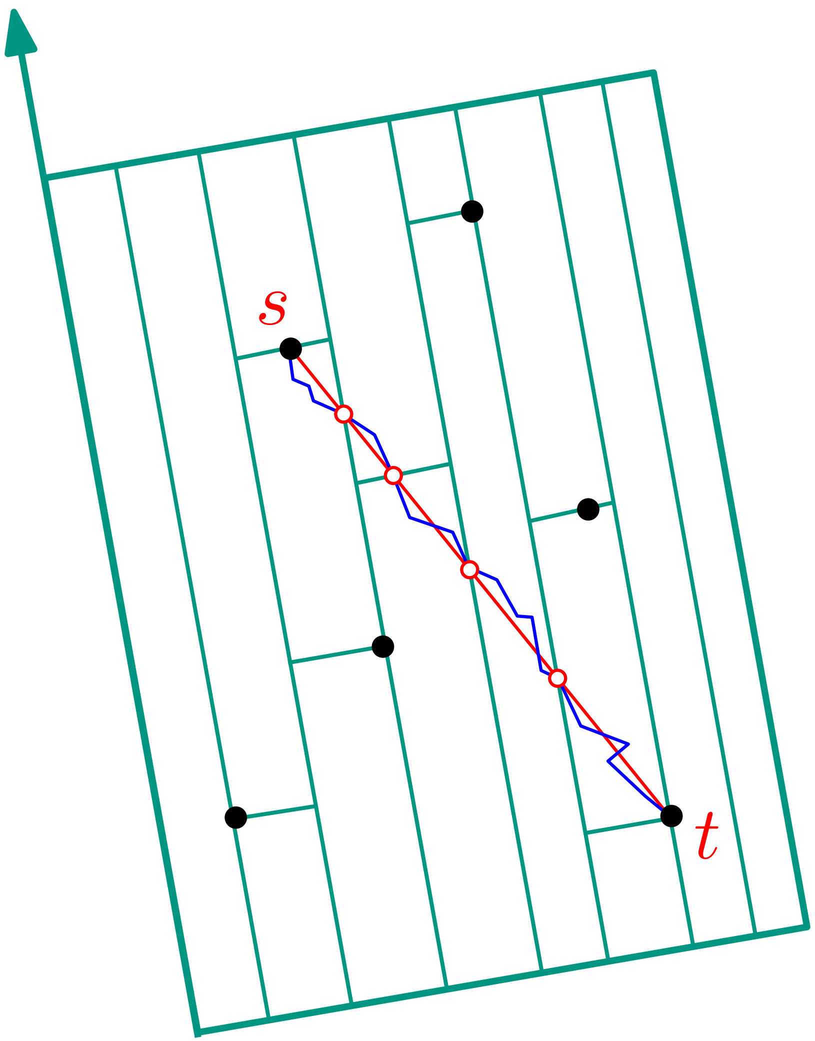

Strategy: for each interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$.

1. Find a tiling of a bounding box of S, of weight $O(\|MST(S)\|)$.
2. For each tile P, construct a directional spanner for a finite point set on the boundary of P, of weight $O(\varepsilon^{-1/2}\,\mathrm{per}(P))$.

The main Lemma

Lemma 1 For a set S of n points in the plane, and for the interval $D = [\frac{\pi}{2} - \sqrt{\varepsilon}, \frac{\pi}{2} + \sqrt{\varepsilon}]$ of directions, there exists a directional $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{-1/2} \| MST(S) \|)$.

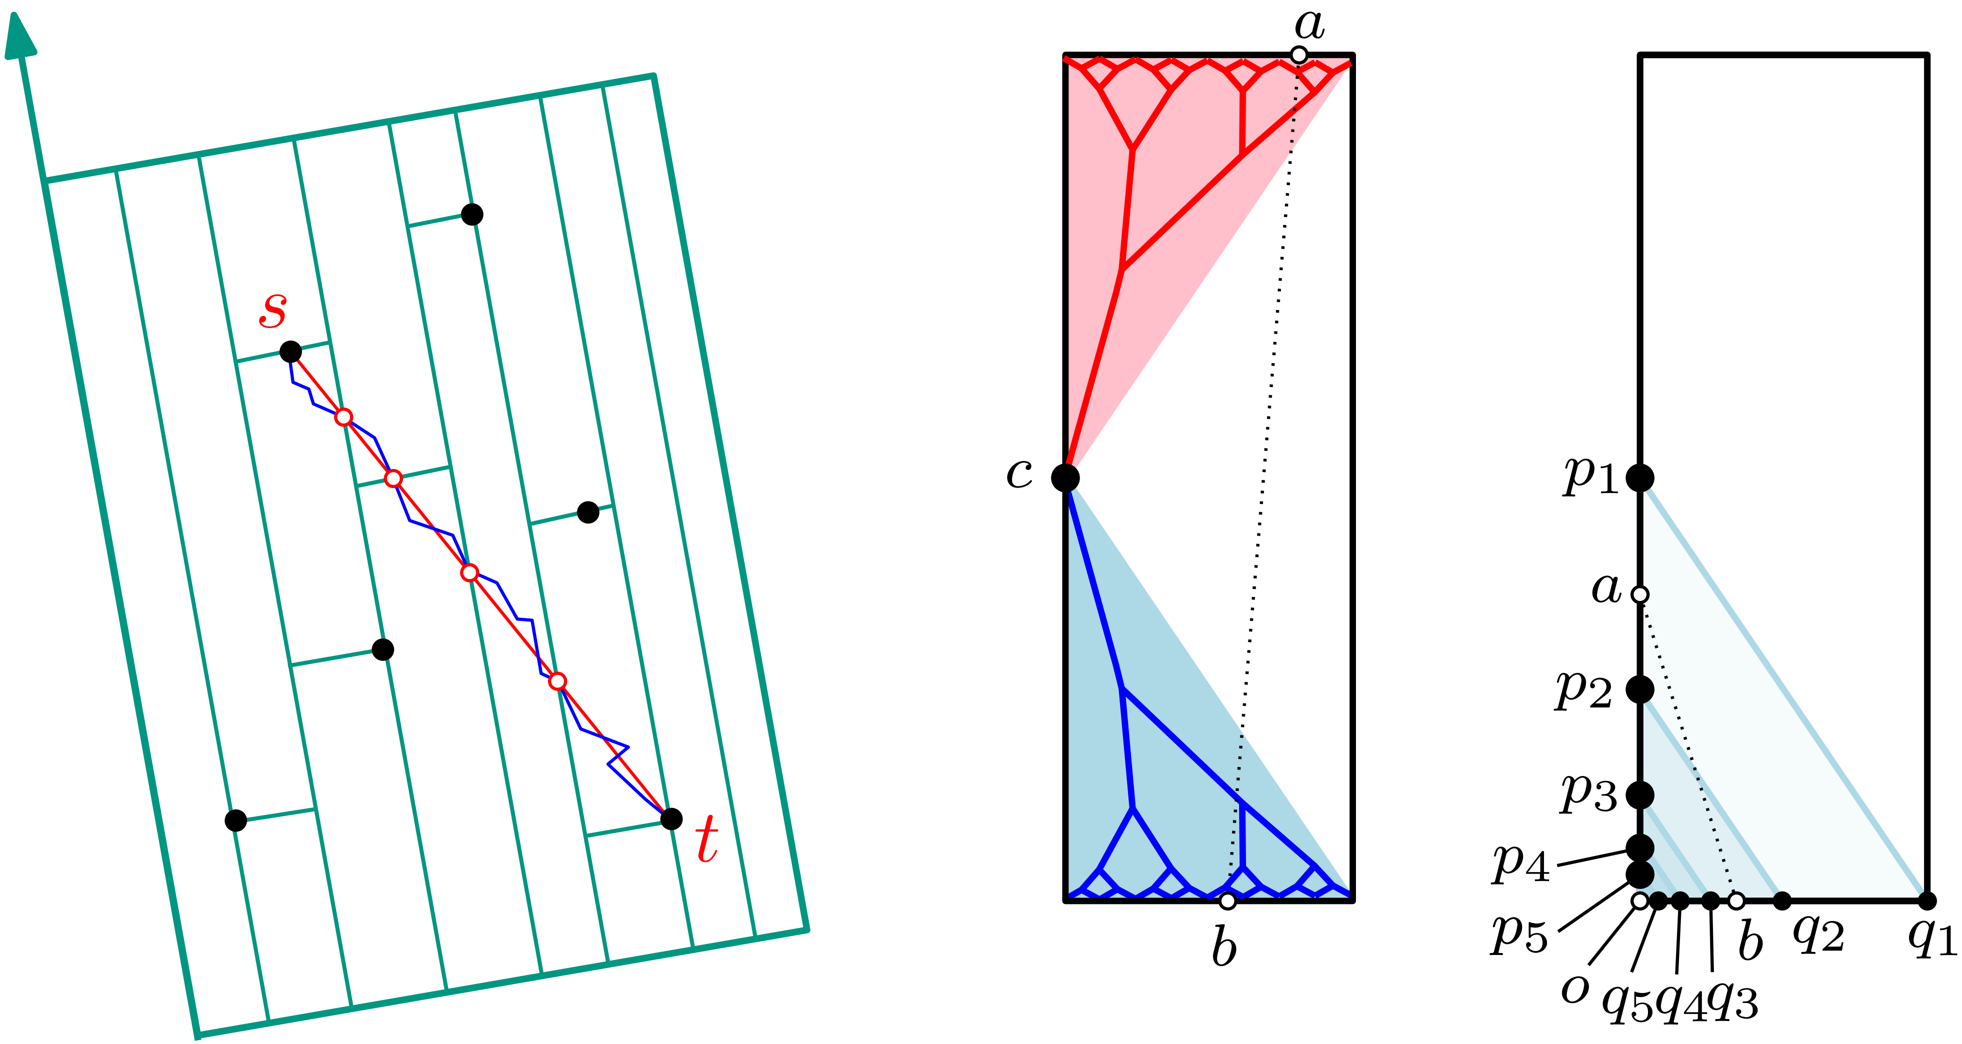

Tiling — Histograms — Window partitioning

Given a set S of n points in $\mathbb{R}^2$,
(1) $BB(S) \cup$ Reclilinear MST(S) gives a tiling.
(2) Each tile is a (weakly) simple rectilinear polygon.

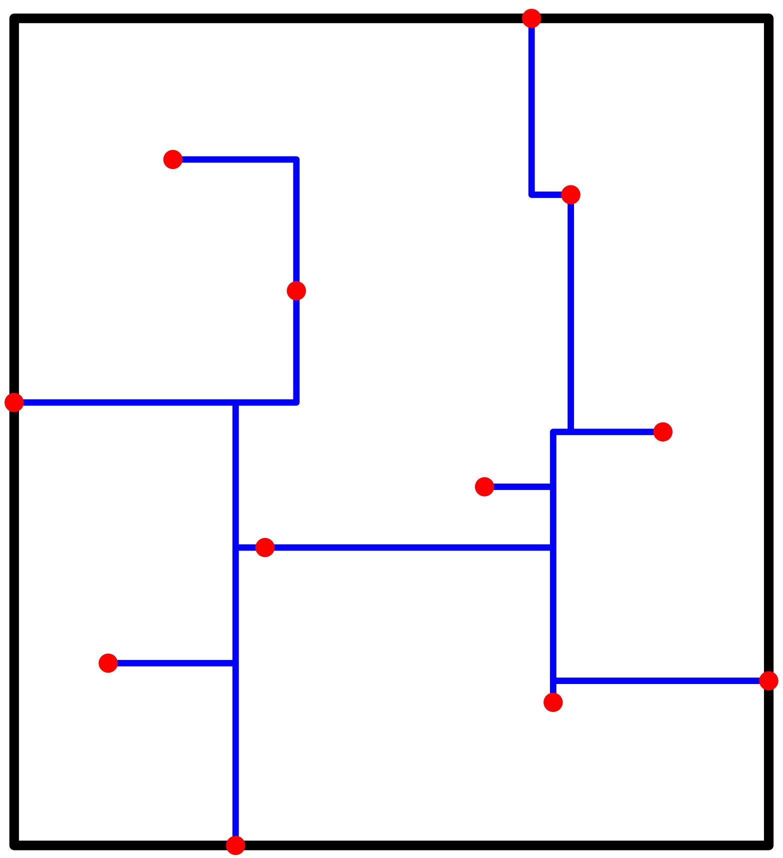

Tiling — Histograms — Window partitioning

Given a set S of n points in $\mathbb{R}^2$,
(1) $BB(S) \cup$ Reclilinear MST(S) gives a tiling.
(2) Each tile is a (weakly) simple rectilinear polygon.
(3) A rectangulation would have weight $O(\|MST(S)\| \log n)$.

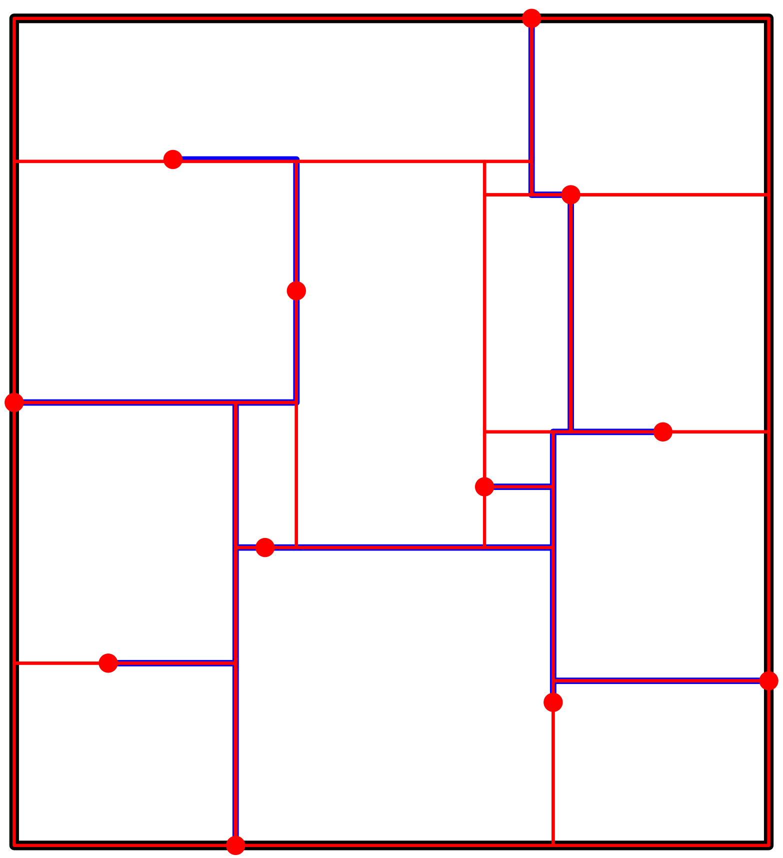

Tiling — Histograms — Window partitioning

Given a set S of n points in $\mathbb{R}^2$,
(1) $BB(S) \cup$ Reclilinear MST(S) gives a tiling.
(2) Each tile is a (weakly) simple rectilinear polygon.
(3) Compute the window-partition into rectilinear histograms.

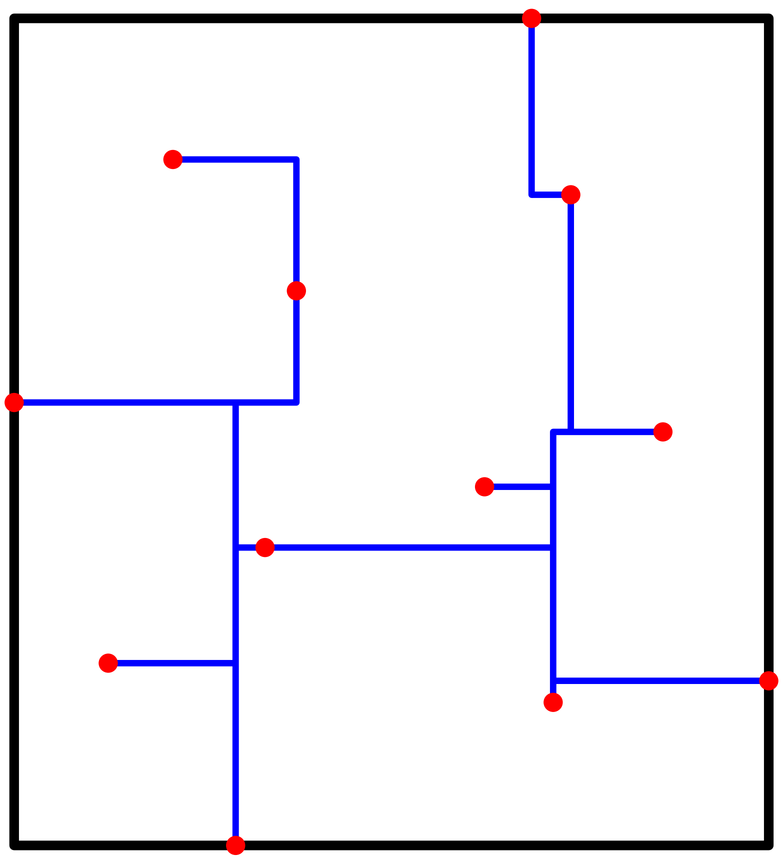

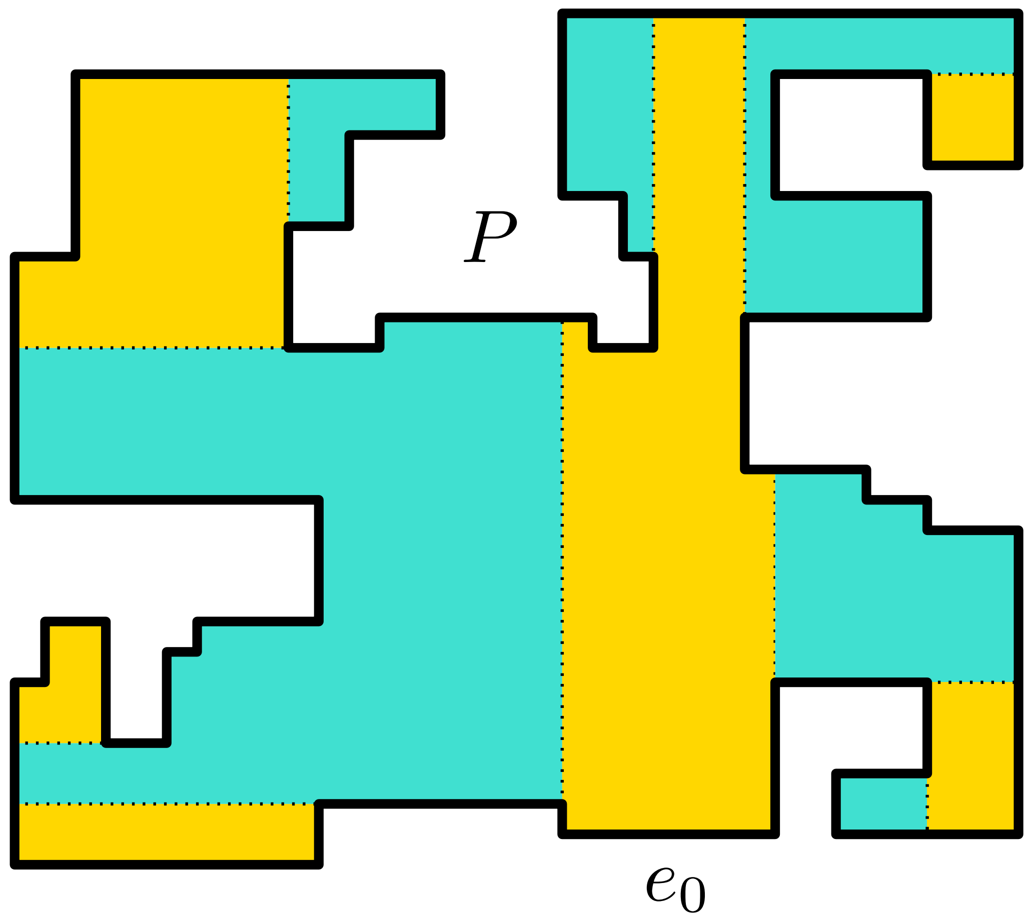

Shallow-Light Trees for Staircases

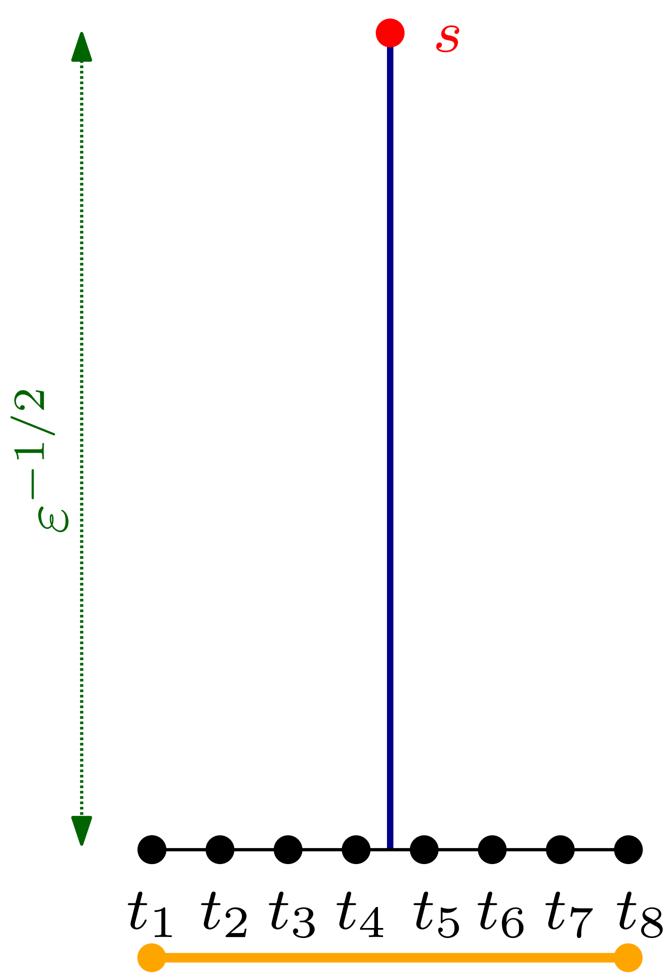

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Shallow-Light Trees for Staircases

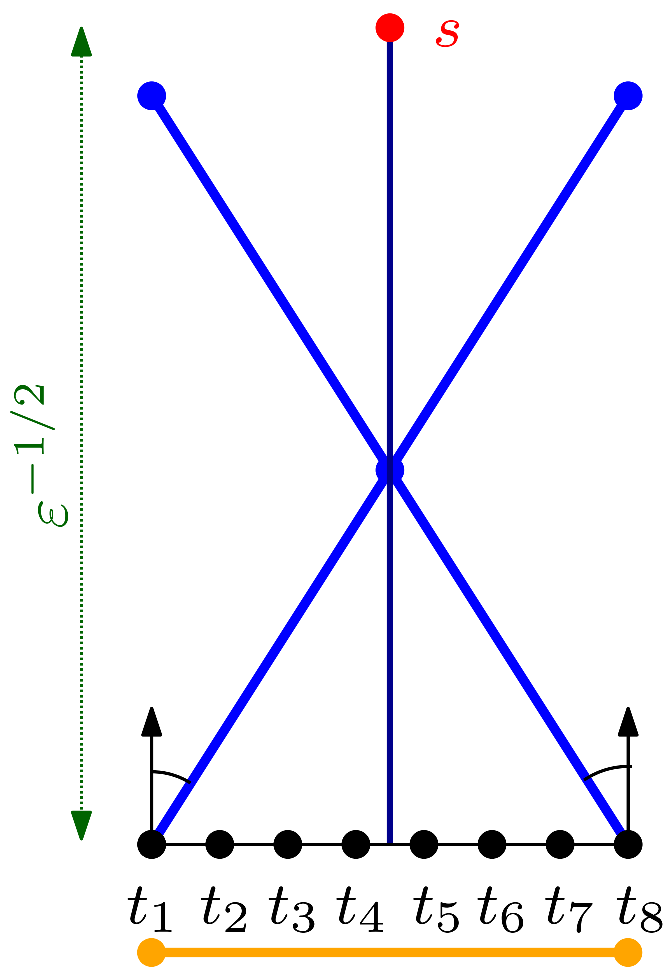

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Shallow-Light Trees for Staircases

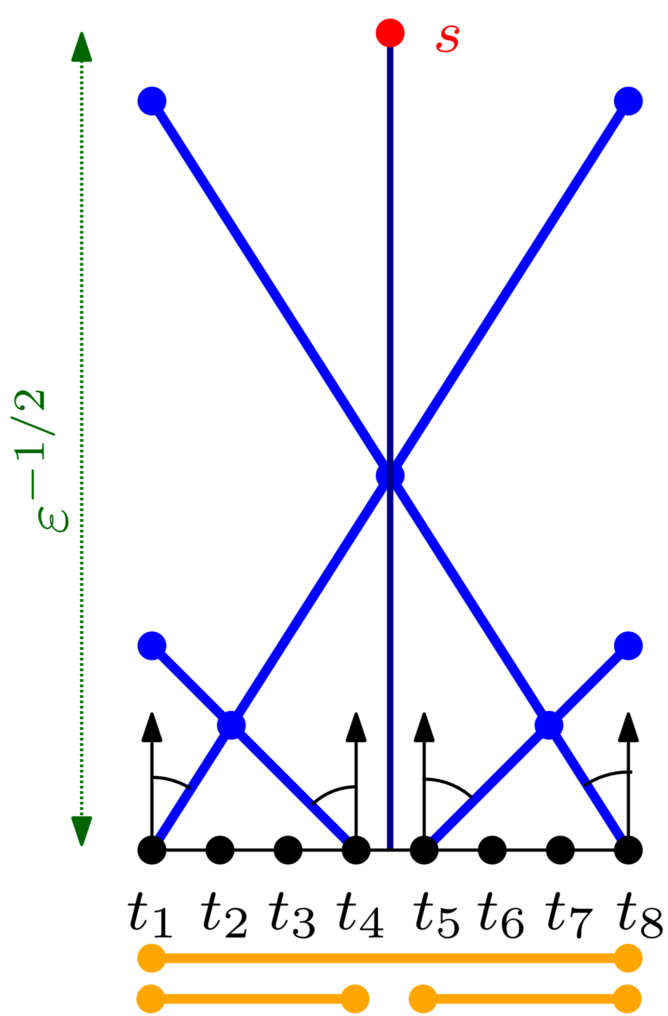

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Shallow-Light Trees for Staircases

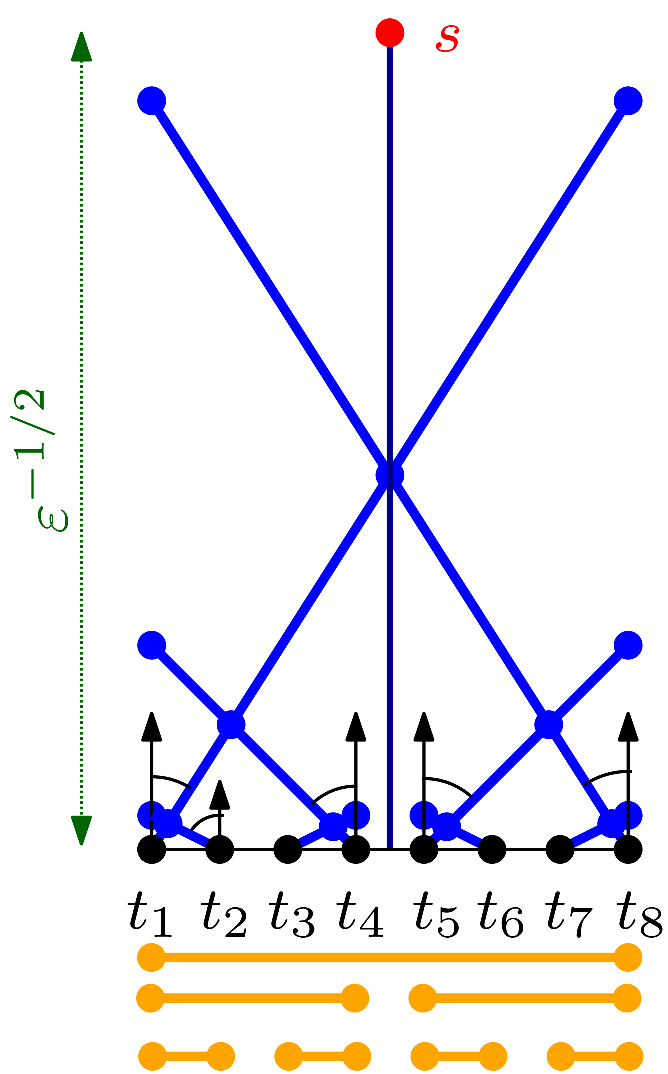

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Shallow-Light Trees for Staircases

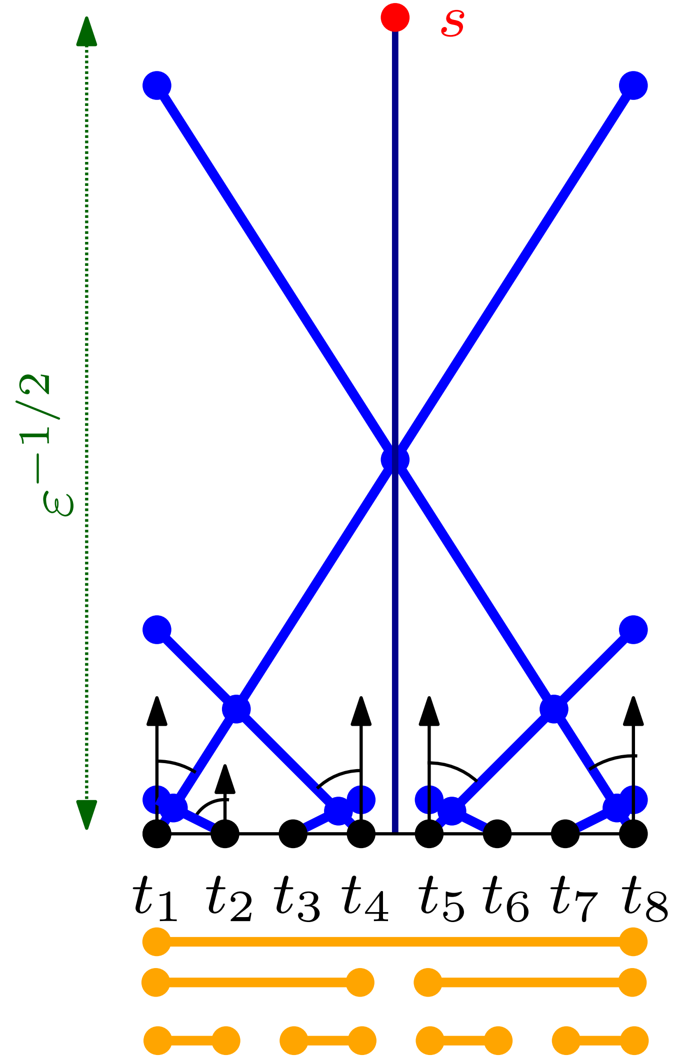

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

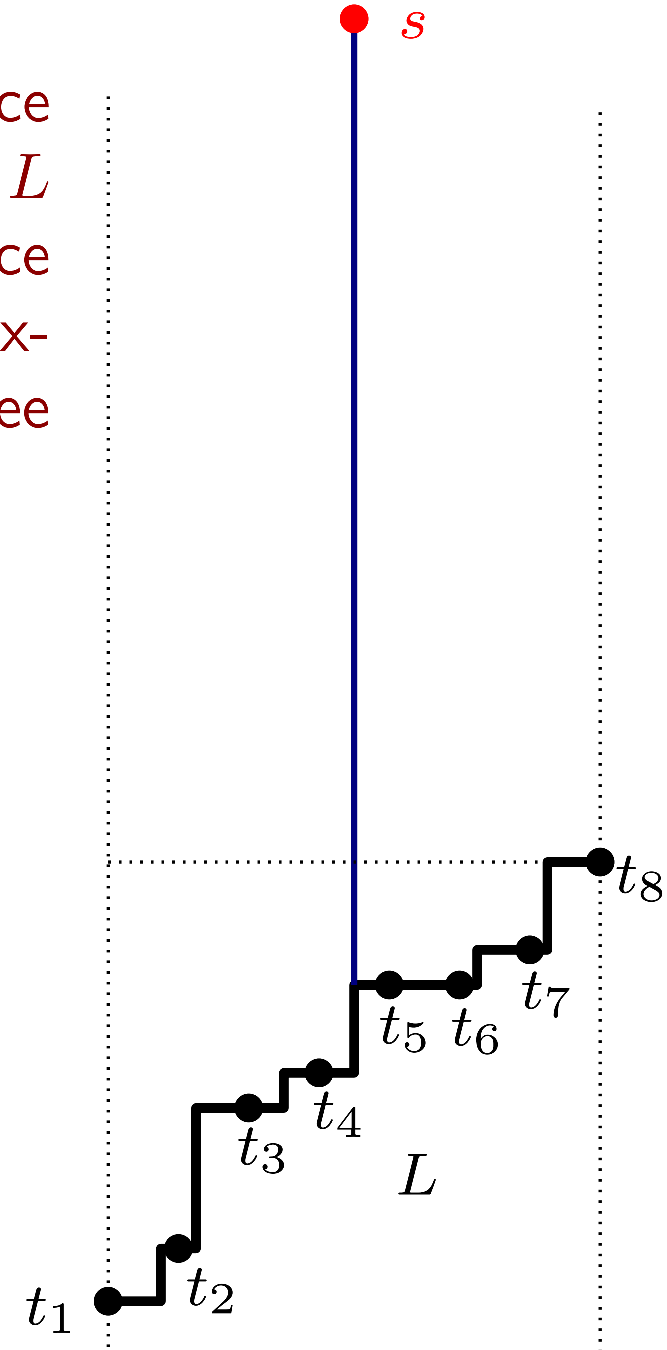

Shallow-Light Trees for Staircases

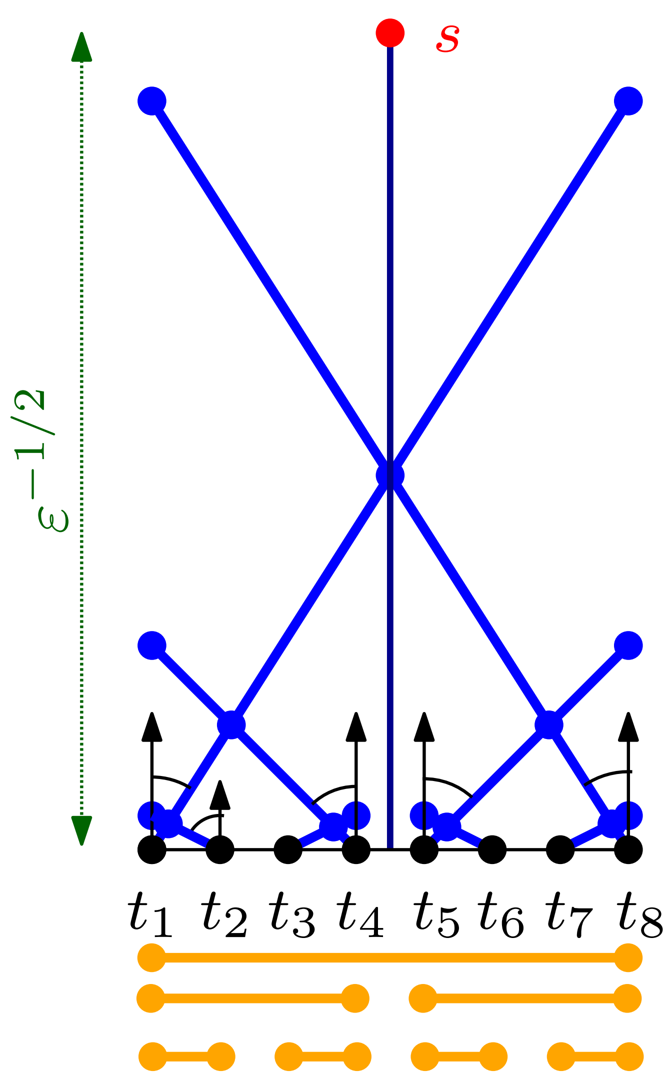

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

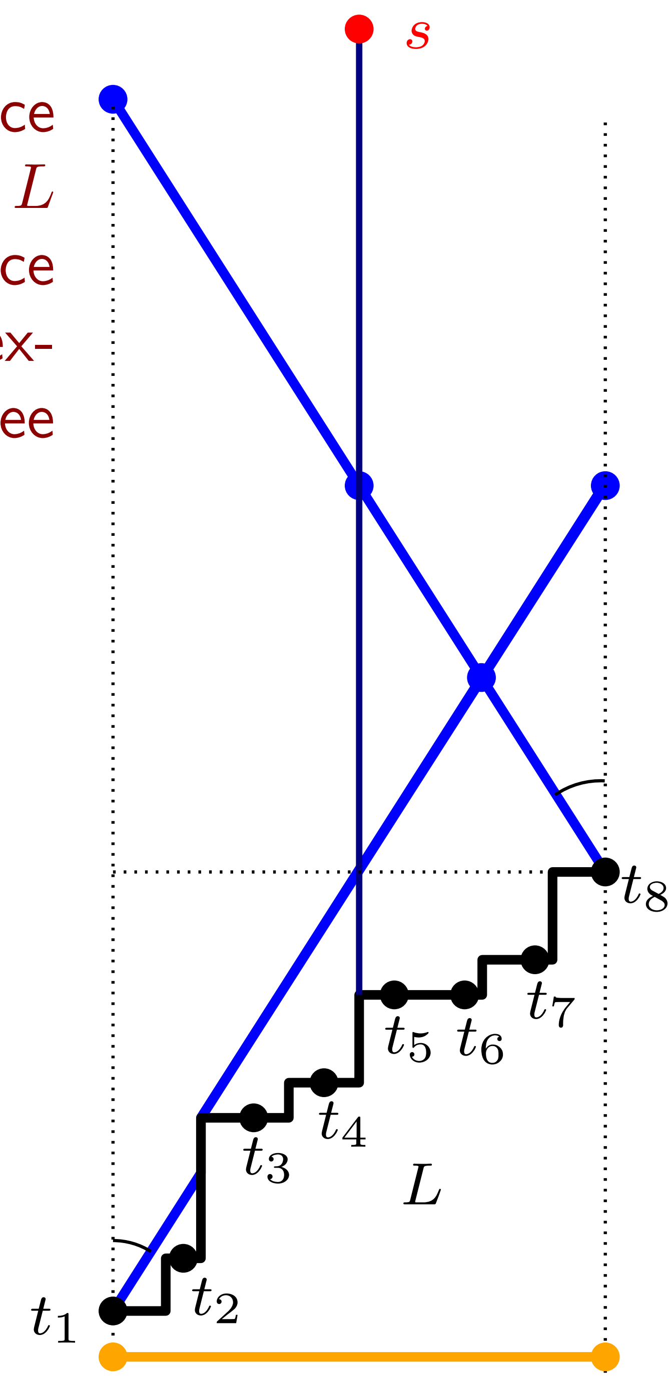

Shallow-Light Trees for Staircases

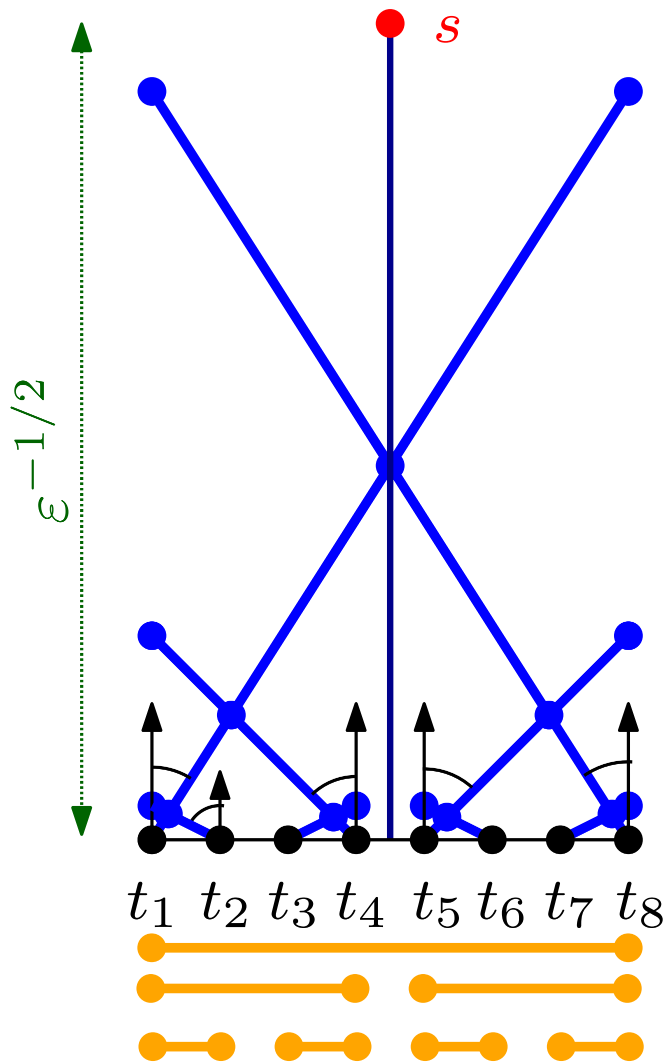

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

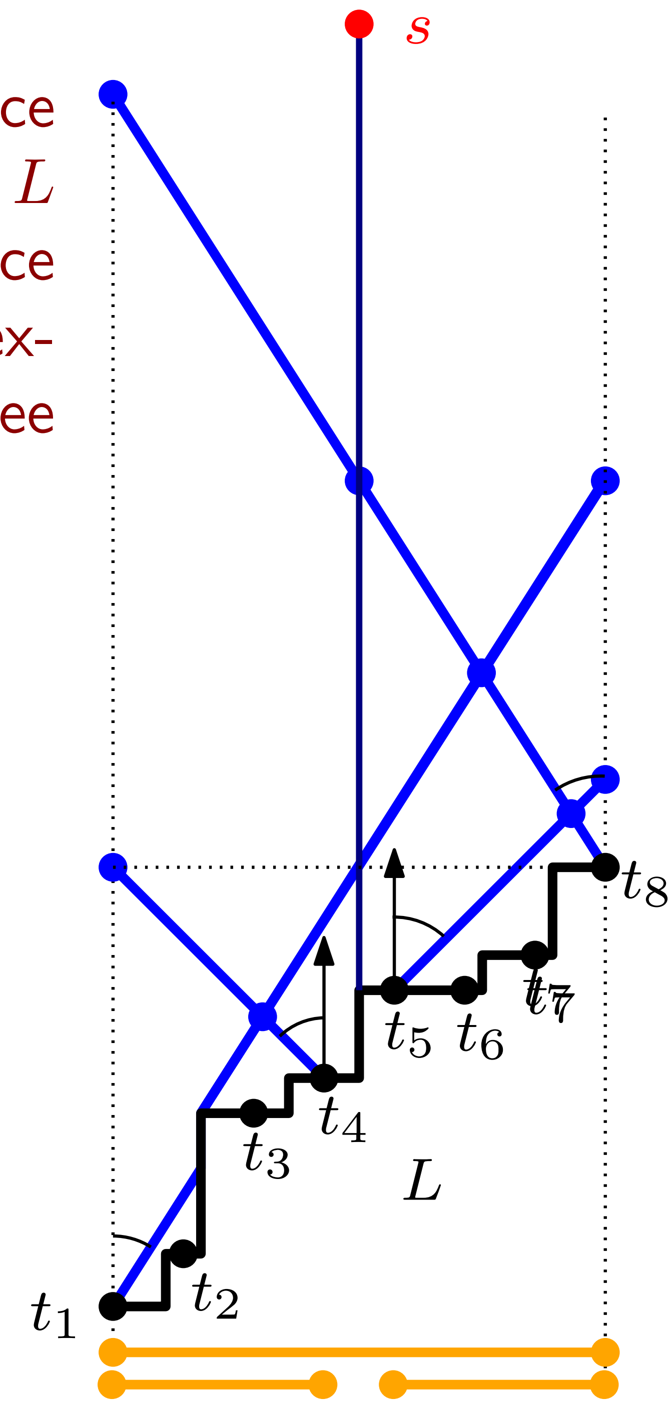

Shallow-Light Trees for Staircases

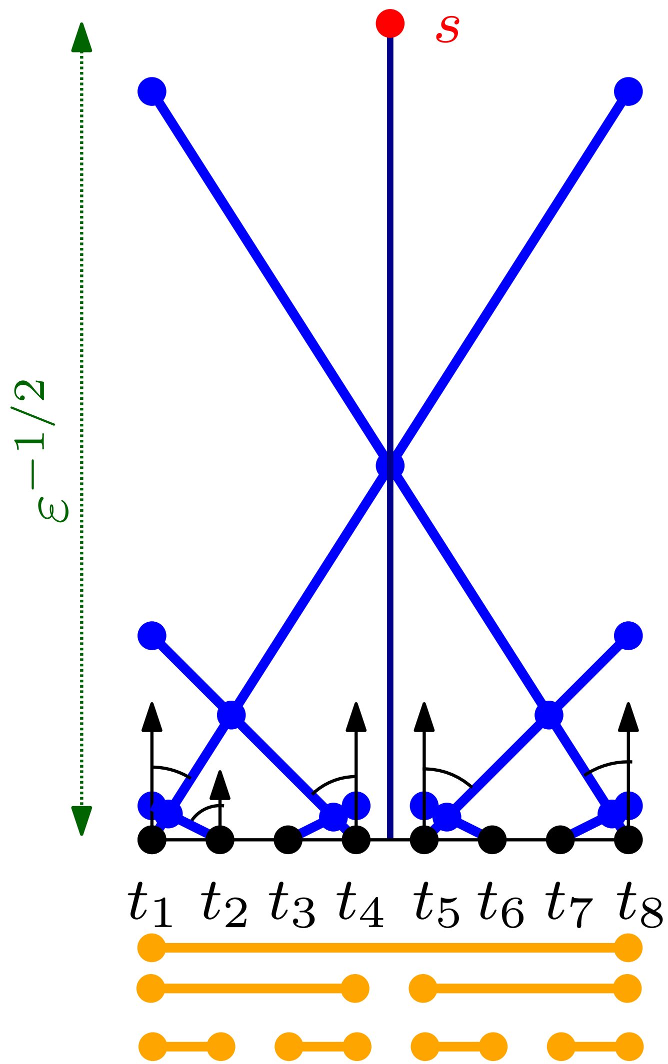

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

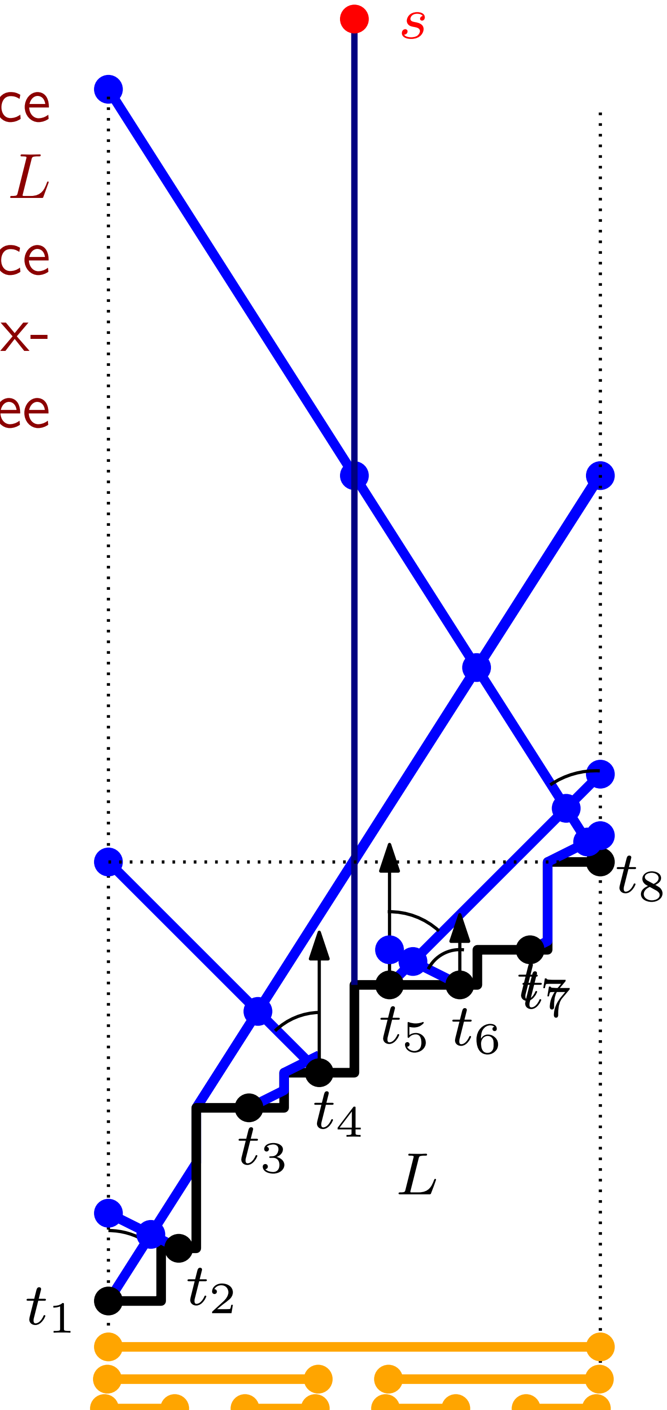

Shallow-Light Trees for Staircases

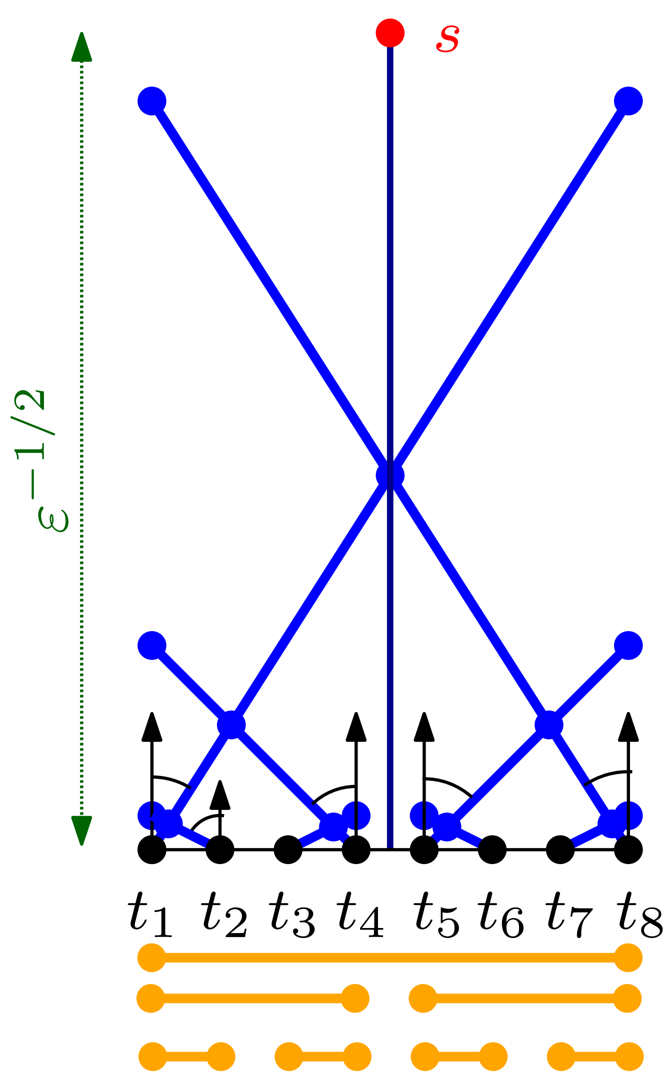

Shallow-light tree from s to a line segment L. [Solomon, JoCG'15]

Solomon: For a source s, and a line segment L of width 1 at distance $\varepsilon^{-1/2}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-1/2})$.

Lemma. For a source s, and a staircase path L of width 1 at distance $\varepsilon^{-\frac{1}{2}}$ from s, there exists a shallow-light tree of weight $O(\varepsilon^{-\frac{1}{2}} + \|L\|)$.

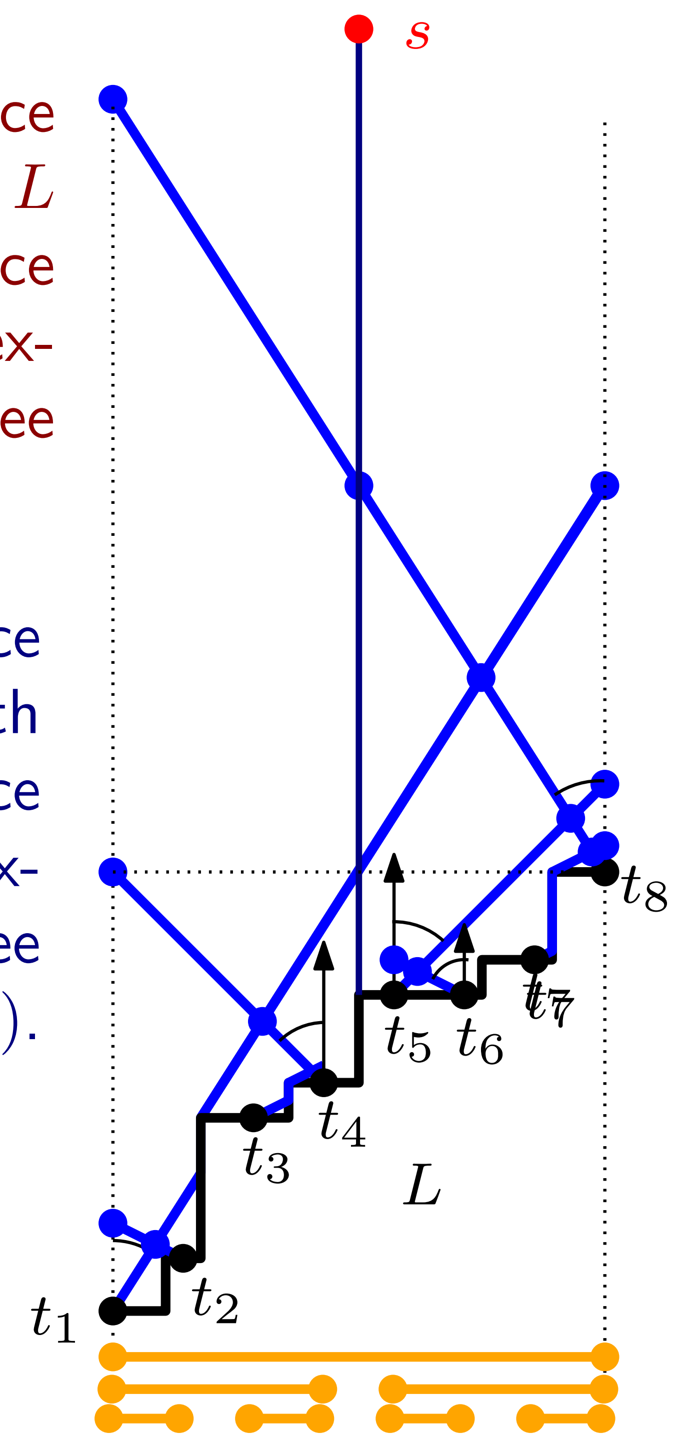

Shallow-Light Trees for Staircases

In a staircase *polygon* P, we place shallow-light trees recursively to construct a directional $(1+\varepsilon)$-spanner for all ab-pairs with $ab \subset P$.
The total weight is $O(\varepsilon^{-1}\,\mathrm{width}(P))$.

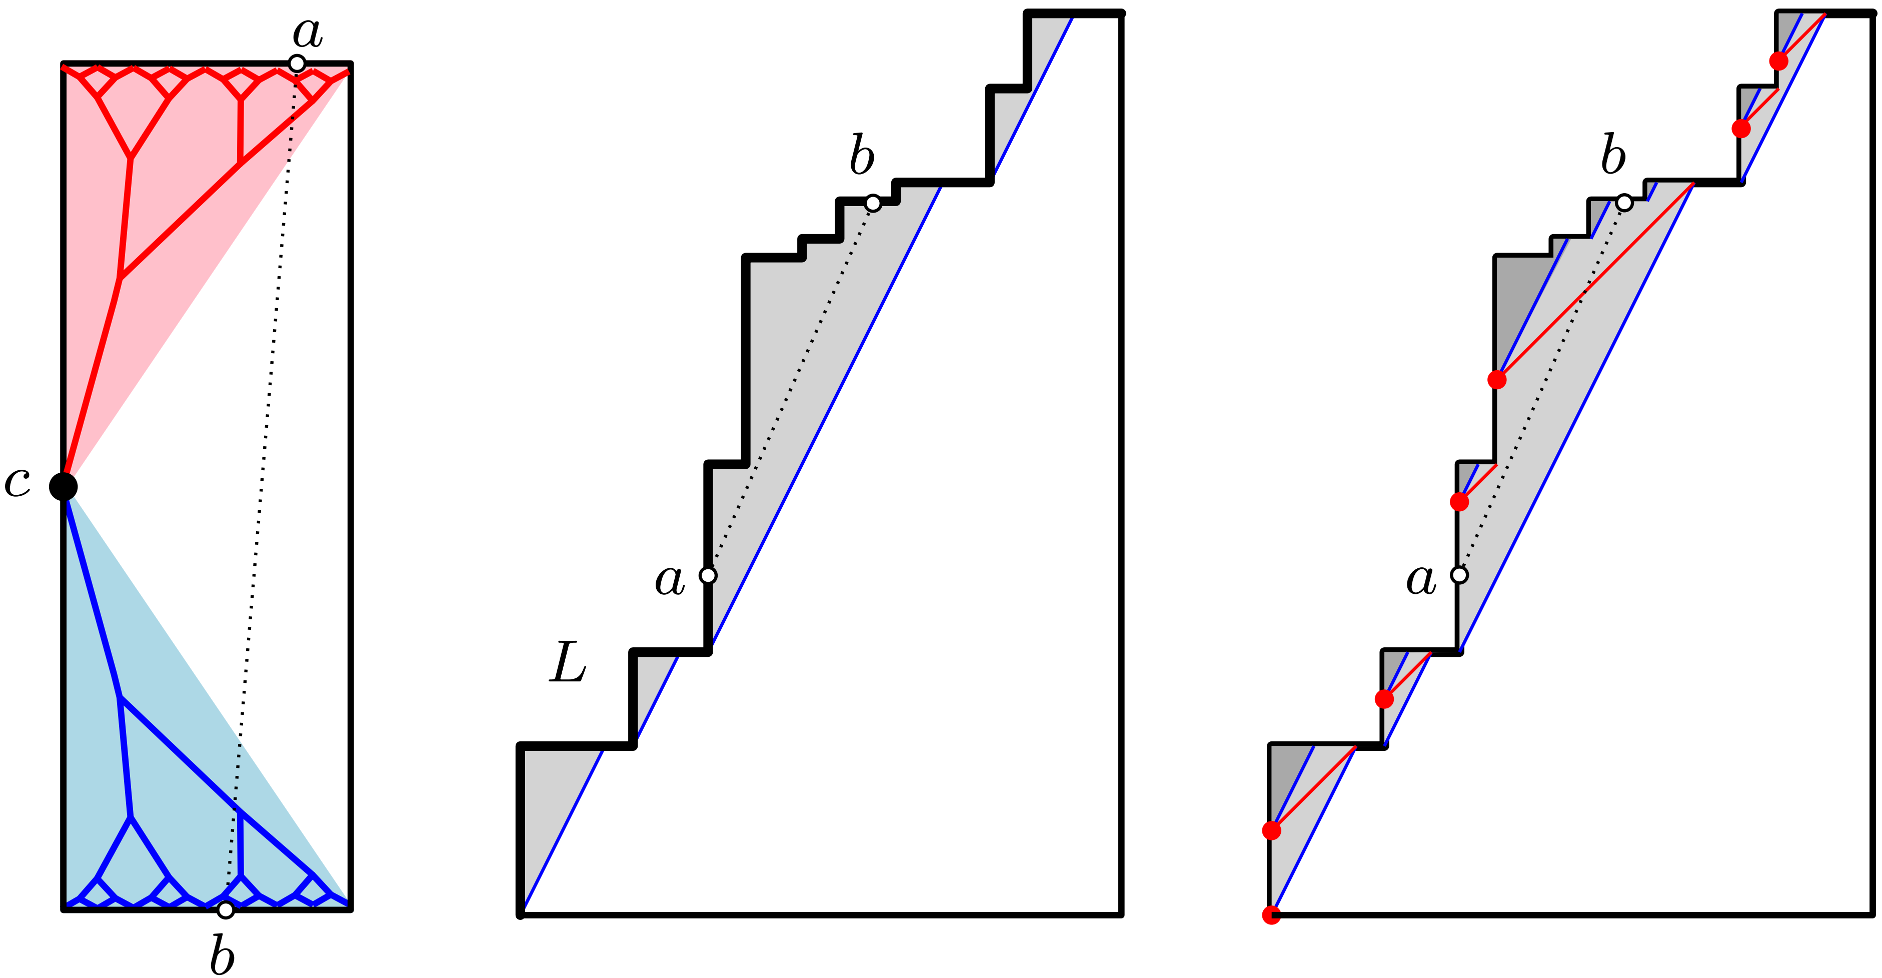

Shallow-Light Trees for Staircases

We can handle rectangle and staircase tiles!

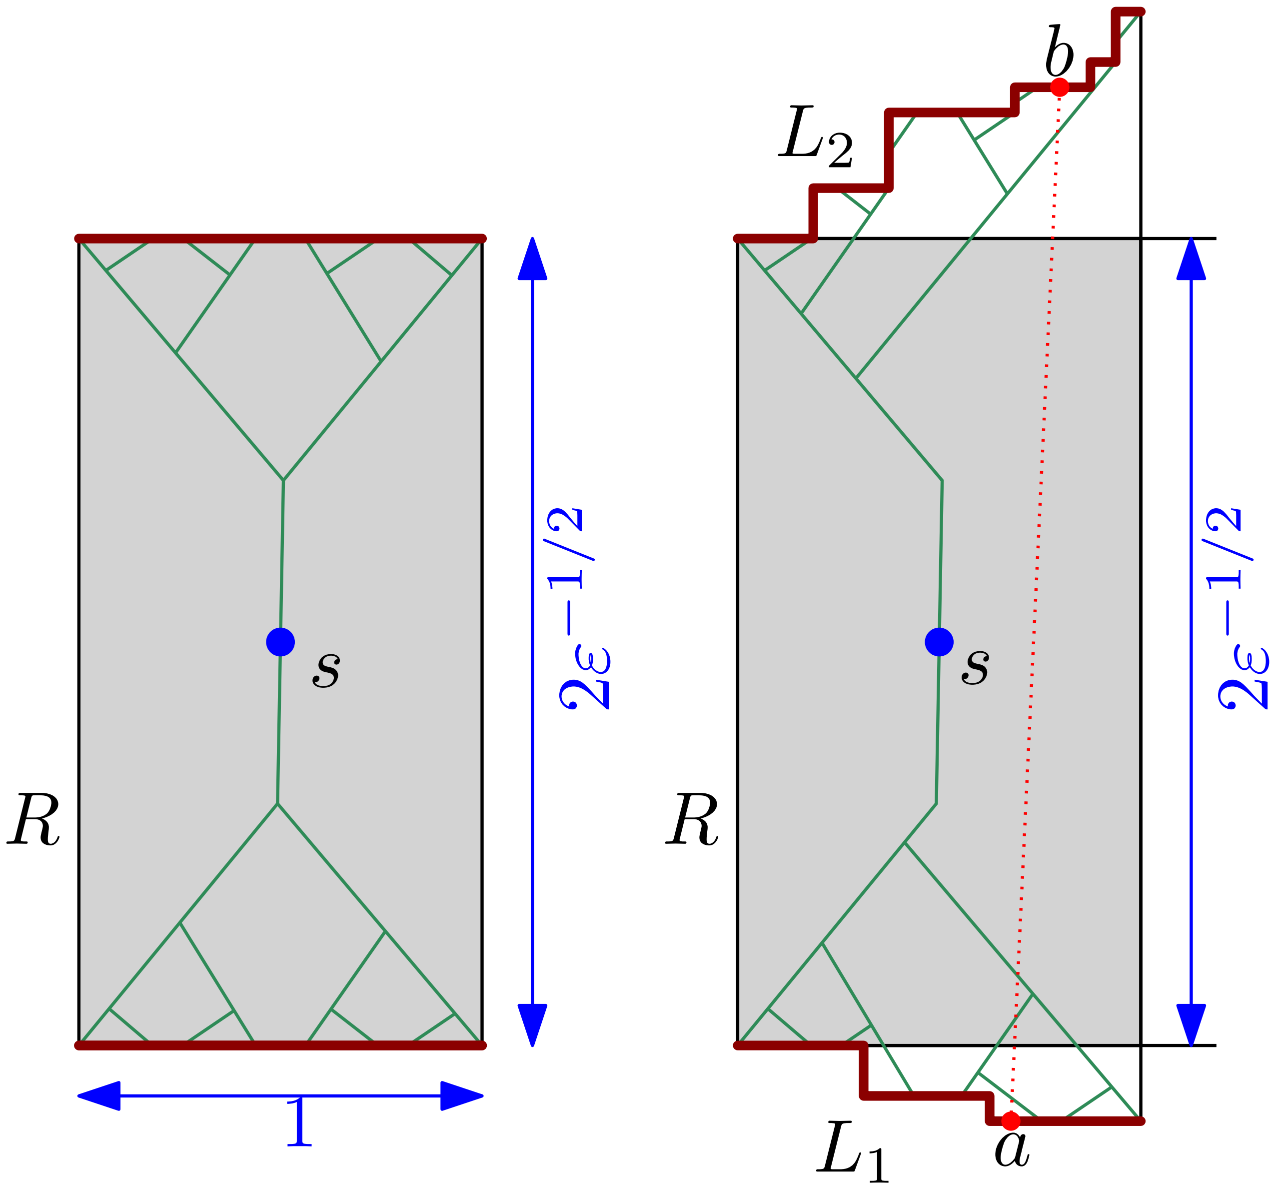

Shallow-Light Trees for Staircases

We can handle rectangle and staircase tiles!

...but not necessarily histograms...

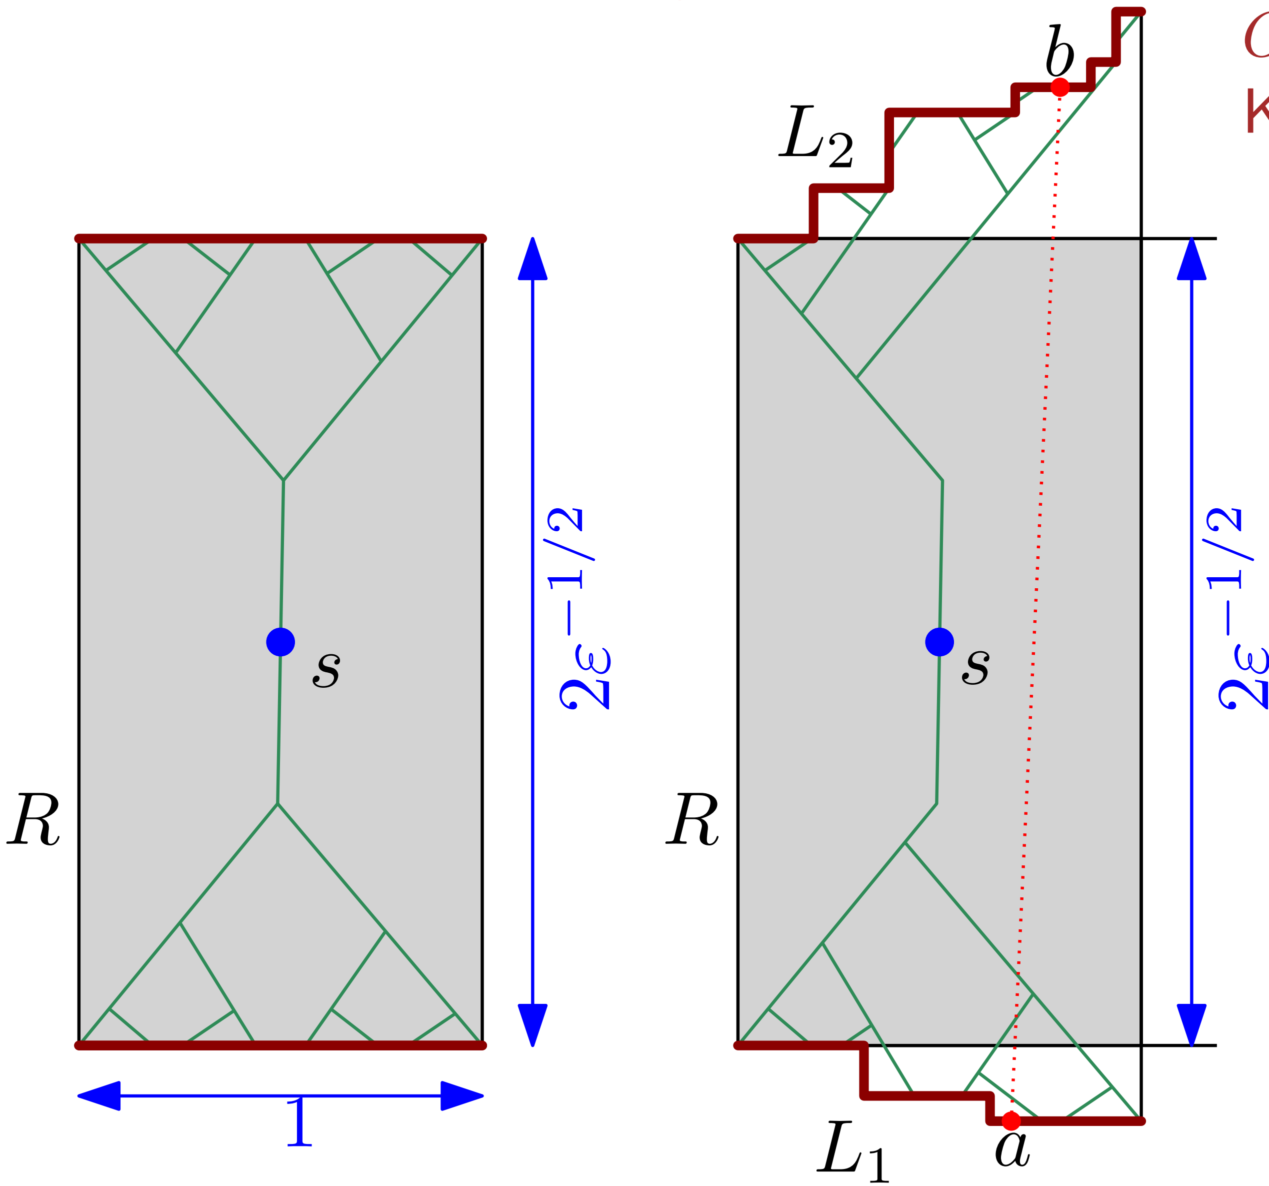

Partitioning a histogram into staircases would cost a $O(\log n)$ factor [de Berg & Kreveld, 1994]

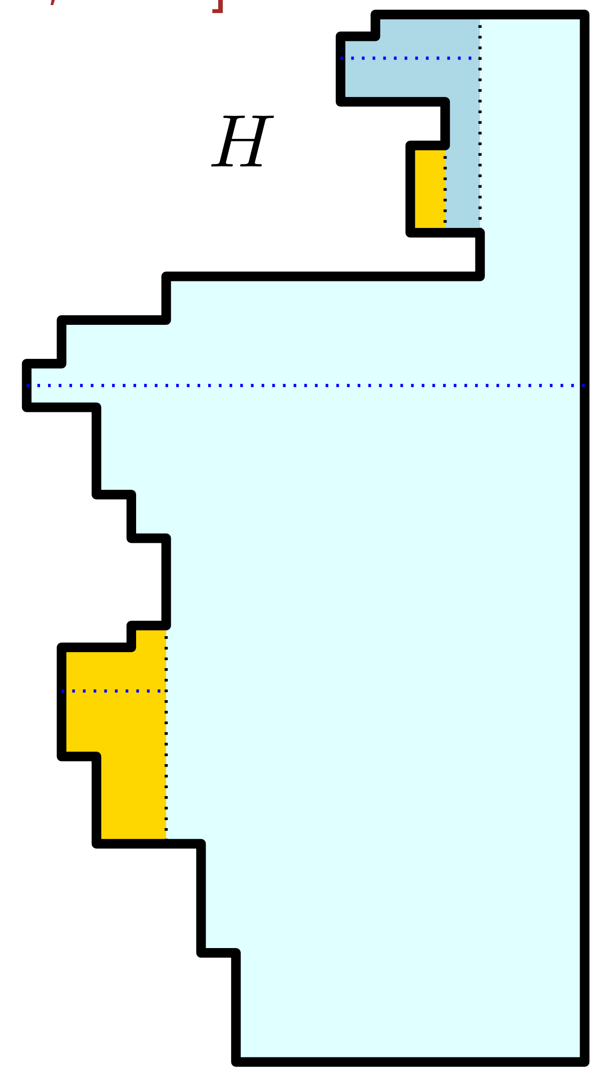

Tiling — Histograms — Window partitioning

(3) Compute the window-partition into rectilinear histograms.

modified

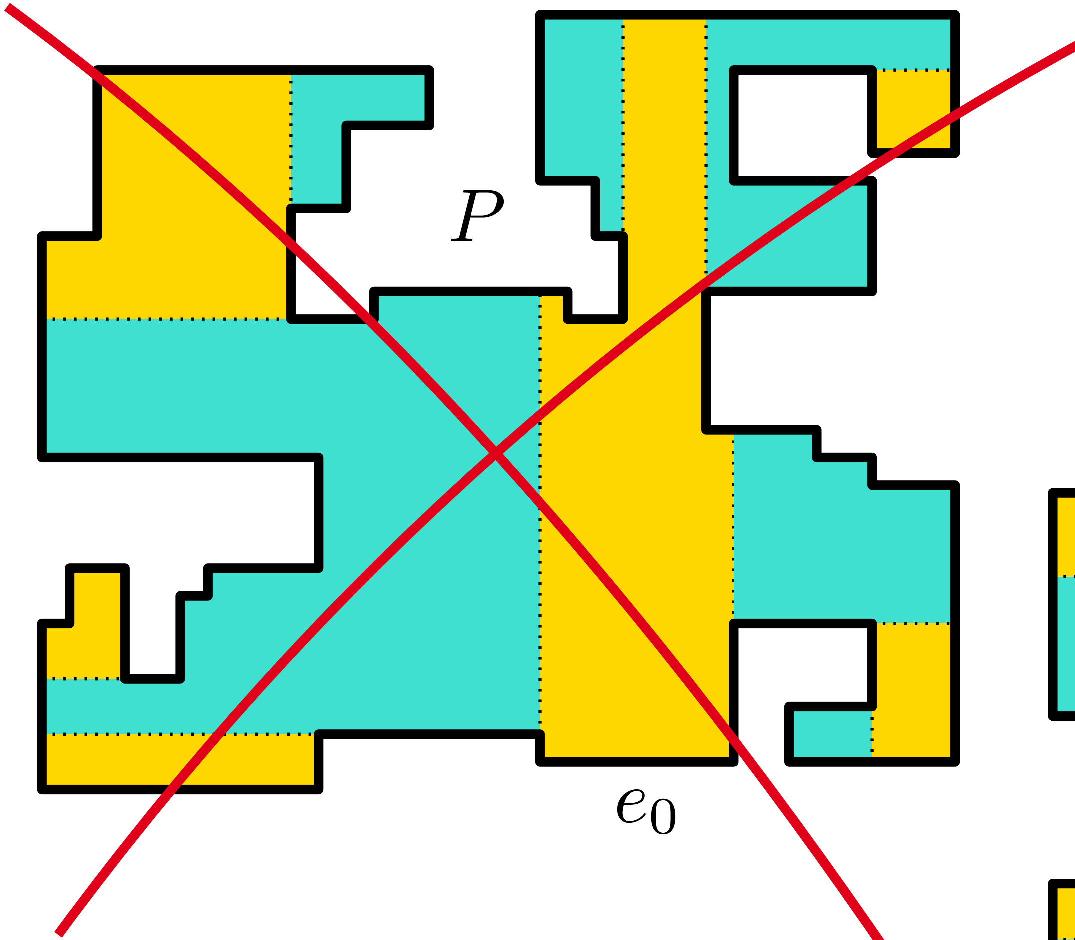

"fuzzy histograms"

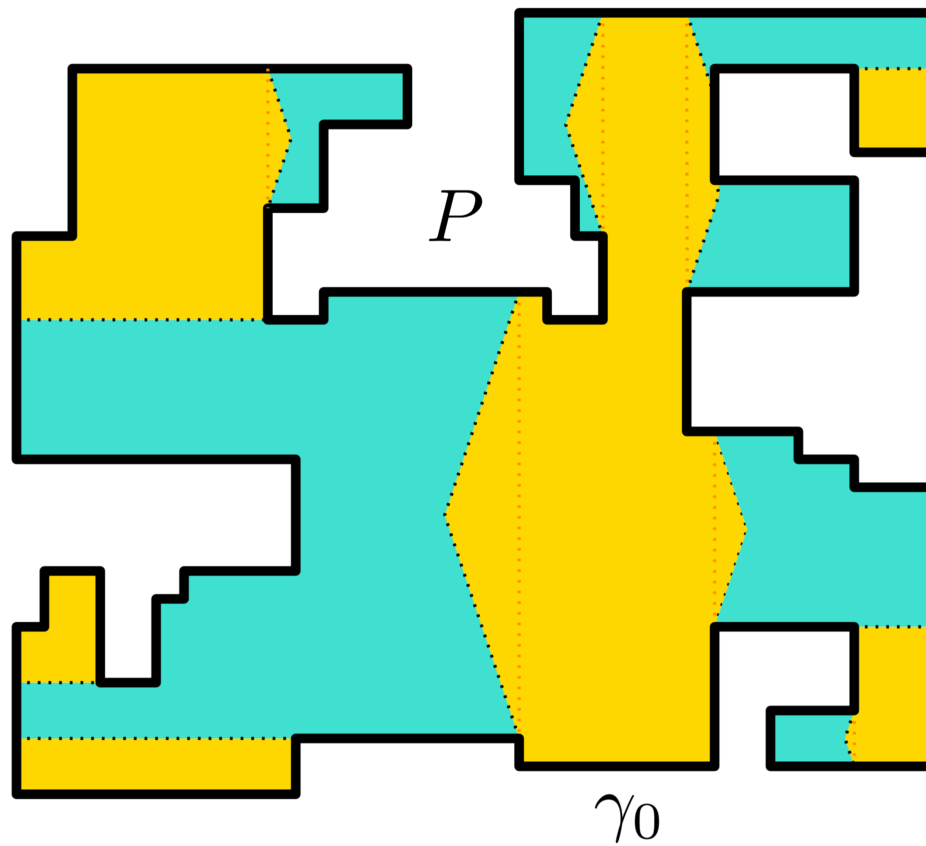

(1) Fuzzy Histograms and Staircases

A *fuzzy histogram* is a simple polygon bounded by a y-monotone rectilinear path L and a path γ of one or two edges of slopes $\pm\Lambda\varepsilon^{-1/2}$; if the latter path has two edges, then its interior vertex is a reflex vertex of the polygon.

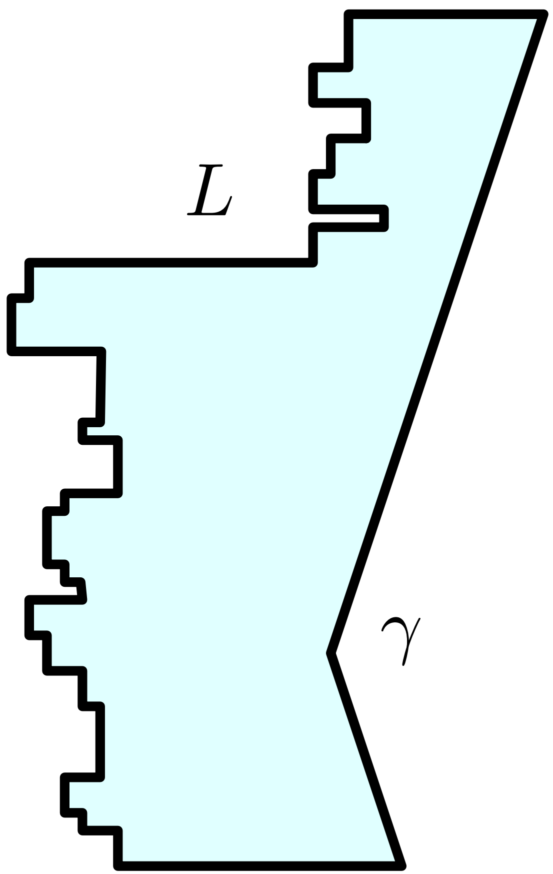

a fuzzy histogram

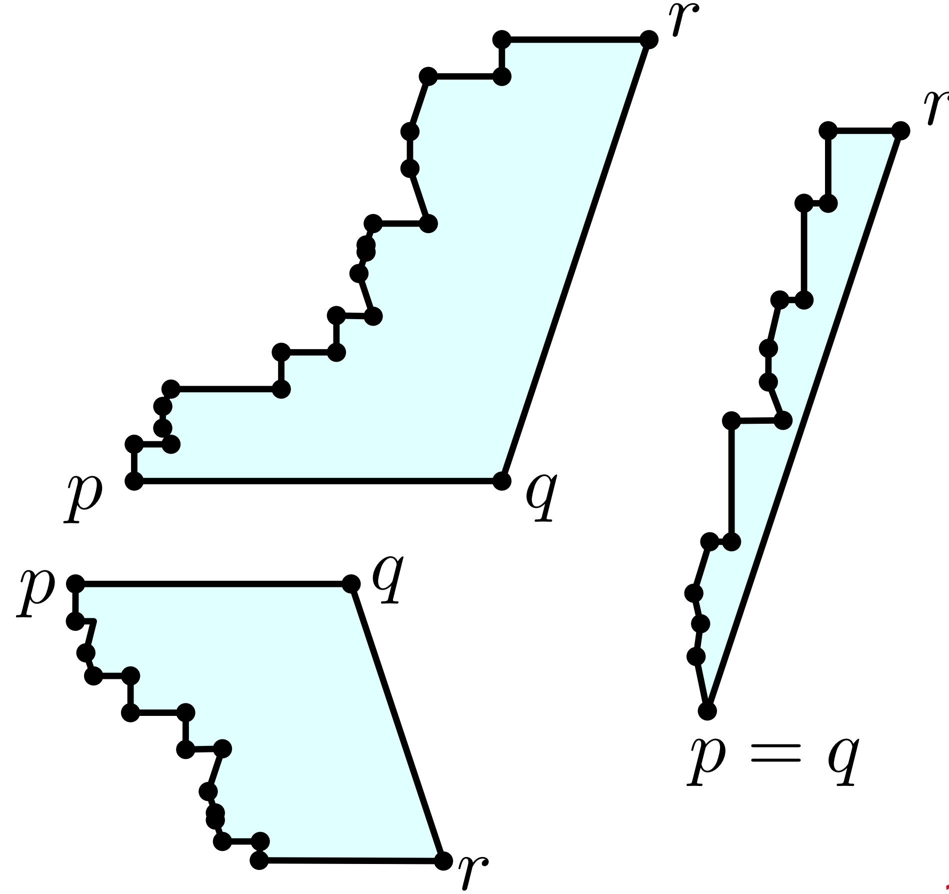

fuzzy staircases

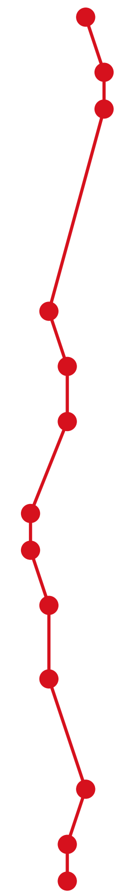

a Λ-path

(1) Fuzzy Histograms and Staircases

A *fuzzy staircase* is a simple polygon bounded by a path pqr, where pq is horizontal and $\text{slope}(qr) = \pm\Lambda\varepsilon^{-1/2}$, and a pr-path obtained from an x- and y-monotone staircase by replacing vertical edges with some Λ-paths.

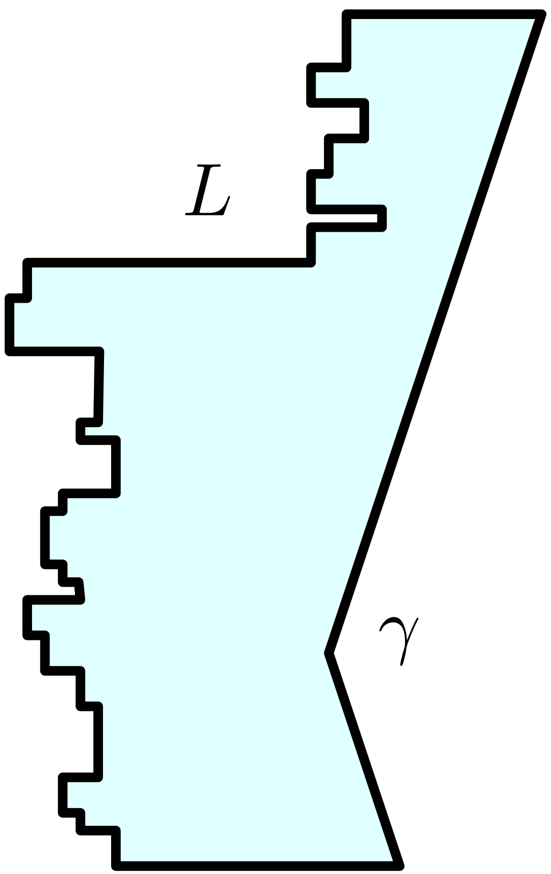

a fuzzy histogram

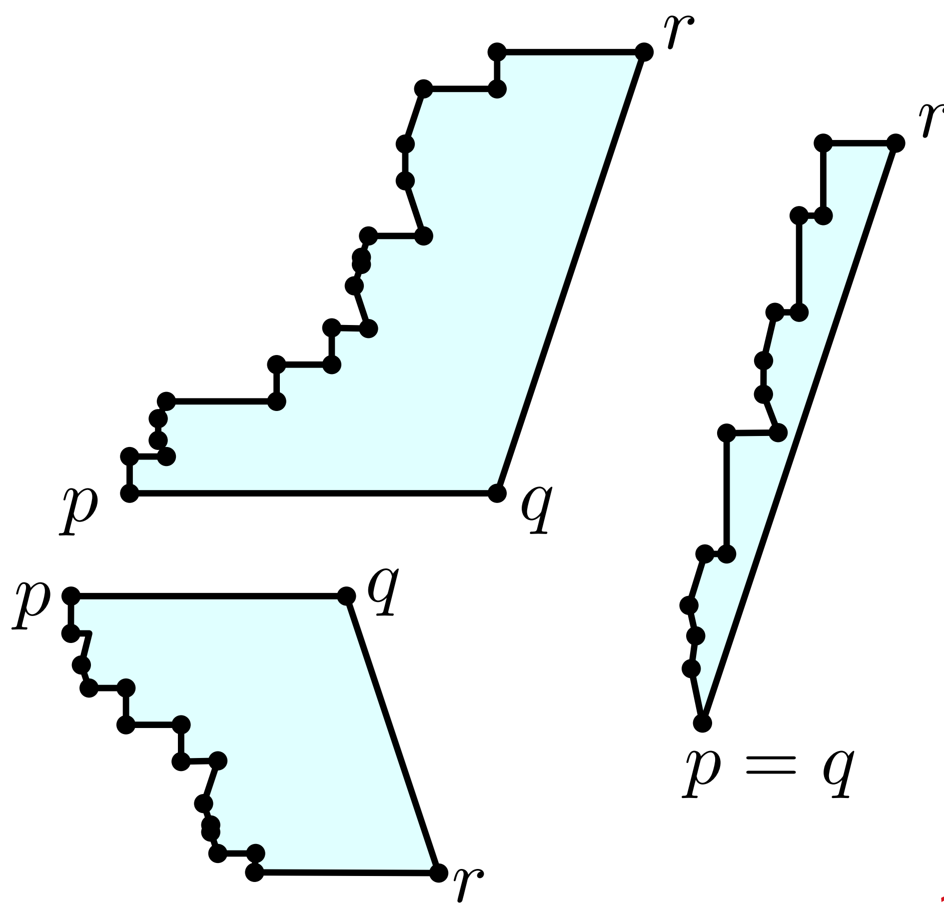

fuzzy staircases

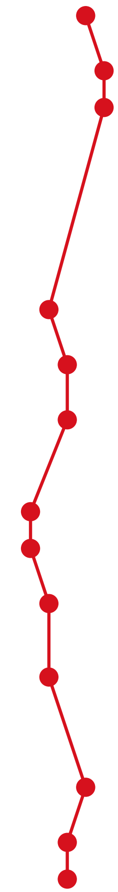

a Λ-path

(1) Partitioning Fuzzy Histograms into Fuzzy Staircases

Partitioning a fuzzy histogram into

- fuzzy staircases and
- tame histograms

increases the weight by only a constant factor (by a simple charging scheme).

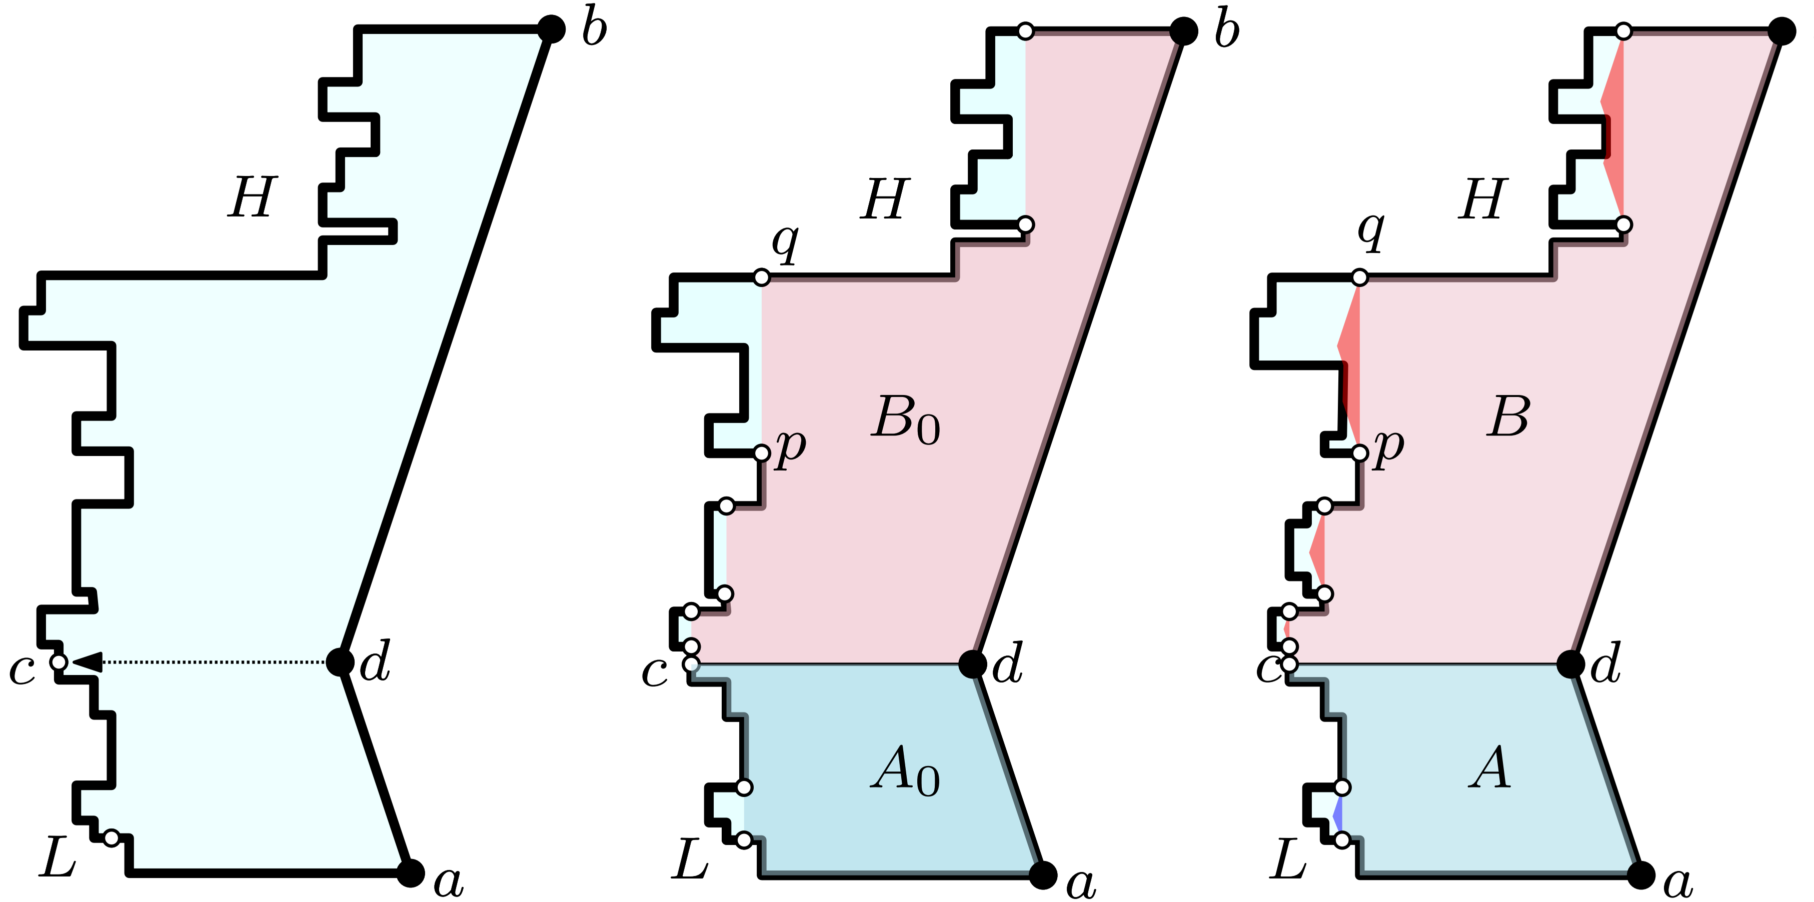

(1) Directions Spanners for Fuzzy Staircases

Generalization from staircases to fuzzy staircases:

L.: We can augment a fuzzy staircase P to a geometric graph of weight $O(\mathrm{per}(P) + \varepsilon^{-1/2}\mathrm{hper}(P))$ that contains, for all $a, b \in \partial P$, a path of weight at most $(1 + O(\varepsilon))\|ab\|$.

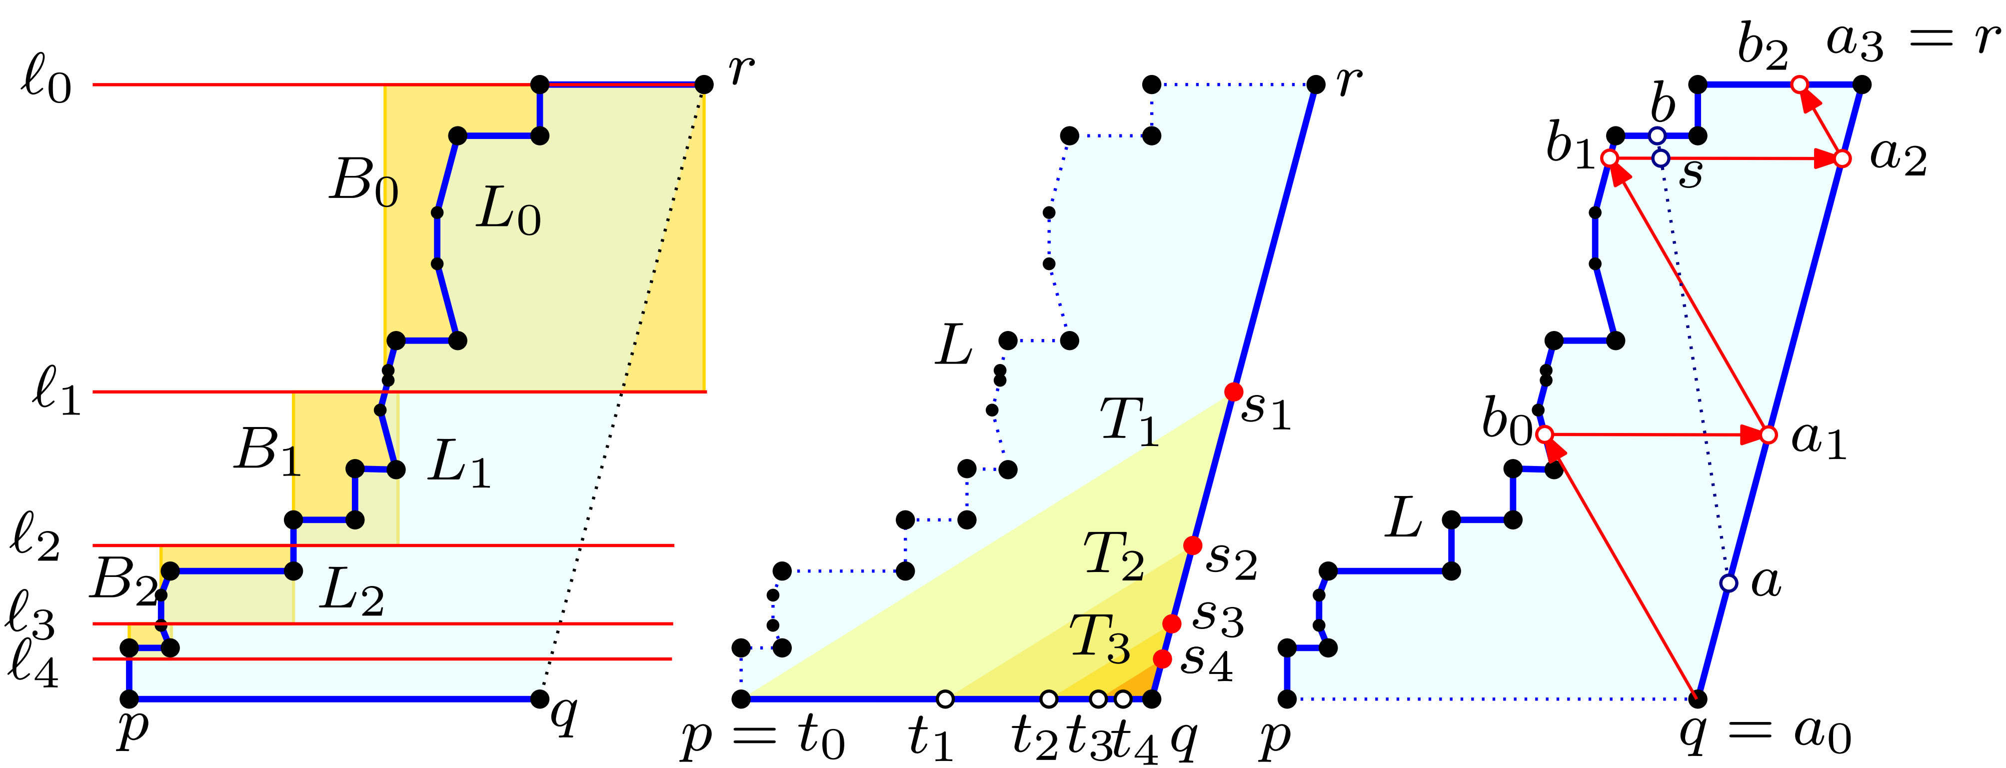

(2) y-Monotone Λ-Histograms

A Λ-*histogram* is a simple polygon obtained from a histogram by replacing each vertical edge with some Λ-path, in which every edge is vertical, or has slope $\pm\Lambda\varepsilon^{-1/2}$, for a constant $\Lambda > 0$.

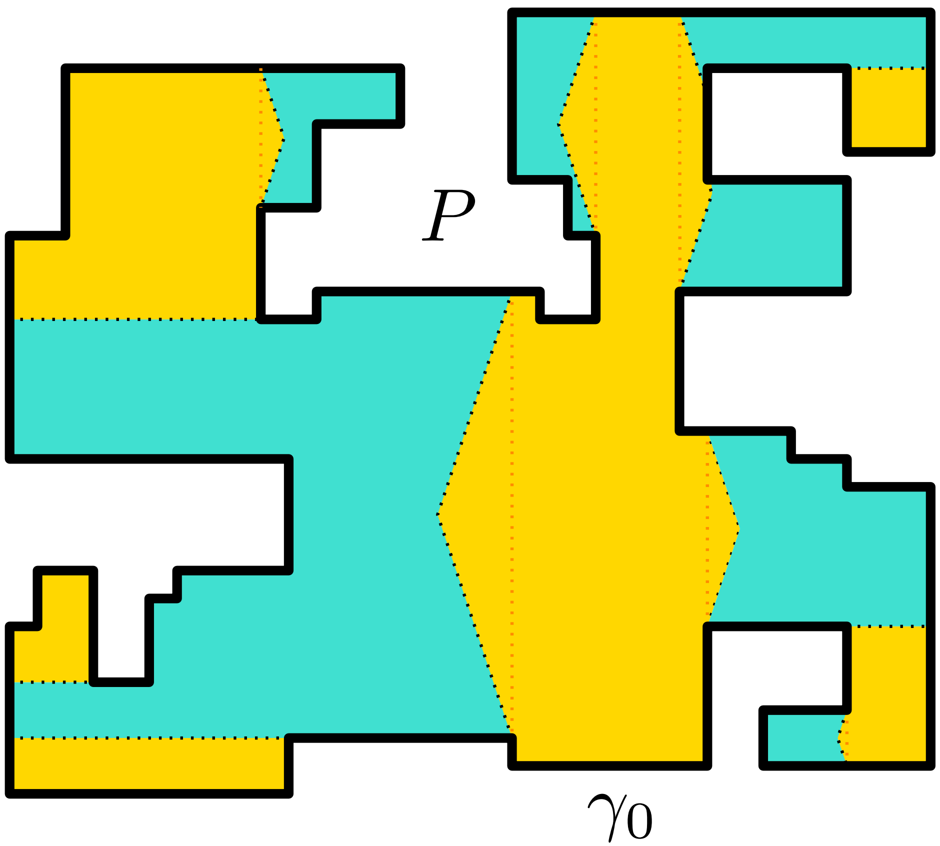

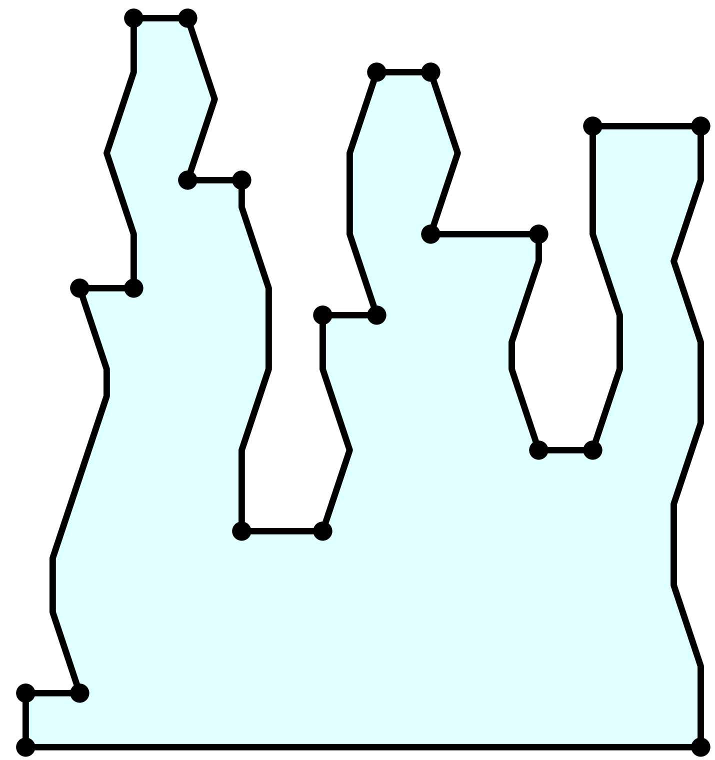

a y-monotone Λ-histogram

(2) Partitioning Λ-Histograms into Tame Histograms

A *tame histogram* is a simple polygon bounded by a horizontal line segment pq and a pq-path that consists of ascending or descending Λ-paths and x-monotone increasing horizontal edges s.t.:
(i) there is no chord between interior points of any two ascending (resp., two descending) Λ-paths; and
(ii) for every horizontal chord ab, with $a, b \in L$, the subpath L_{ab} of L between a and b satisfies $\|L_{ab}\| \leq 2\|ab\|$.

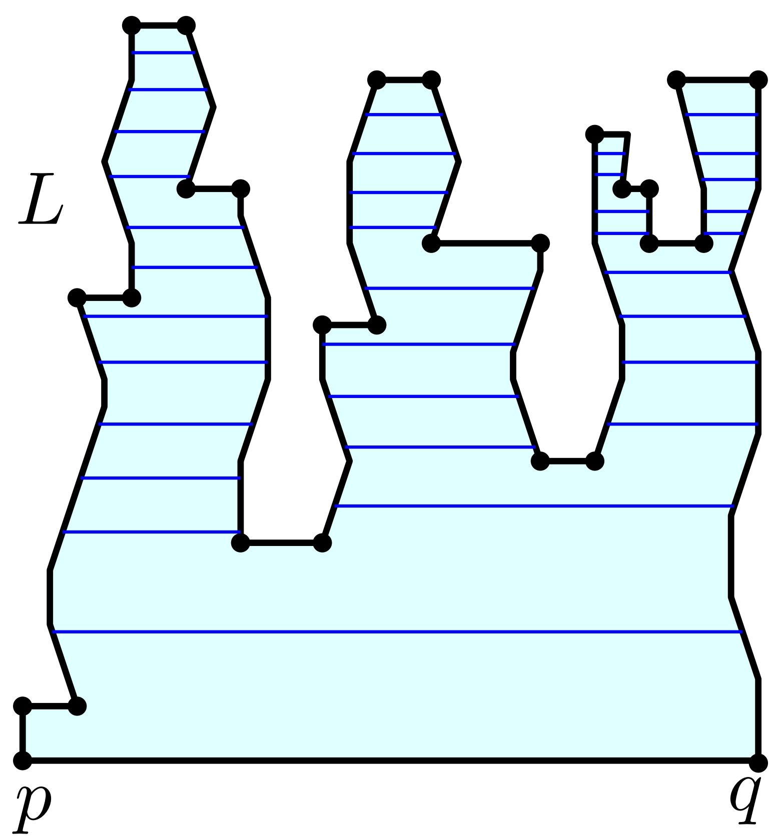

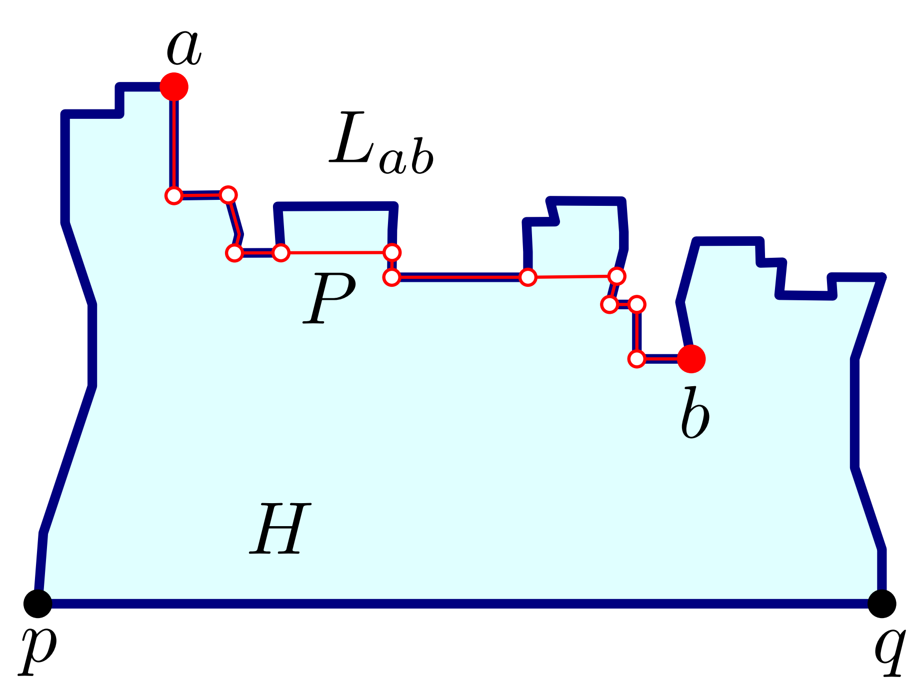

Shallow-light Trees for Tame Histograms

The SLT construction generalizes to tame histograms:

L.: We can augment a tame histogram ∂P to a geometric graph of weight $O(\varepsilon^{-1/2}\text{hper}(P))$ that contains, for all $a, b \in \partial P$, a path of weight at most $(1 + O(\varepsilon))\|ab\|$.

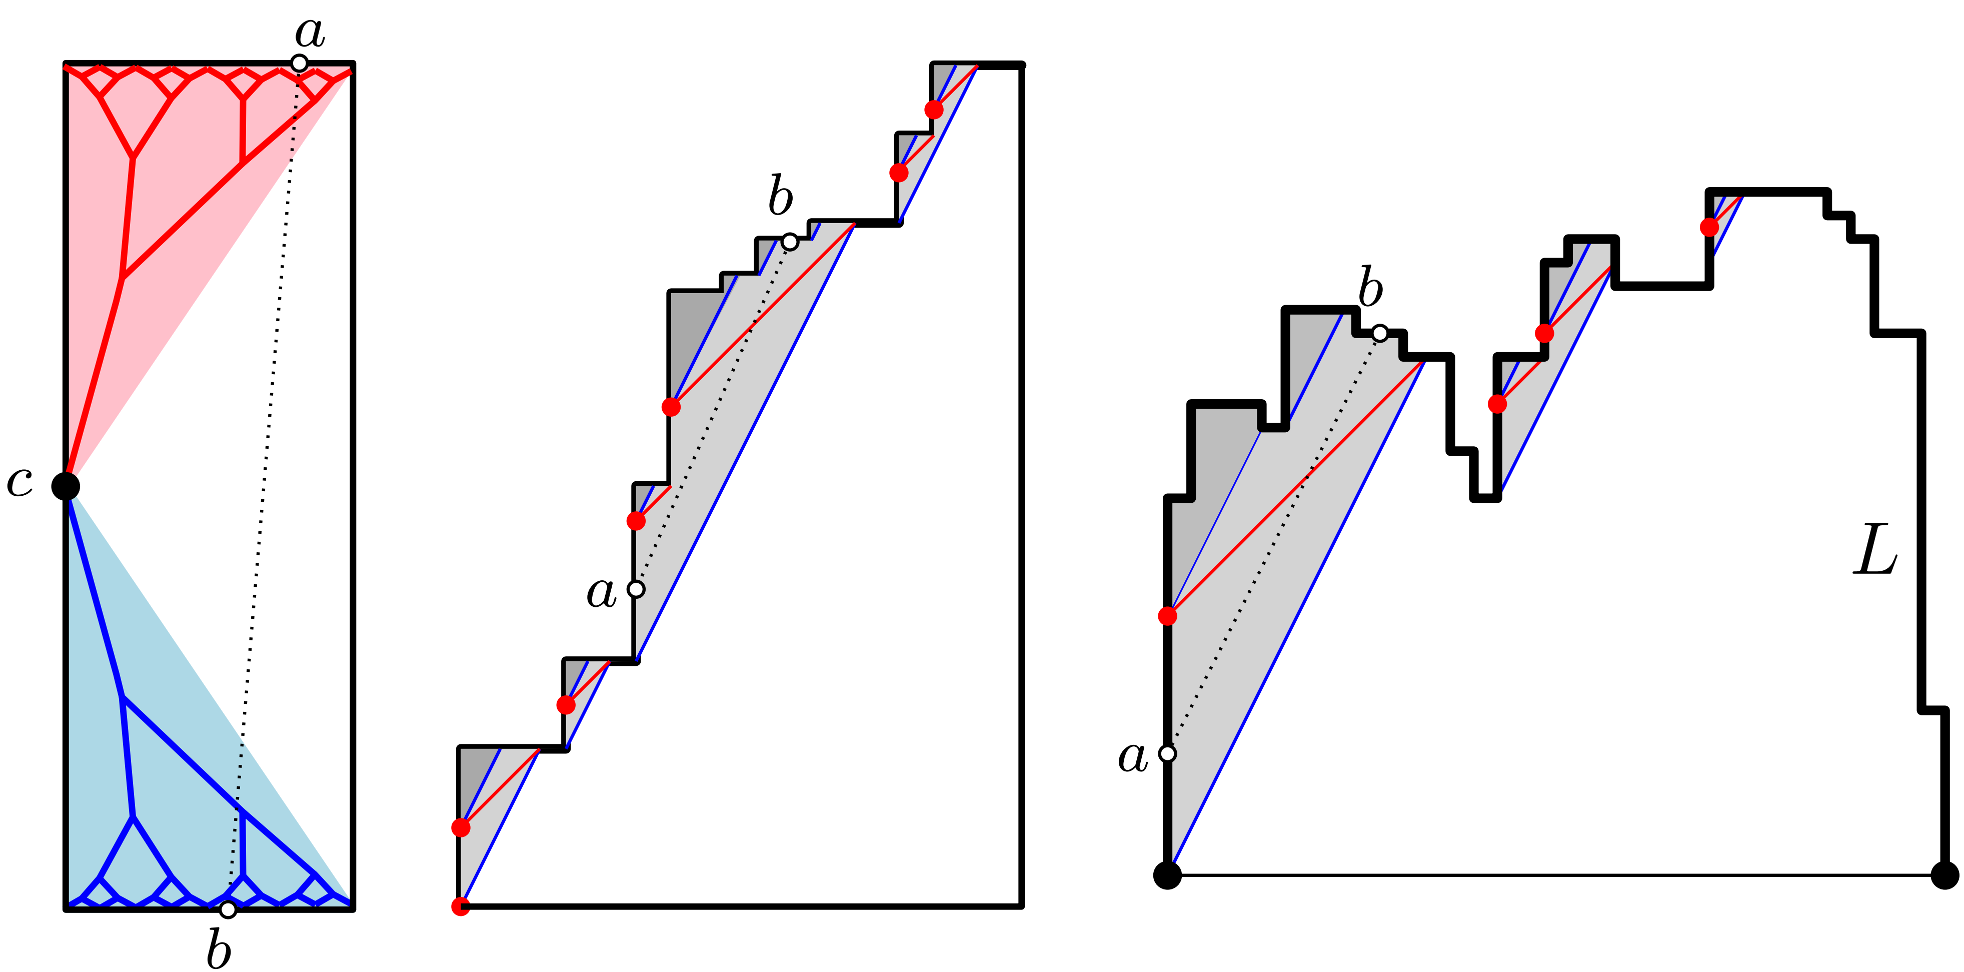

Tame histograms

The construction generalizes to tame histogram:
We can augment a tame histogram ∂P to a geometric graph of weight $O(\varepsilon^{-1/2}\text{hper}(P))$ that contains, for all $a, b \in \partial P$, a path of weight at most $(1 + O(\varepsilon))\|ab\|$.

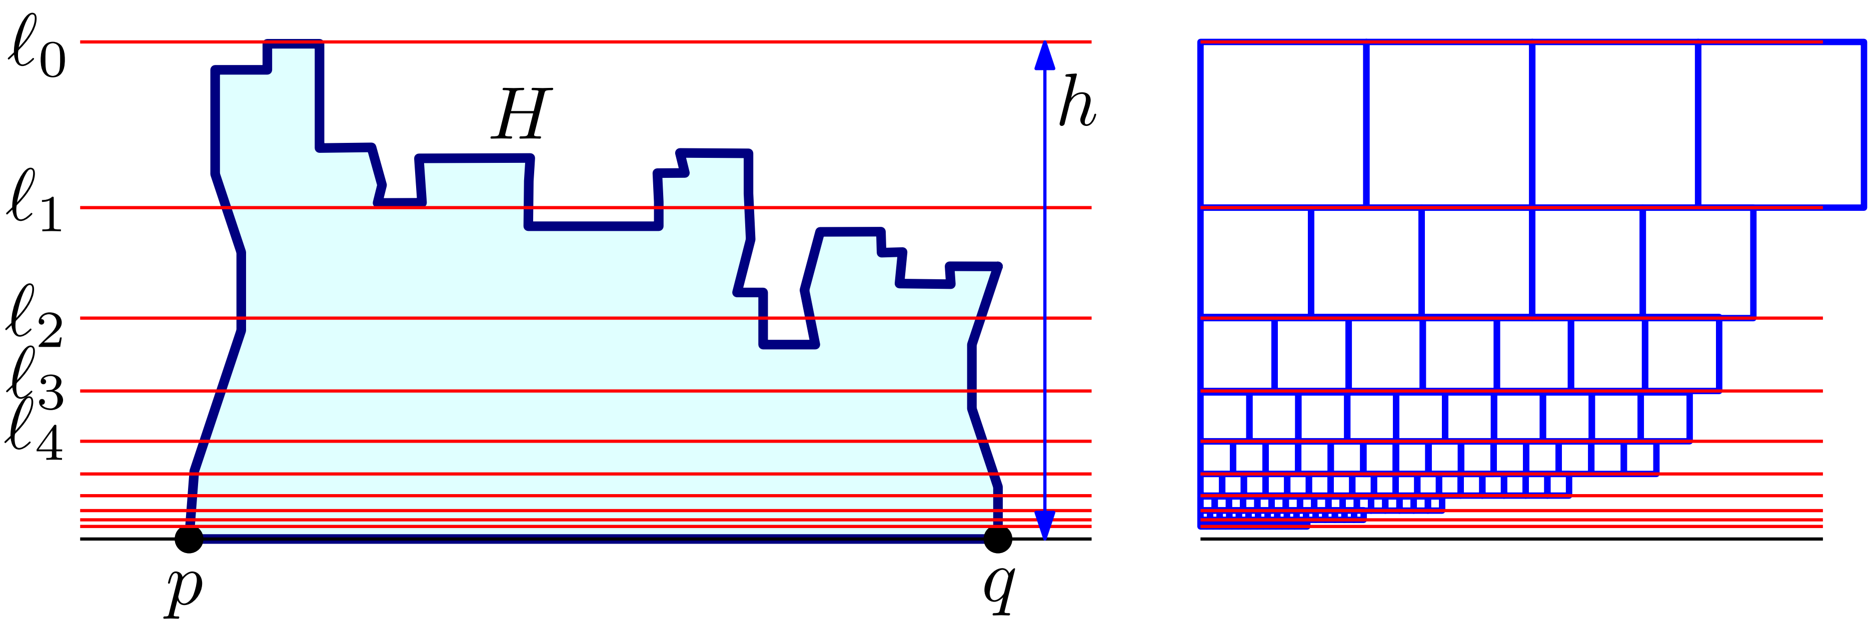

Upper bound - Theorem

Theorem. For every set S of n points in Euclidean plane, there exists a Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$.

Lemma. We can subdivide a (weakly) simple rectilinear polygon P into a collection $\mathcal{F}$ of fuzzy staircases and tame histograms of total perimeter $\sum_{F\in\mathcal{F}} \mathrm{per}(F) \leq O(\varepsilon^{-1/2}\mathrm{per}(P))$ and total horizontal perimeter $\sum_{F\in\mathcal{F}} \mathrm{hper}(F) \leq O(\mathrm{per}(P))$.

Lemma. Let F be a fuzzy staircase or a tame histogram, $S \subset \partial F$ a finite point set, $\varepsilon > 0$, and $D = [\frac{\pi-\sqrt{\varepsilon}}{2}, \frac{\pi+\sqrt{\varepsilon}}{2}]$ an interval of nearly vertical directions. Then there exists a geometric graph of weight

$$O(\mathrm{per}(F) + \varepsilon^{-1/2}\,\mathrm{hper}(F))$$

such that for all $a, b \in S$, if ab is a chord of F and $\mathrm{dir}(ab) \in D$, then G contains an ab-path of weight at most $(1+O(\varepsilon))\|ab\|$.

Summary

	Sparsity	Lightness
Lower Bound	■ $\Omega(\varepsilon^{-1/2} / \log \varepsilon^{-1})$ [Le&Solomon, FOCS'19] ■ $\Omega(\varepsilon^{(1-d)/2})$ [Bhore &Tóth, SIDMA'22]	■ $\Omega(\varepsilon^{-1} / \log \varepsilon^{-1})$, for $d = 2$ [Le&Solomon, FOCS'19] ■ $\Omega(\varepsilon^{-d/2})$ [Bhore &Tóth, SIDMA'22]
Upper Bound	■ $O(\varepsilon^{(1-d)/2})$ for d-space [Le & Solomon, FOCS'19]	■ $\Omega(\varepsilon^{-1} \log \Delta)$, for $d = 2$ [Le & Solomon, ESA'20] ■ $\tilde{O}(\varepsilon^{-(d+1)/2})$, for $d \geq 3$ [Le & Solomon, ArXiv'20] ■ $O(\varepsilon^{-1})$, for $d = 2$ [Bhore &Tóth, SoCG'21]

- All bounds are for Euclidean Steiner $(1 + \varepsilon)$-spanners

Future directions.

- Question: Does there exist Euclidean Steiner $(1+\varepsilon)$-spanners for a finite set of points in $\mathbb{R}^d$, of lightness $O(\varepsilon^{-d/2})$, for $d \geq 3$?

The *Steiner ratio for Euclidean $(1+\varepsilon)$-spanners* in $\mathbb{R}^d$ is the supremum ratio between the min-weight $(1+\varepsilon)$-spanners and the min-weight Steiner $(1+\varepsilon)$-spanners over all finite point sets $S \subset \mathbb{R}^d$.

Future directions.

- Question: Does there exist Euclidean Steiner $(1+\varepsilon)$-spanners for a finite set of points in $\mathbb{R}^d$, of lightness $O(\varepsilon^{-d/2})$, for $d \geq 3$?

The *Steiner ratio for Euclidean $(1+\varepsilon)$-spanners* in $\mathbb{R}^d$ is the supremum ratio between the min-weight $(1+\varepsilon)$-spanners and the min-weight Steiner $(1+\varepsilon)$-spanners over all finite point sets $S \subset \mathbb{R}^d$.

- **Conjecture:** A Euclidean Steiner $(1+\varepsilon)$-spanner cannot simultaneously attain both lower bounds, that is, both $O(\varepsilon^{-1})$ lightness and $O(\varepsilon^{-1/2})$ sparsity in Euclidean plane.
- Explore the trade-offs between lightness and sparsity in $\mathbb{R}^d$.

Future directions.

- Question: Does there exist Euclidean Steiner $(1+\varepsilon)$-spanners for a finite set of points in $\mathbb{R}^d$, of lightness $O(\varepsilon^{-d/2})$, for $d \geq 3$?

The *Steiner ratio for Euclidean $(1+\varepsilon)$-spanners* in $\mathbb{R}^d$ is the supremum ratio between the min-weight $(1+\varepsilon)$-spanners and the min-weight Steiner $(1+\varepsilon)$-spanners over all finite point sets $S \subset \mathbb{R}^d$.

- **Conjecture:** A Euclidean Steiner $(1+\varepsilon)$-spanner cannot simultaneously attain both lower bounds, that is, both $O(\varepsilon^{-1})$ lightness and $O(\varepsilon^{-1/2})$ sparsity in Euclidean plane.

- Explore the trade-offs between lightness and sparsity in $\mathbb{R}^d$.

– Every Steiner spanner can be converted into a plane spanner.

– A simple planarization procedure could cost lots (quadratic number of) of Steiner points.

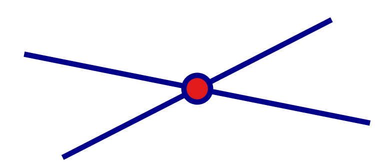

- Question: Bound the sparsity of a plane Steiner $(1+\varepsilon)$-spanner for n points in Euclidean plane, as a function of n and ε.

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

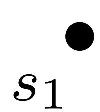

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

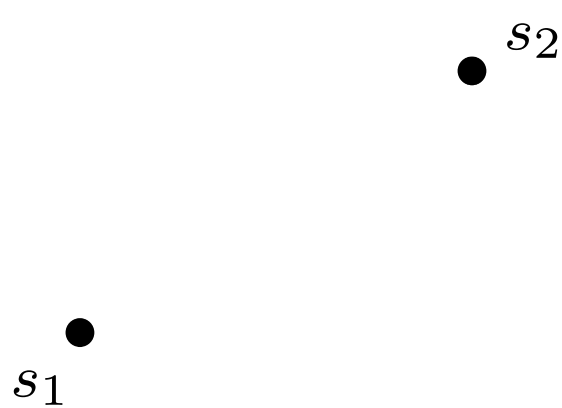

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

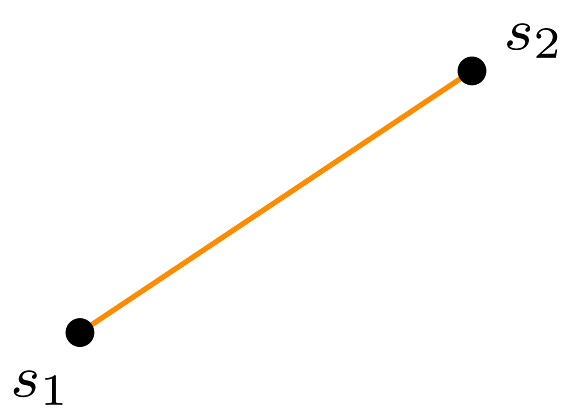

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

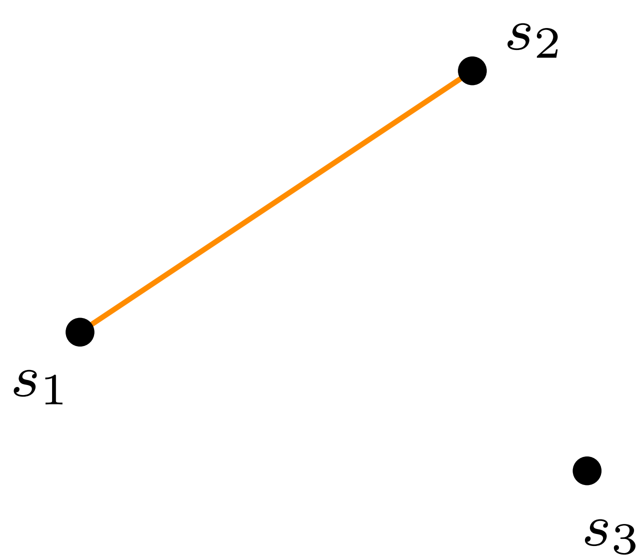

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

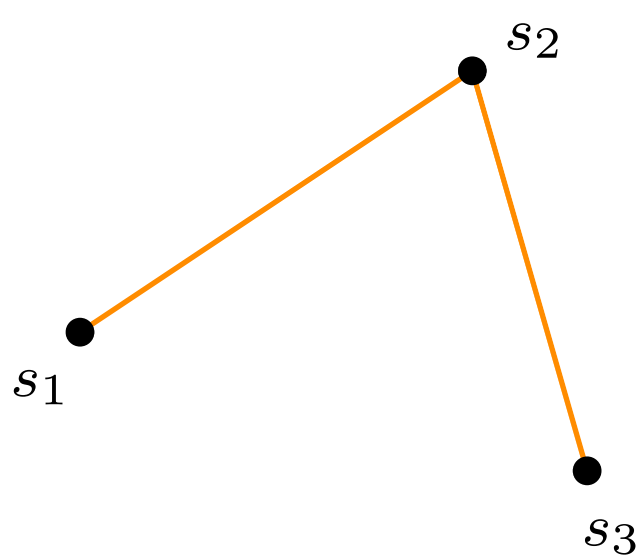

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

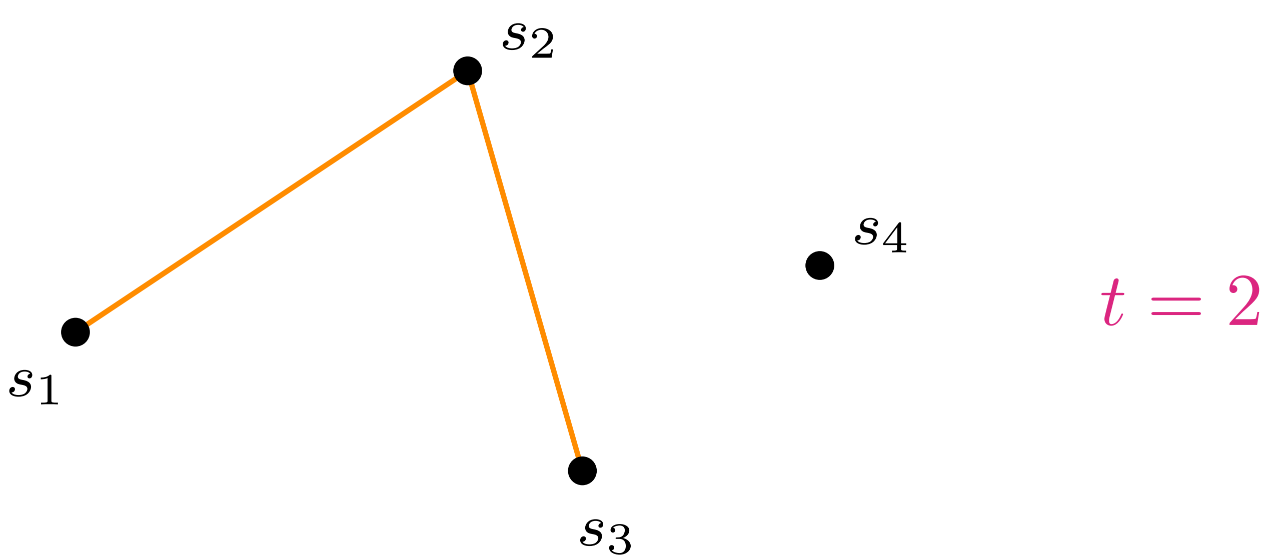

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \dots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \dots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \dots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

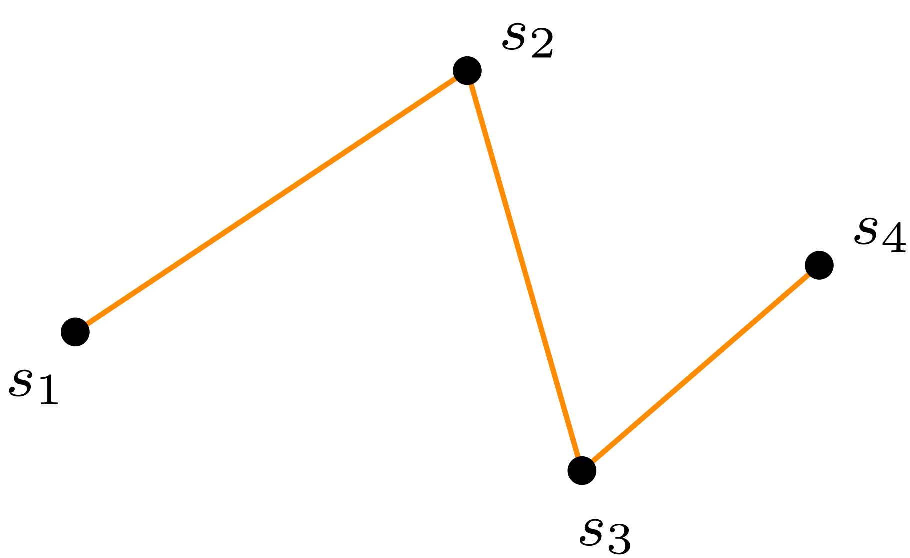

$t = 2$

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

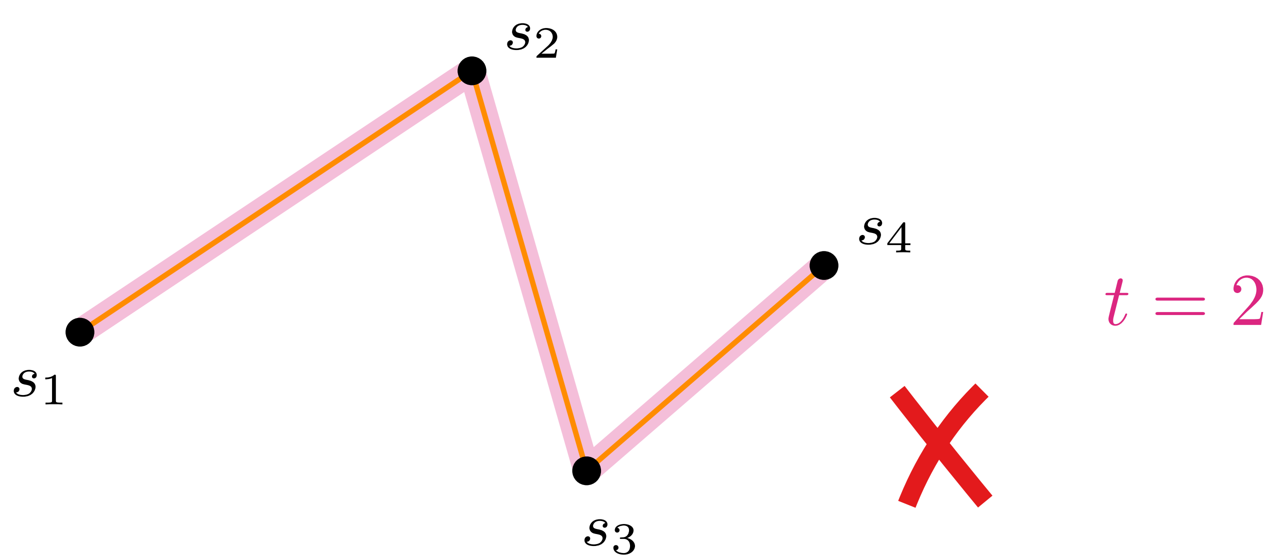

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \dots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \dots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \dots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

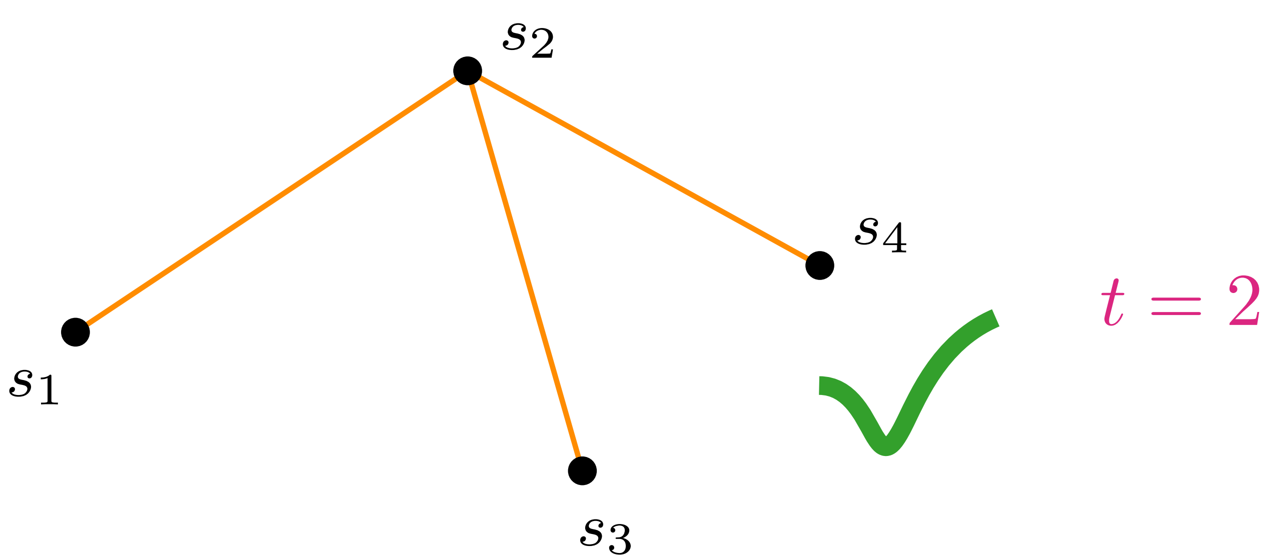

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \dots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \dots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \dots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

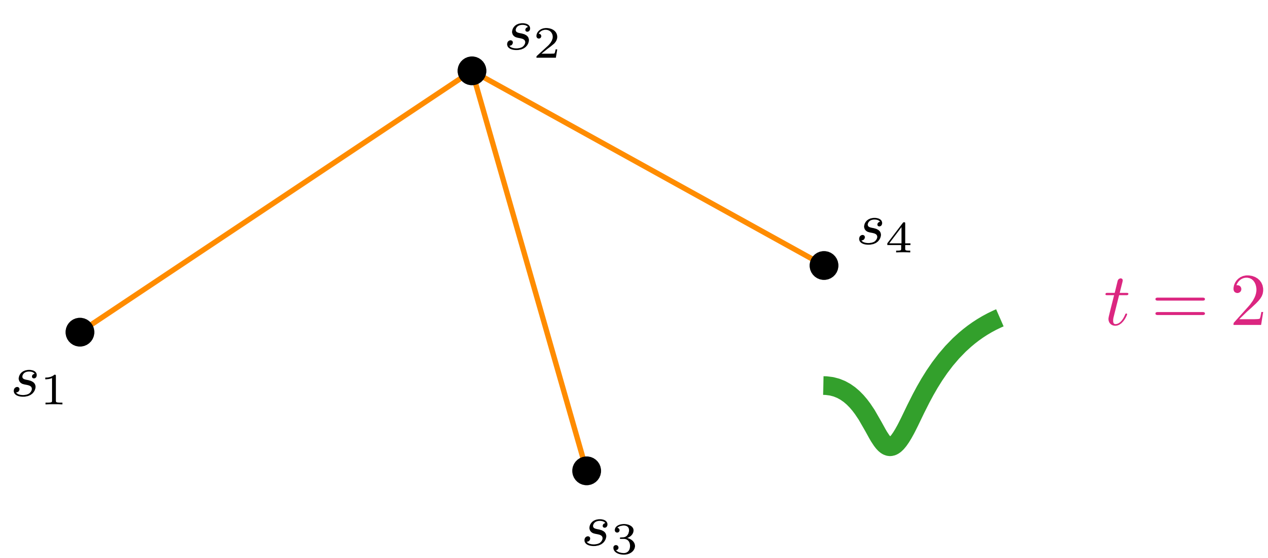

- **Performance** of an online algortihm **ALG** is measured by comparing it to the offline optimum **OPT** using the standard notion of competitive ratio.

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

- **Competitive Ratio** of an online t-spanner algorithm ALG is defined as $\sup_\sigma \frac{ALG(\sigma)}{OPT(\sigma)}$, where the supremum is taken over all input sequences σ, $OPT(\sigma)$ is the minimum weight of a t-spanner for the (unordered) set of points in σ, and $ALG(\sigma)$ denotes the weight of the t-spanner produced by ALG for this input sequence.

Going Online ...

Model -

- Input: We are given sequence of n points $(s_1, s_2, \ldots, s_n)$ in a metric space, where point s_i is presented in step i for $i = 1, \ldots, n$.
- Objective: Maintain a geometric t-spanner on $S_i = \{s_1, \ldots, s_i\}$ for each step i. The algorithm is allowed to add edges but not delete edges.

- **Competitive Ratio** of an online t-spanner algorithm ALG is defined as $\sup_\sigma \frac{ALG(\sigma)}{OPT(\sigma)}$, where the supremum is taken over all input sequences σ, $OPT(\sigma)$ is the minimum weight of a t-spanner for the (unordered) set of points in σ, and $ALG(\sigma)$ denotes the weight of the t-spanner produced by ALG for this input sequence.

Problem. Determine the best possible bounds for the competitive ratios for the weight and the number of edges of online t-spanners, for $t \geq 1$.

A bit of history ...

- Computing $(1+\varepsilon)$-spanner of minimum weight is NP-hard.
- Several approximation algorithms known - they also approximate the lightness.
- Online Steiner tree problem was studied by Imase and Waxman [SODA'1991] and they gave $\Theta(\log n)$.
- Alon and Azar [DCG'1993] studied minimum Steiner trees for points in the Euclidean plane, and gave improved bound of $\Omega(\log n / \log\log n)$.
- A large body of work done on dynamic spanners.

A bit of history ...

- Computing $(1+\varepsilon)$-spanner of minimum weight is NP-hard.
- Several approximation algorithms known - they also approximate the lightness.
- Online Steiner tree problem was studied by Imase and Waxman [SODA'1991] and they gave $\Theta(\log n)$.
- Alon and Azar [DCG'1993] studied minimum Steiner trees for points in the Euclidean plane, and gave improved bound of $\Omega(\log n / \log \log n)$.
- A large body of work done on dynamic spanners.

Online Steiner spanners.

- It is allowed to use auxiliary points (Steiner points) which are not part of input sequence of points.
- An online algorithm is allowed to add Steiner points and subdivide existing edges with Steiner points at each time step.

Effects of Irrevocability

- The algorithm is allowed to add edges but not delete edges.

Effects of Irrevocability

- The algorithm is allowed to add edges but not delete edges.
- The value of OPT is not necessarily monotone!!!

Effects of Irrevocability

- The algorithm is allowed to add edges but not delete edges.
- The value of OPT is not necessarily monotone!!!

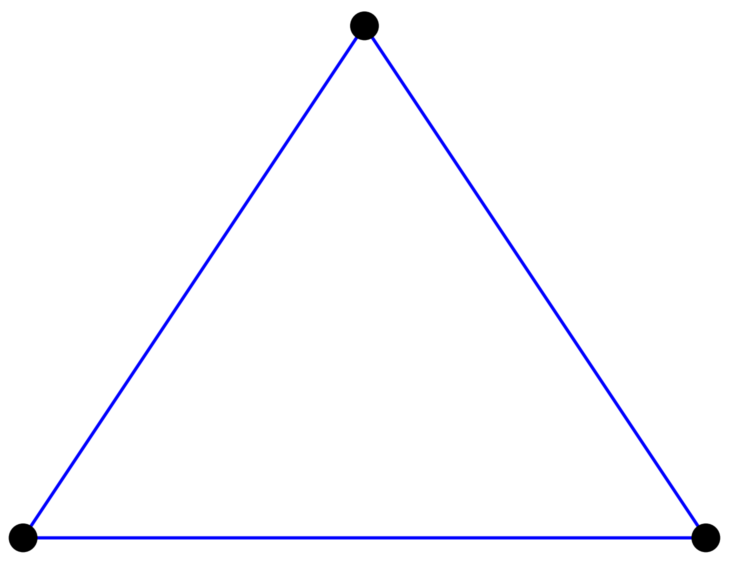

An optimum $\frac{3}{2}$-spanner on three points with all edges of unit length

Effects of Irrevocability

- The algorithm is allowed to add edges but not delete edges.
- The value of OPT is not necessarily monotone!!!

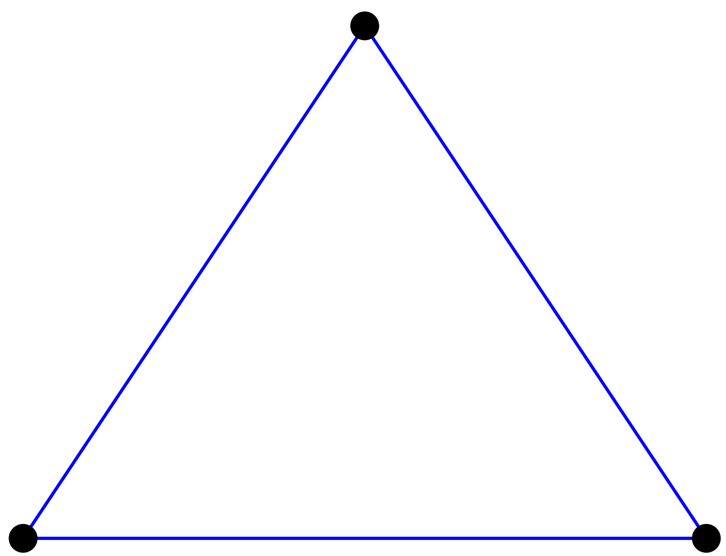

An optimum $\frac{3}{2}$-spanner on three points with all edges of unit length

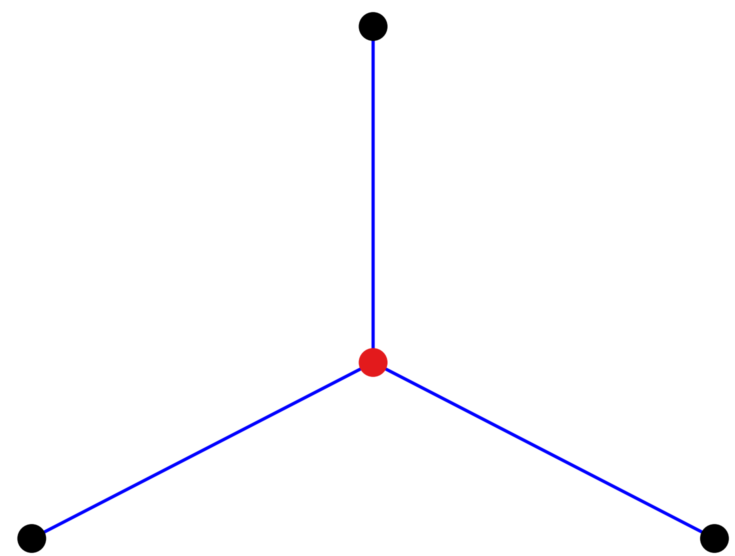

After inserting a point at the center, the cost decreases.

Let's start with one dimension.

- Lower Bound

Let's start with one dimension.

- Lower Bound

  - Adversarial strategy - start with two points $p_0 = 0$ and $q_0 = 1$. Then, successively places points $p_i = i \cdot \frac{\varepsilon}{2}$, for $i = 1, \ldots, n$ so that all points remain in the interval $[0, \frac{1}{2}]$.

Let's start with one dimension.

- Lower Bound

  - Adversarial strategy - start with two points $p_0 = 0$ and $q_0 = 1$. Then, successively places points $p_i = i \cdot \frac{\varepsilon}{2}$, for $i = 1, \ldots, n$ so that all points remain in the interval $[0, \frac{1}{2}]$.

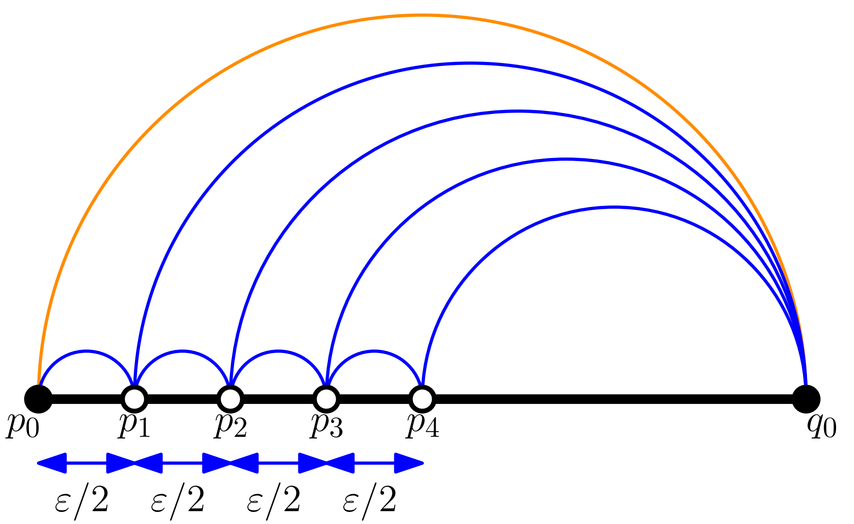

Let's start with one dimension.

- Lower Bound

  - Adversarial strategy - start with two points $p_0 = 0$ and $q_0 = 1$. Then, successively places points $p_i = i \cdot \frac{\varepsilon}{2}$, for $i = 1, \ldots, n$ so that all points remain in the interval $[0, \frac{1}{2}]$.

  - Repeats the same strategy in every subinterval.

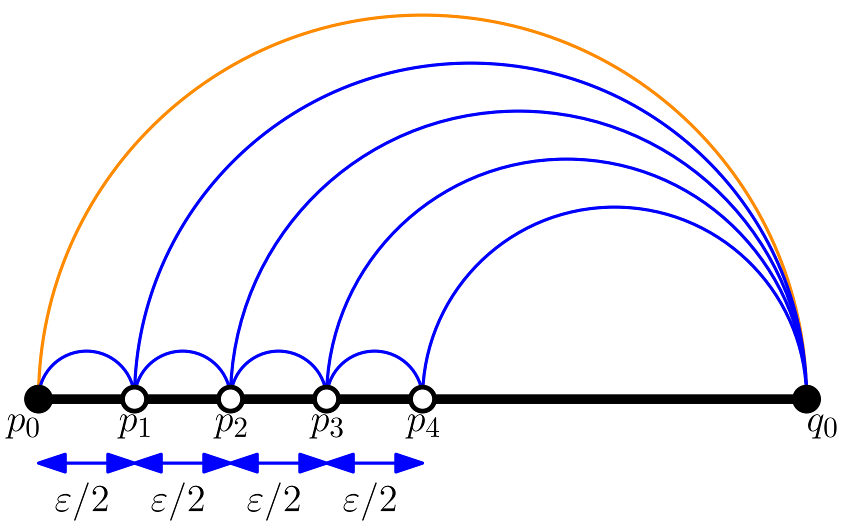

Let's start with one dimension.

Lower Bound

- Adversarial strategy - start with two points $p_0 = 0$ and $q_0 = 1$. Then, successively places points $p_i = i \cdot \frac{\varepsilon}{2}$, for $i = 1, \ldots, n$ so that all points remain in the interval $[0, \frac{1}{2}]$.

- Repeats the same strategy in every subinterval.

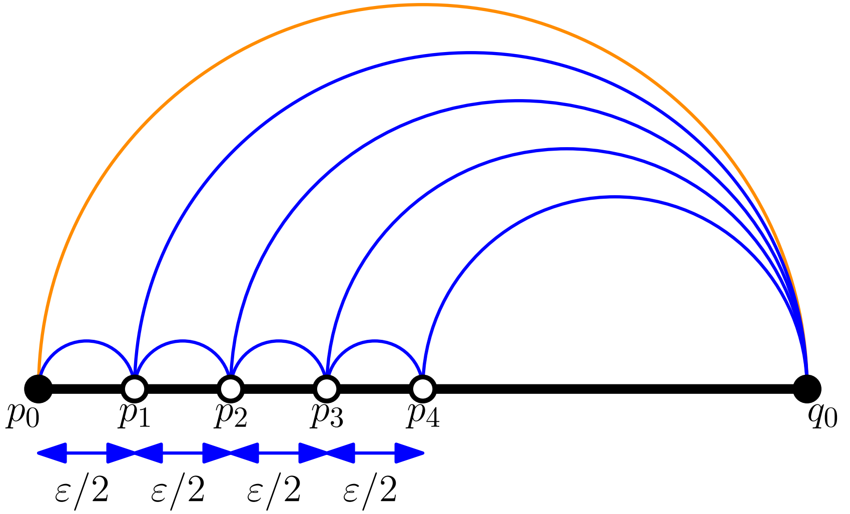

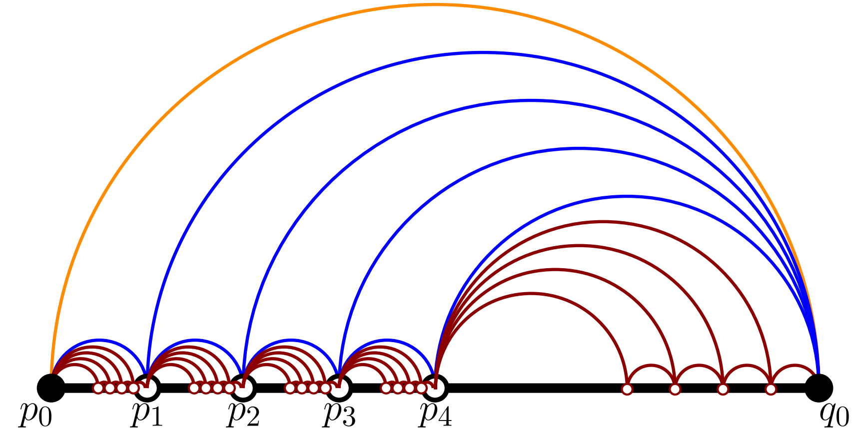

One dimension - Cont.

Upper Bound

Algorithm -

- For all $i = 1, \dots, n$, we maintain a spanning graph G_i on $S_i = \{s_1, \dots, s_i\}$ and the x-monotone path P_i between the leftmost and the rightmost points in $S_i = \{s_1, \dots, s_i\}$.
- When s_i arrives,
  - If s_i is left (resp., right) of all previous points - add to the closest point.
  - Else, consider the interval inside which s_i appeared, and join to the endpoints.

One dimension - Cont.

Upper Bound

Algorithm -

- For all $i = 1, \ldots, n$, we maintain a spanning graph G_i on $S_i = \{s_1, \ldots, s_i\}$ and the x-monotone path P_i between the leftmost and the rightmost points in $S_i = \{s_1, \ldots, s_i\}$.
- When s_i arrives,
  - If s_i is left (resp., right) of all previous points - add to the closest point.
  - Else, consider the interval inside which s_i appeared, and join to the endpoints.

Theorem. Competitive ratio of any online algorithm for $(1+\varepsilon)$-spanners for a sequence of points on a line is $\Omega(\varepsilon^{-1} \log n / \log \varepsilon^{-1})$. Moreover, there is an online algorithm that maintains a $(1+\varepsilon)$-spanner with competitive ratio $O(\varepsilon^{-1} \log n / \log \varepsilon^{-1})$.

Higher Dimensions under the L_2-norm

Theorem. For every $\varepsilon > 0$, an online algorithm can maintain, for a sequence of $n \in \mathbb{N}$ points in $\mathbb{R}^d$, a Euclidean Steiner $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{(1-d)/2} \log n) \cdot \mathrm{OPT})$.

Higher Dimensions under the L_2-norm

Theorem. For every $\varepsilon > 0$, an online algorithm can maintain, for a sequence of $n \in \mathbb{N}$ points in $\mathbb{R}^d$, a Euclidean Steiner $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{(1-d)/2} \log n) \cdot \mathrm{OPT})$.

Online Algorithm in two layers:

1. DefSpanner algorithm [Gao et al., 2006]: a $(1+\varepsilon)$-spanner G_1 of weight $O(\varepsilon^{-(d+1)} \log n \cdot \mathrm{OPT})$, without Steiner points.
2. We maintain a $(1+\varepsilon)$-spanner G_2 for G_1 of weight $O(\varepsilon^{(1-d)/2} \log n \cdot \mathrm{OPT})$ with Steiner points.

Higher Dimensions under the L_2-norm

Theorem. For every $\varepsilon > 0$, an online algorithm can maintain, for a sequence of $n \in \mathbb{N}$ points in $\mathbb{R}^d$, a Euclidean Steiner $(1+\varepsilon)$-spanner of weight $O(\varepsilon^{(1-d)/2} \log n) \cdot \mathrm{OPT})$.

Online Algorithm in two layers:

1. DefSpanner algorithm [Gao et al., 2006]: a $(1+\varepsilon)$-spanner G_1 of weight $O(\varepsilon^{-(d+1)} \log n \cdot \mathrm{OPT})$, without Steiner points.
2. We maintain a $(1+\varepsilon)$-spanner G_2 for G_1 of weight $O(\varepsilon^{(1-d)/2} \log n \cdot \mathrm{OPT})$ with Steiner points.

→ Stretch factor: $(1+\varepsilon)(1+\varepsilon) = (1+O(\varepsilon))$

→ Key idea: nearly parallel edges, of similar lengths in G_1 are replaced by a **Shalow-Light Tree (SLT)** in G_2.

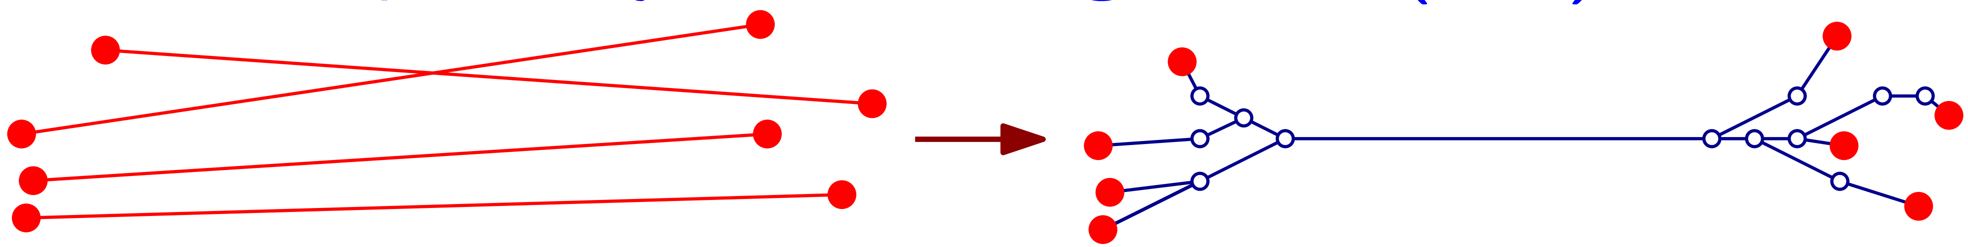

Higher Dimensions under the L_2-norm

Overview of DefSpanner algorithm by Gao et al [2006].

1. **Hierarchical Clustering:** Maintain a **quadtree** for the points incrementally (online)
2. **Well-Separated Pair Decomposition (WSPD)**: At each level of the quadtree, with cubes of side length ℓ, add an edge between any two nonempty cells at distance $O(\ell/\varepsilon)$.

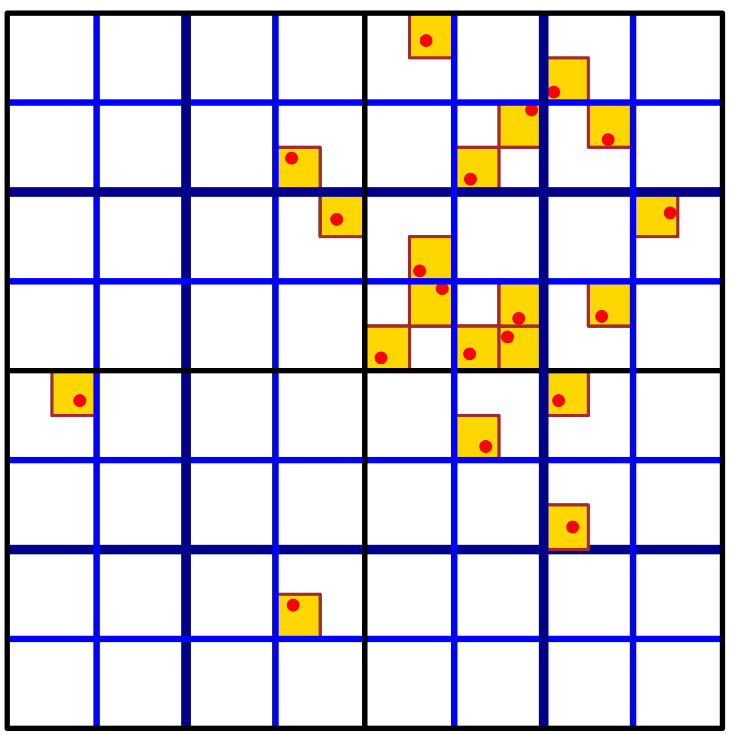

Q_0

Higher Dimensions under the L_2-norm

Overview of DEFSPANNER algorithm by Gao et al [2006].

1. **Hierarchical Clustering:** Maintain a **quadtree** for the points incrementally (online)
2. **Well-Separated Pair Decomposition (WSPD)**: At each level of the quadtree, with cubes of side length ℓ, add an edge between any two nonempty cells at distance $O(\ell/\varepsilon)$.

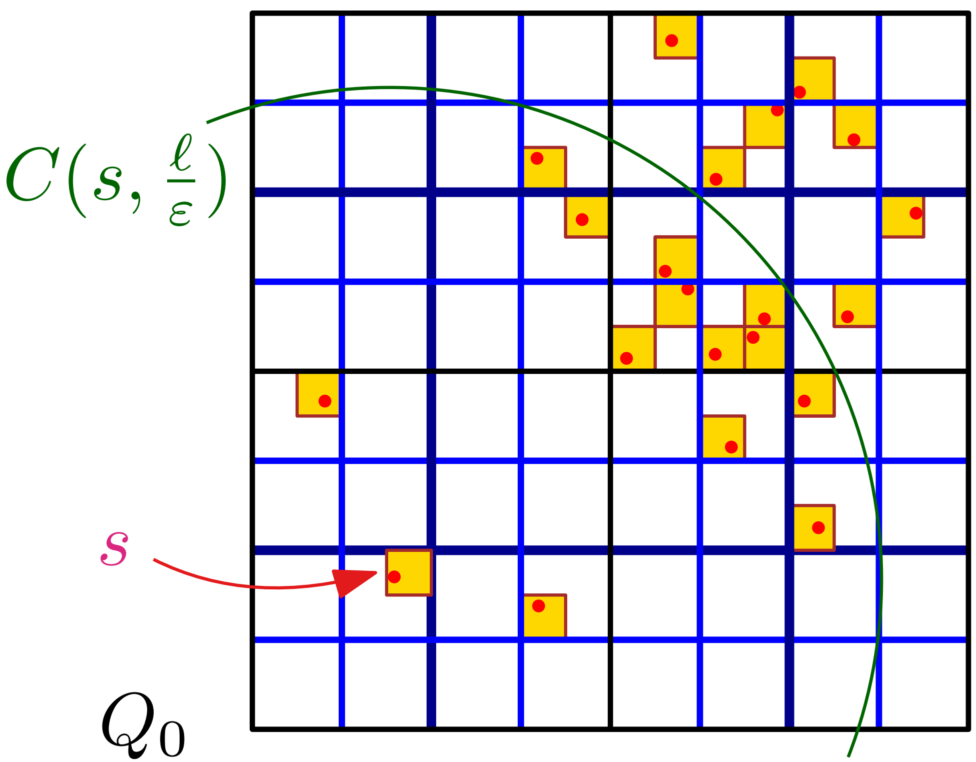

When a new point s arrives:

- Insert s into the Quadtree,
- If s is the first point in a cell, which has side length ℓ, add edges between s and a representative of other cells within distance $O(\ell/\varepsilon)$.

Higher Dimensions under the L_2-norm

Overview of DefSpanner algorithm by Gao et al [2006].

1. **Hierarchical Clustering:** Maintain a **quadtree** for the points incrementally (online)
2. **Well-Separated Pair Decomposition (WSPD)**: At each level of the quadtree, with cubes of side length ℓ, add an edge between any two nonempty cells at distance $O(\ell/\varepsilon)$.

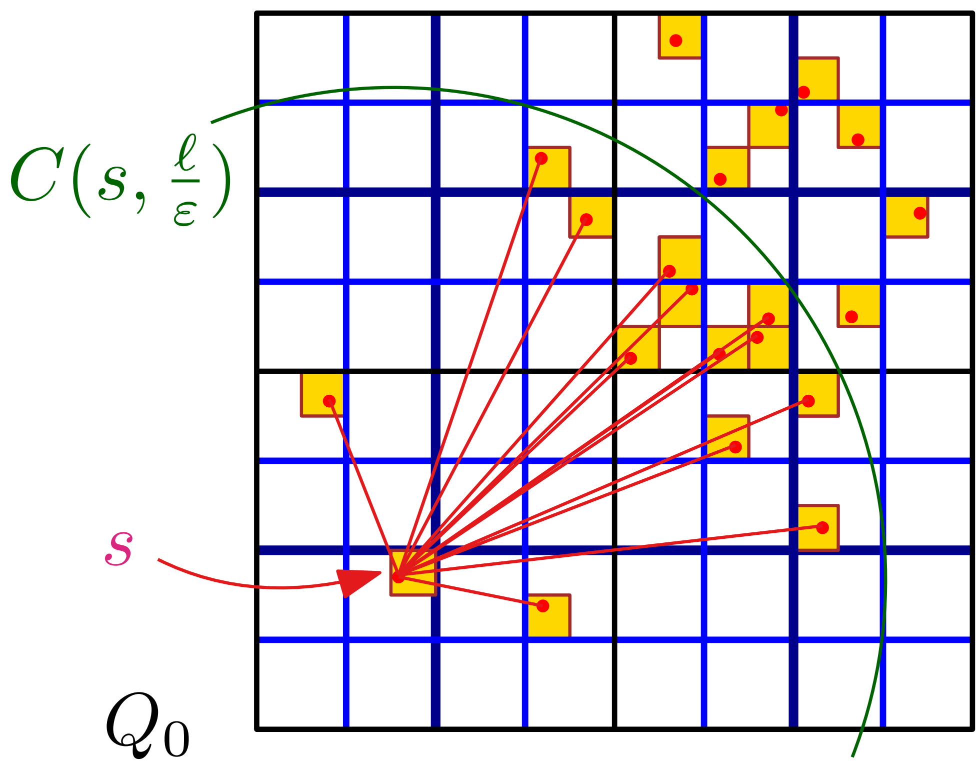

When a new point s arrives:

- Insert s into the Quadtree,
- If s is the first point in a cell, which has side length ℓ, add edges between s and a representative of other cells within distance $O(\ell/\varepsilon)$.

Higher Dimensions under the L_2-norm

Online Steiner spanner algorithm.

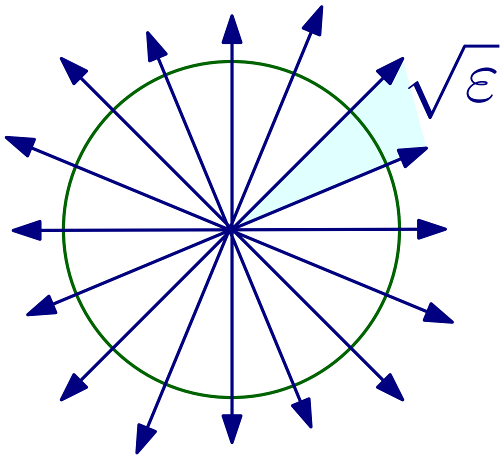

- Partition the sphere of directions into $\Theta(\varepsilon^{(1-d)/2})$ cones of aperture $\sqrt{\varepsilon}$).
- We constuction a Steiner spanner for each cone of directions.

Higher Dimensions under the L_2-norm

Online Steiner spanner algorithm.

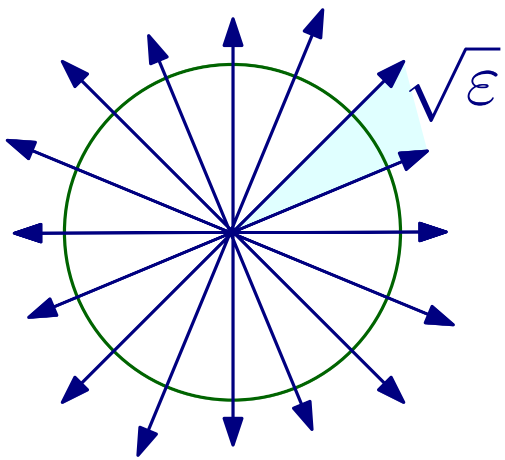

- Partition the sphere of directions into $\Theta(\varepsilon^{(1-d)/2})$ cones of aperture $\sqrt{\varepsilon}$).
- We constuction a Steiner spanner for each cone of directions.

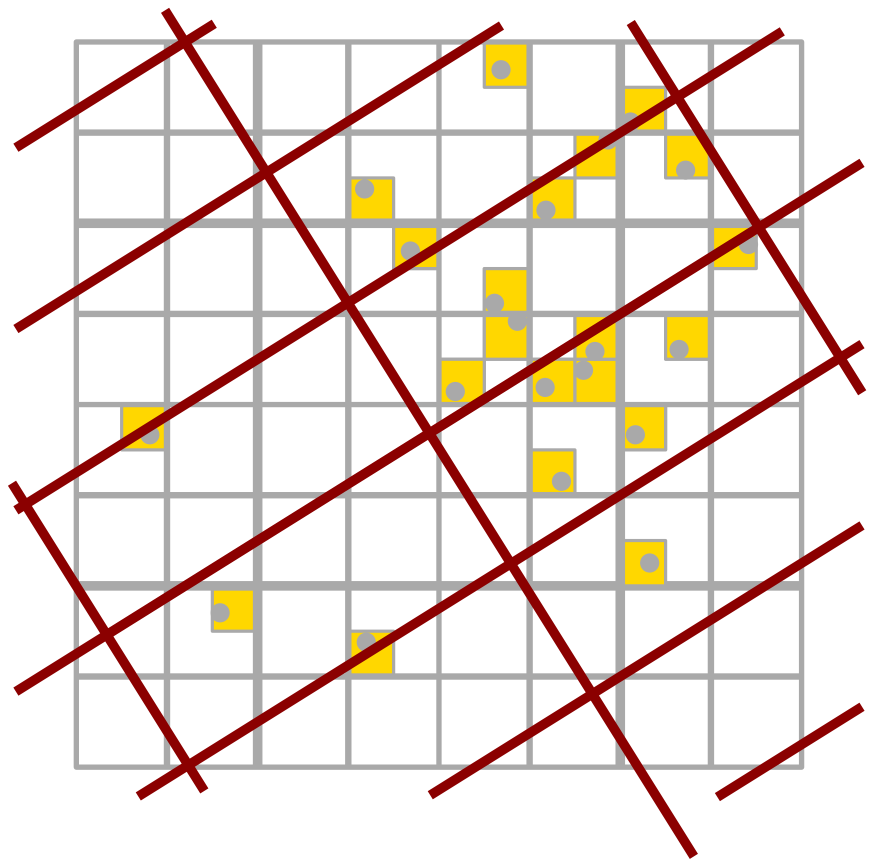

- For each level of the quadtree, create a covering cylinders of width $\sqrt{\varepsilon} \cdot \ell$
- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges of the same direction.

Higher Dimensions under the L_2-norm

Online Steiner spanner algorithm.

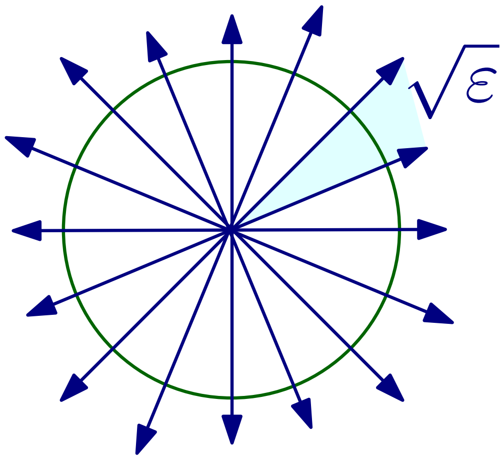

- Partition the sphere of directions into $\Theta(\varepsilon^{(1-d)/2})$ cones of aperture $\sqrt{\varepsilon}$).
- We constuction a Steiner spanner for each cone of directions.

- For each level of the quadtree, create a covering cylinders of width $\sqrt{\varepsilon} \cdot \ell$
- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges of the same direction.

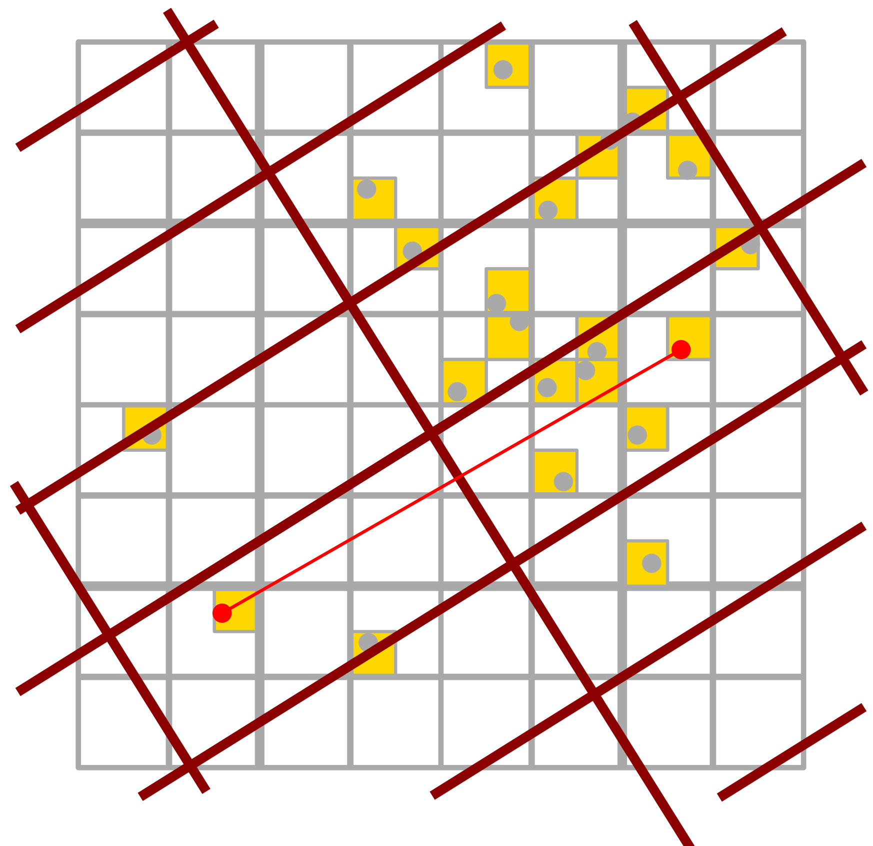

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

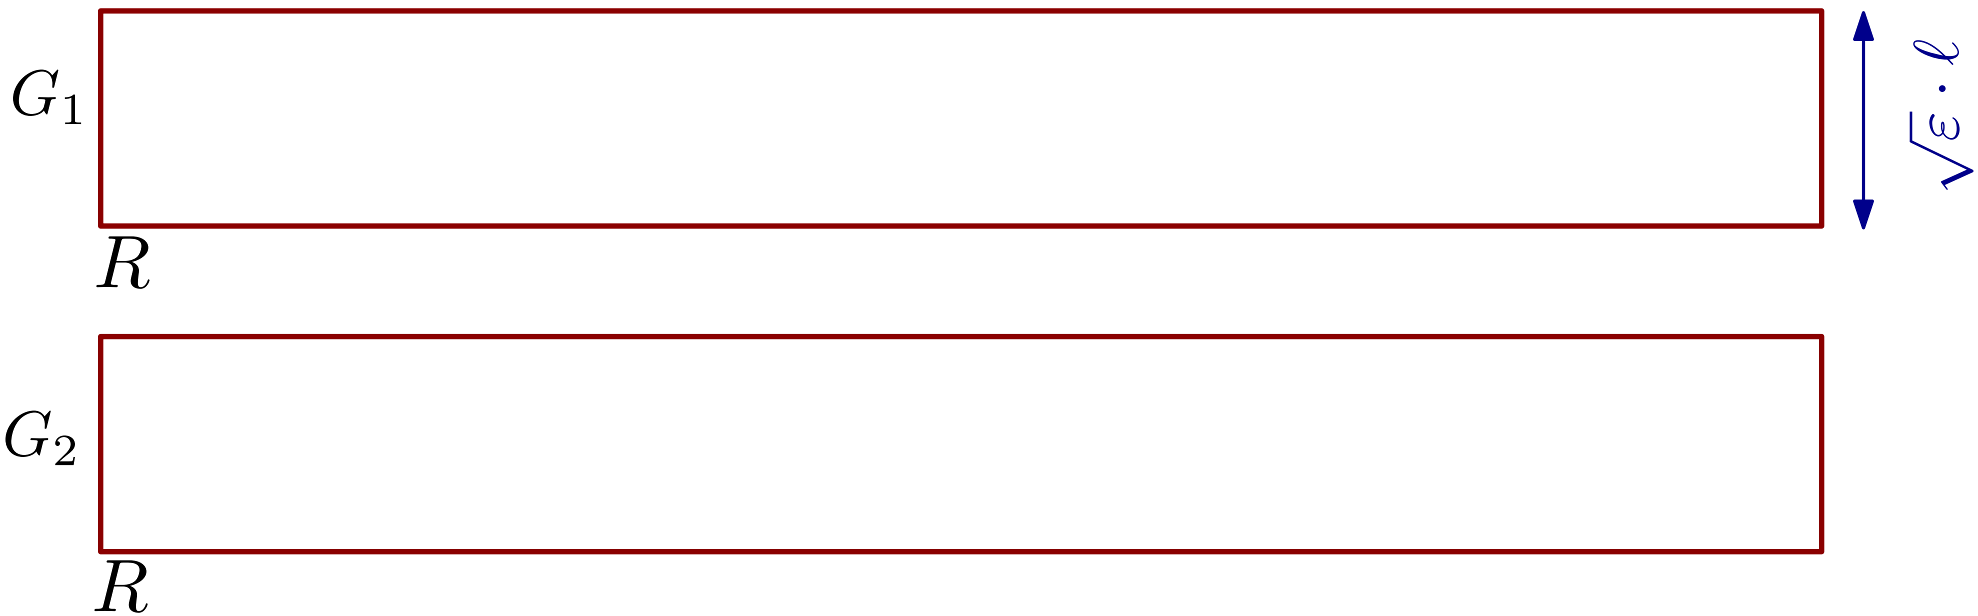

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

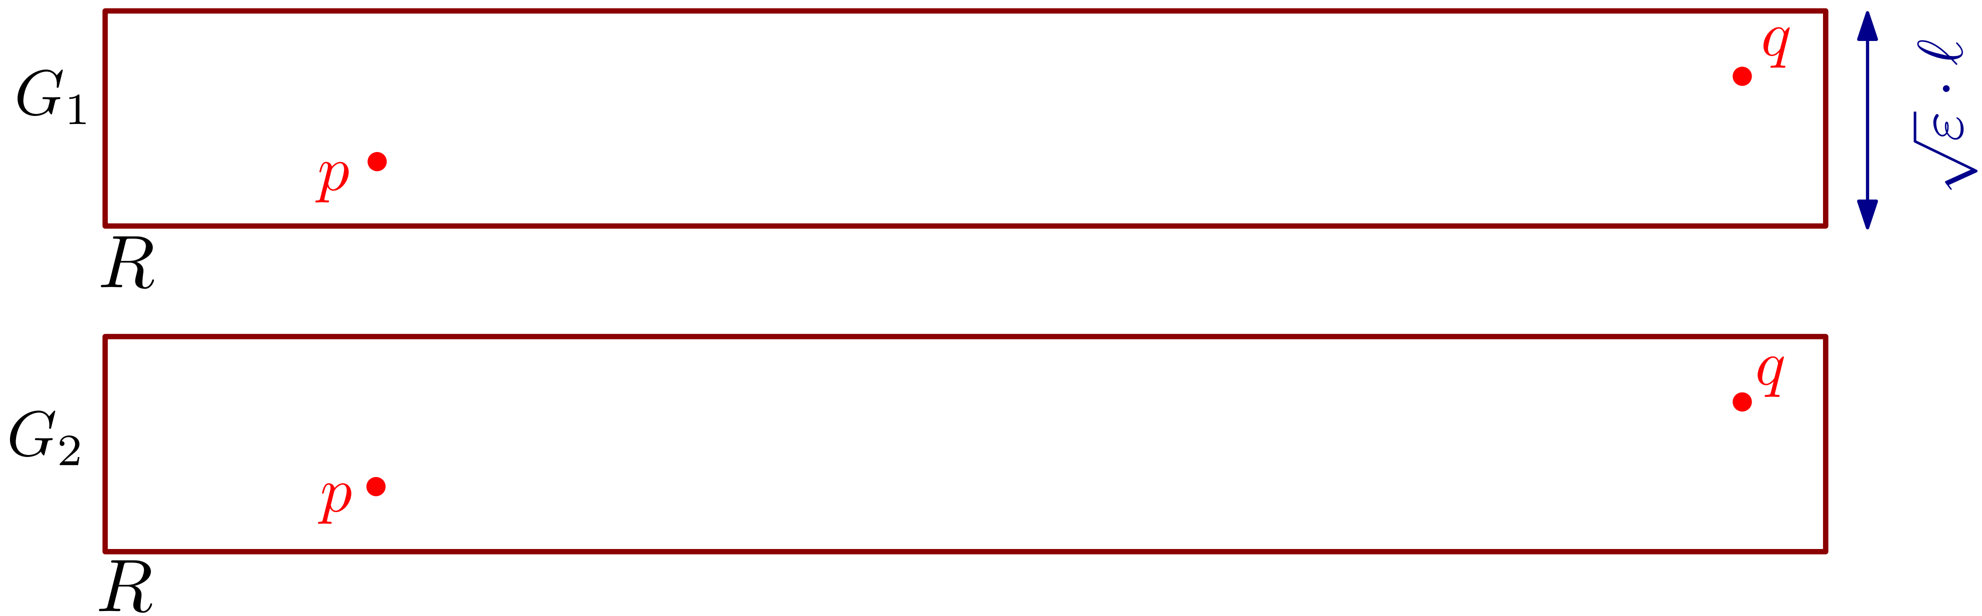

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

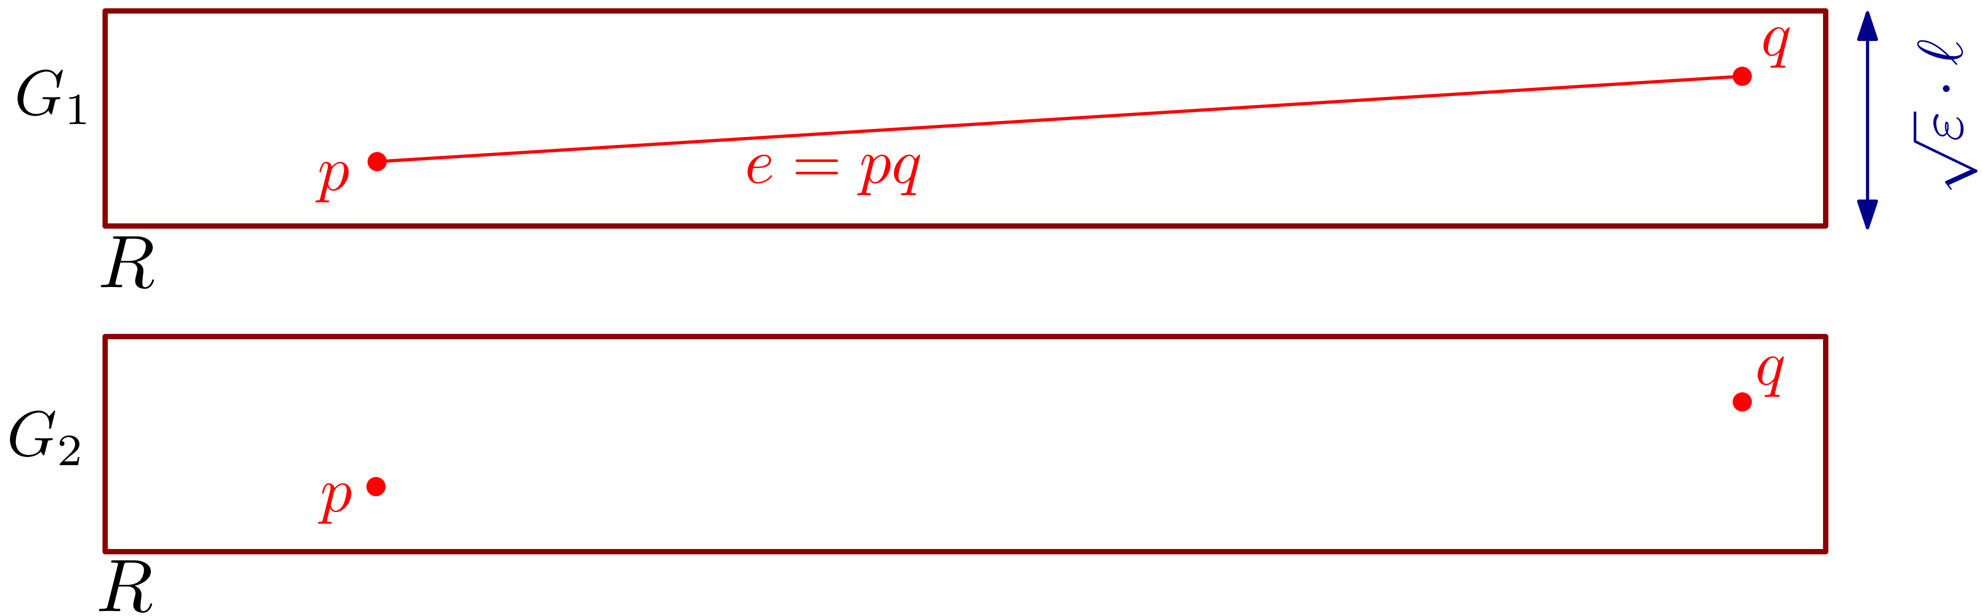

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points;
- add a SLT between the two grids.

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

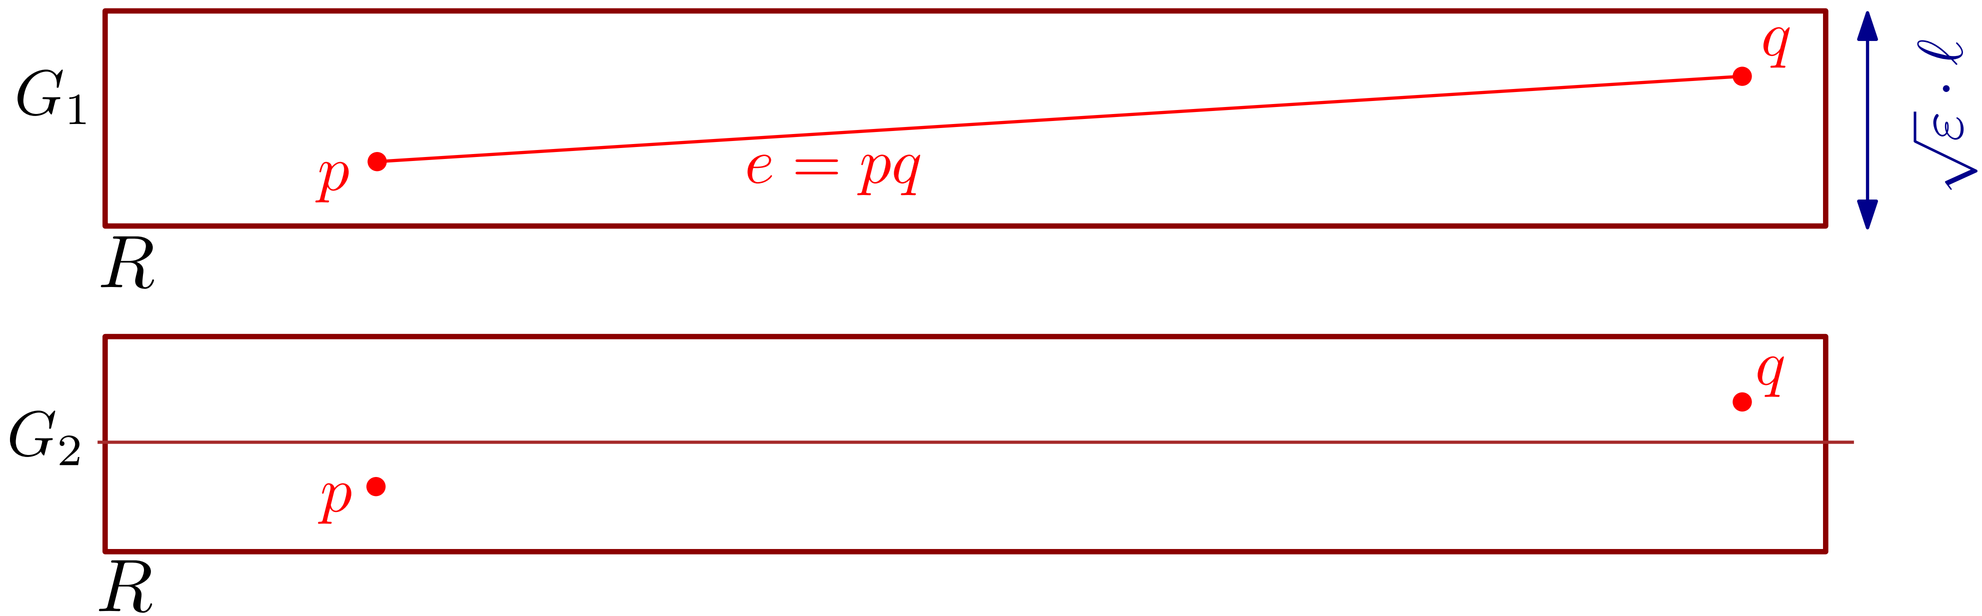

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points;
- add a SLT between the two grids.

Higher Dimensions under the L_2-norm

- When DEFSPANNER inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

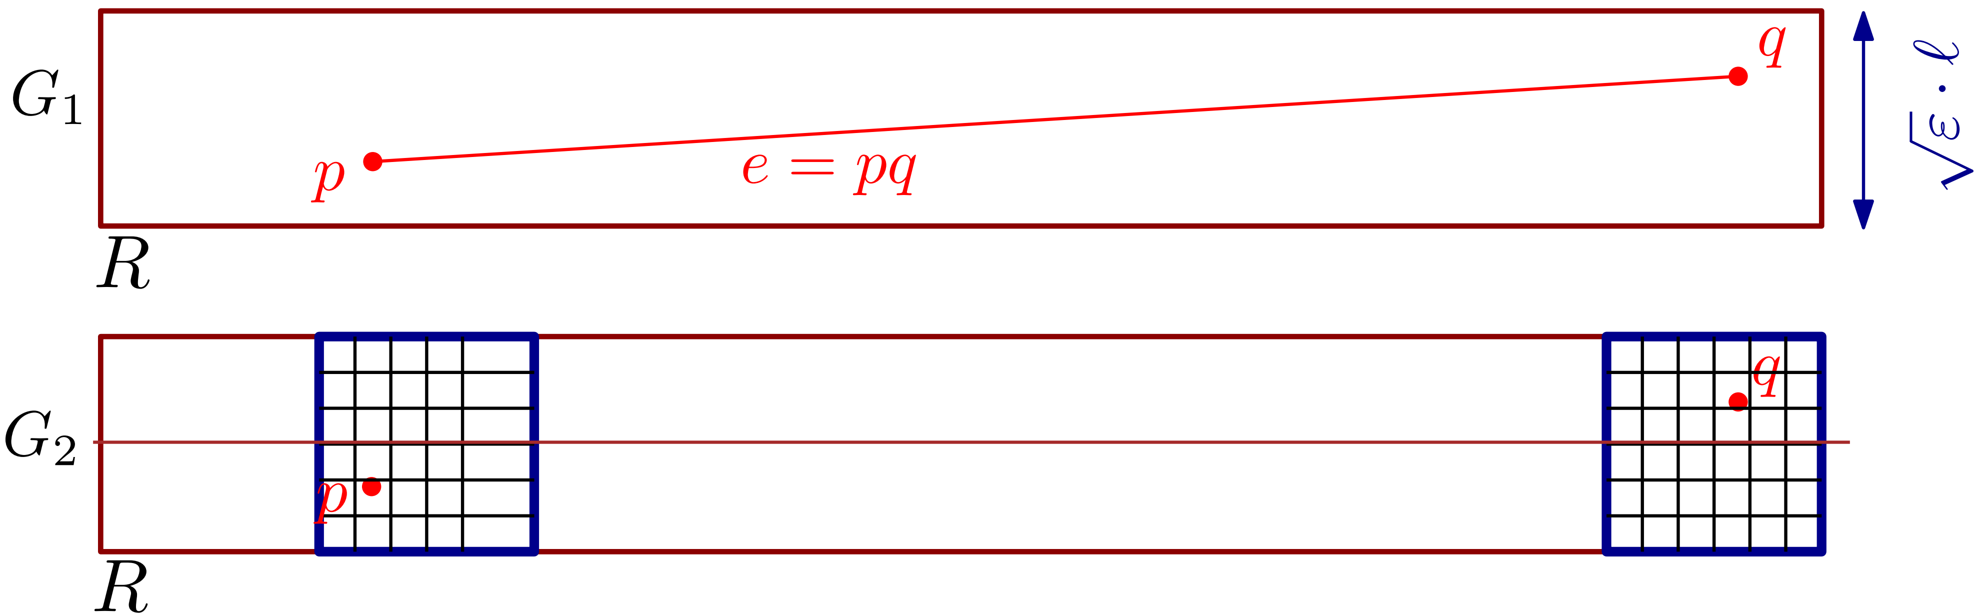

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points;
- add a SLT between the two grids.

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

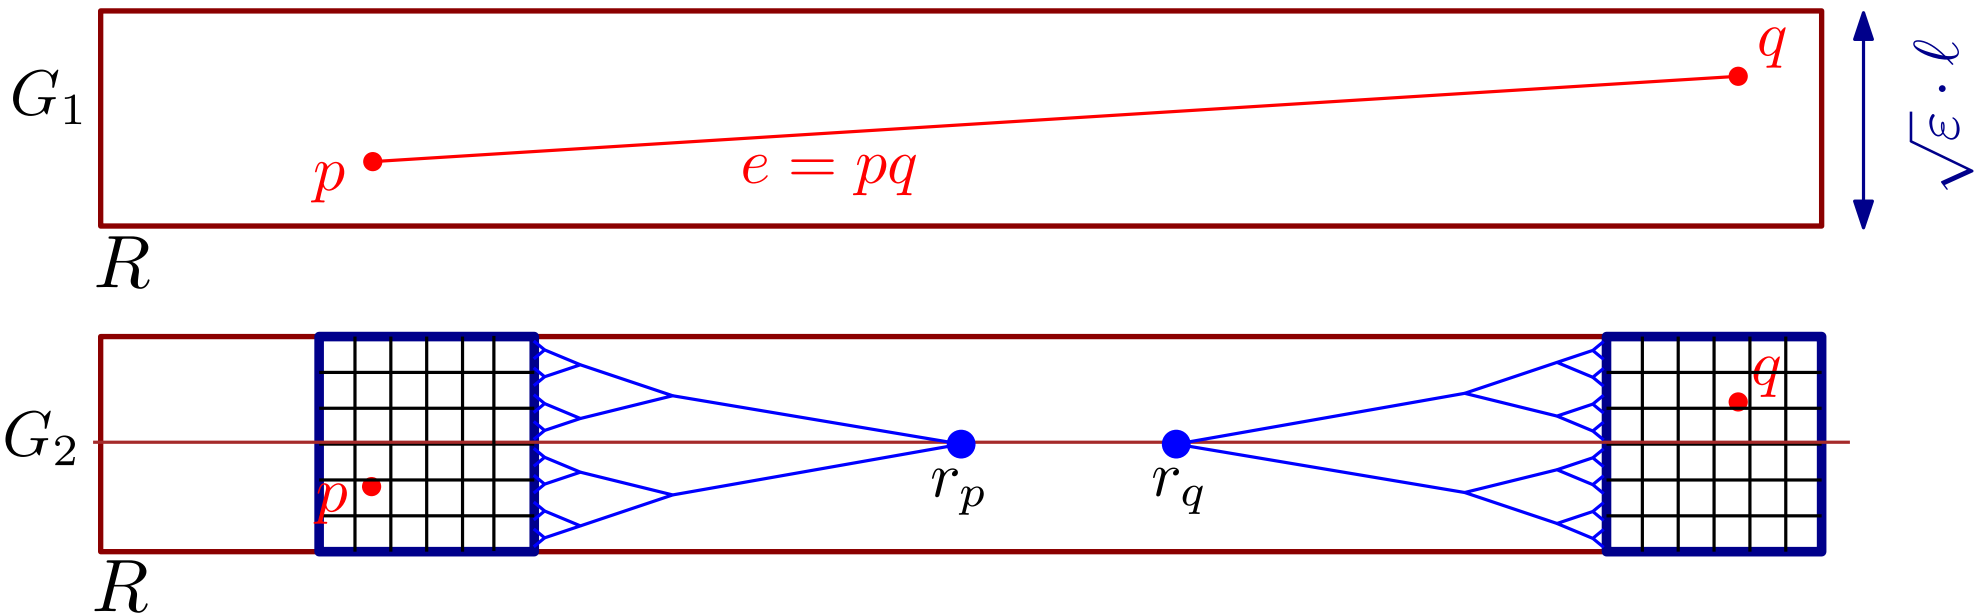

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points;
- add a SLT between the two grids.

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

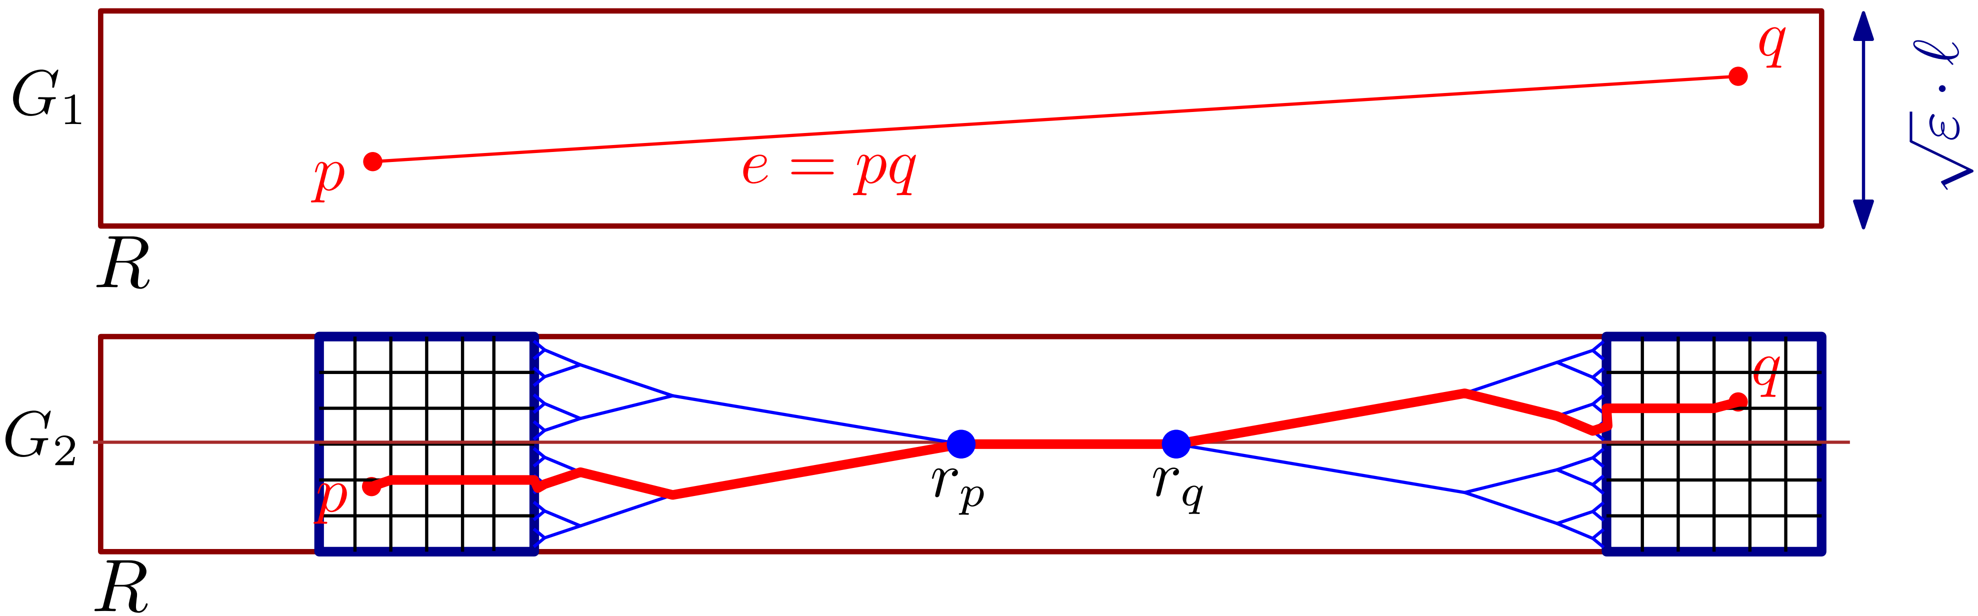

For the first edge $e = pq$,

- add a “backbone” line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points;
- add a SLT between the two grids.

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

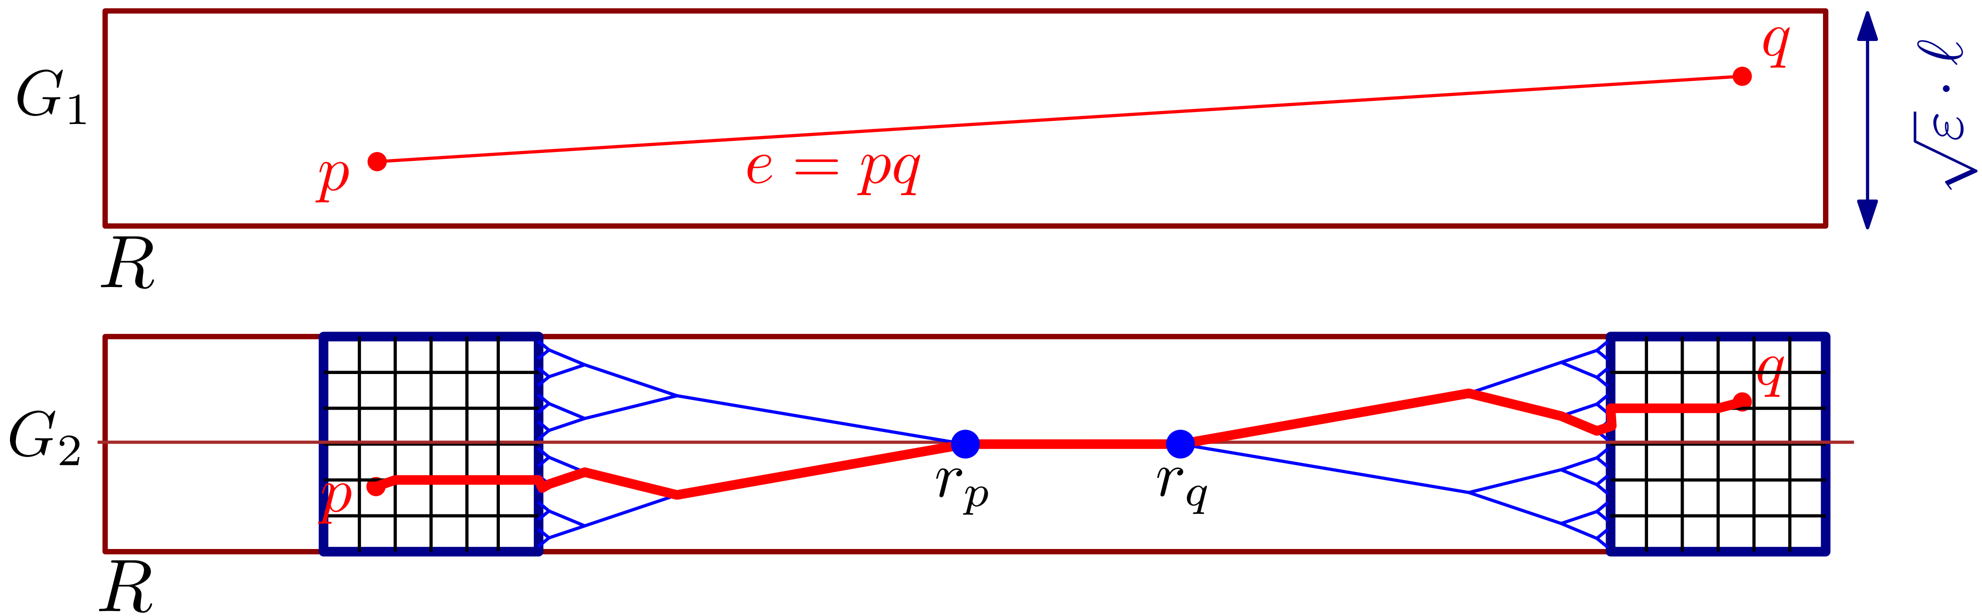

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points; ← cost for additional edges
- add a SLT between the two grids.

Future edges in the same cylinder can use the same infrastructure.
The extra cost is $O(\varepsilon\ell)$ for the connection to the closest grid points.

Higher Dimensions under the L_2-norm

- When DefSpanner inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

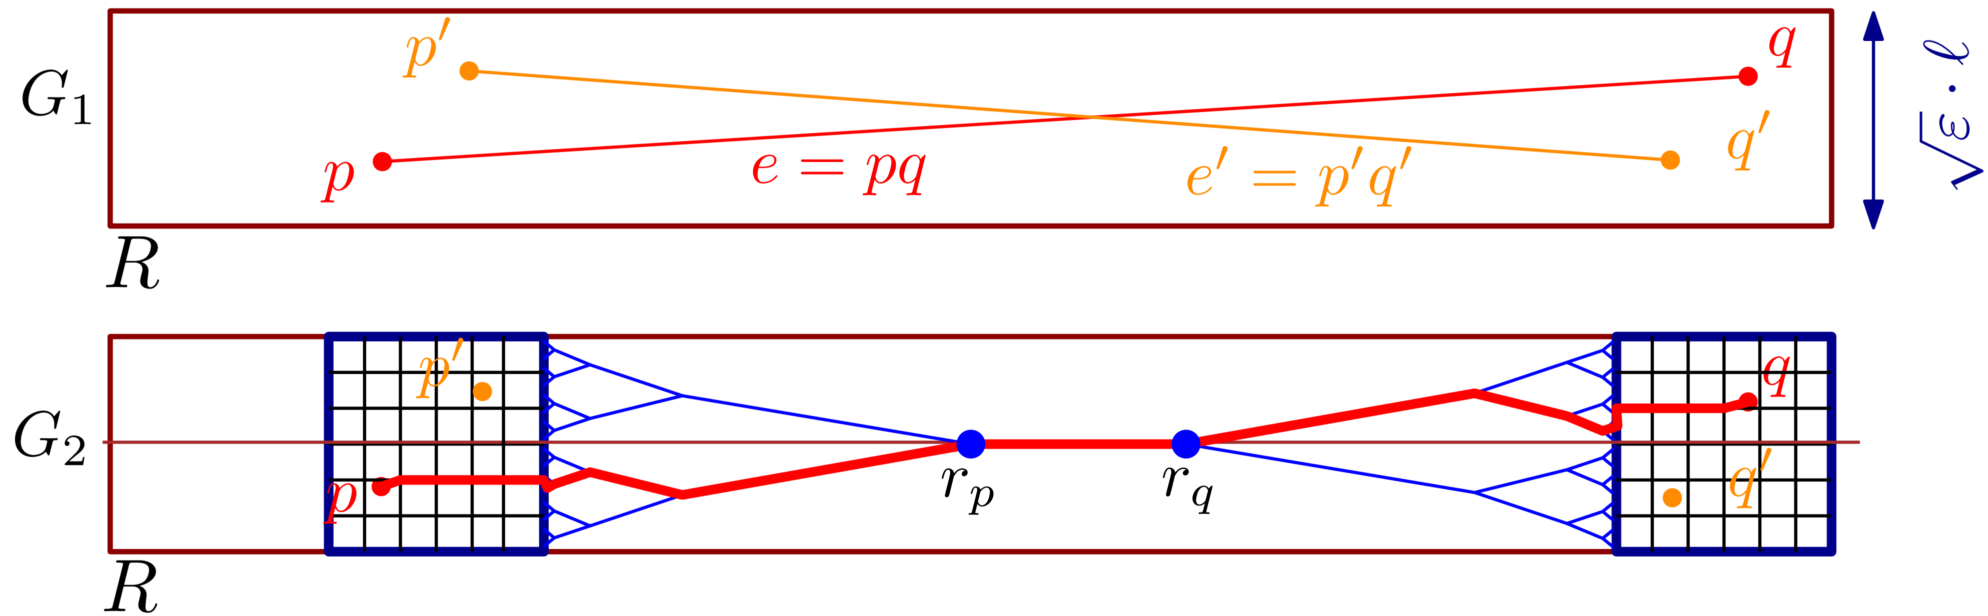

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points; ← cost for additional edges
- add a SLT between the two grids.

Future edges in the same cylinder can use the same infrastructure.
The extra cost is $O(\varepsilon\ell)$ for the connection to the closest grid points.

Higher Dimensions under the L_2-norm

- When DEFSPANNER inserts an edge e in a cyclinder, we construct an SLT that can accomodate future edges in the same cyclinder.

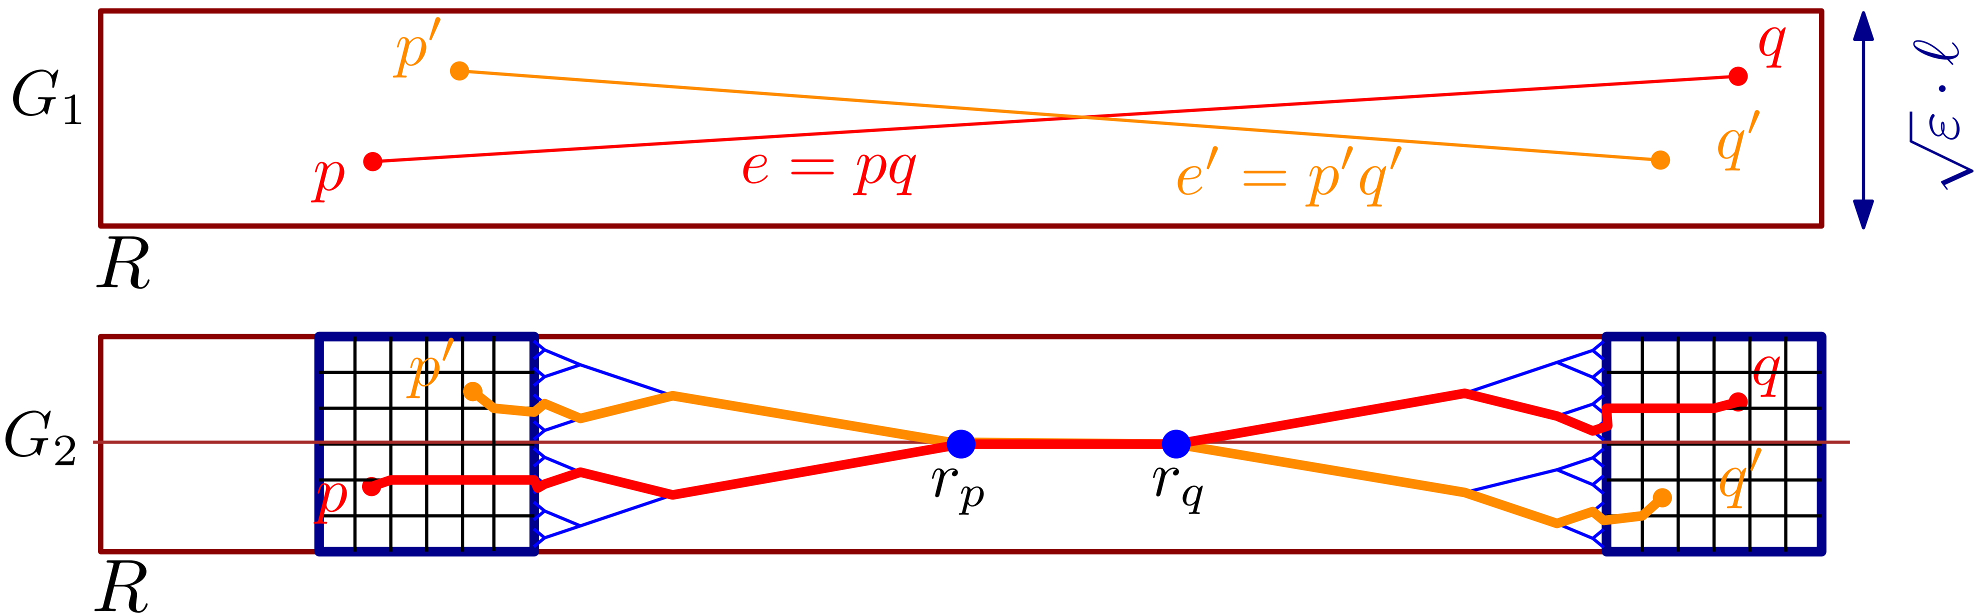

For the first edge $e = pq$,

- add a "backbone" line in the center of the cyclinder;
- add a grid of cell-size $\varepsilon\ell$ around p and q;
- connect p and q to the nearest grid points; ← cost for additional edges
- add a SLT between the two grids.

Future edges in the same cylinder can use the same infrastructure.
The extra cost is $O(\varepsilon\ell)$ for the connection to the closest grid points.

Higher Dimensions under the L_2-norm

Competitive analysis. In each cyclinder, we charge the weight of the Steiner graph (backbone, grid, and SLT) to the weight of OPT.

For every $p, q \in S$, an OPT spanner contains a pq-path of weight at most $(1+\varepsilon)\|pq\|$ in an ellipse B_{pq} with foci p & q.

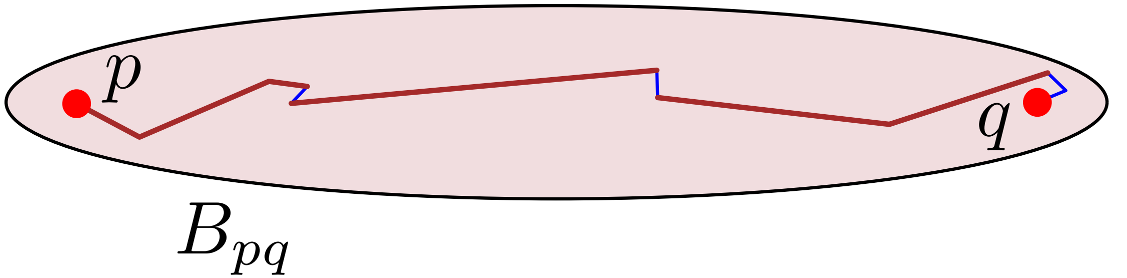

Lemma (B&T, STACS 2021). In a pq-path of weight $\leq (1+\varepsilon)\|pq\|$, the edges e such that $\angle(pq, e) \leq \sqrt{\varepsilon}$ have total weight at least $\frac{1}{2}\|pq\|$.

Higher Dimensions under the L_2-norm

Competitive analysis. In each cyclinder, we charge the weight of the Steiner graph (backbone, grid, and SLT) to the weight of OPT.

For every $p, q \in S$, an OPT spanner contains a pq-path of weight at most $(1+\varepsilon)\|pq\|$ in an ellipse B_{pq} with foci p & q.

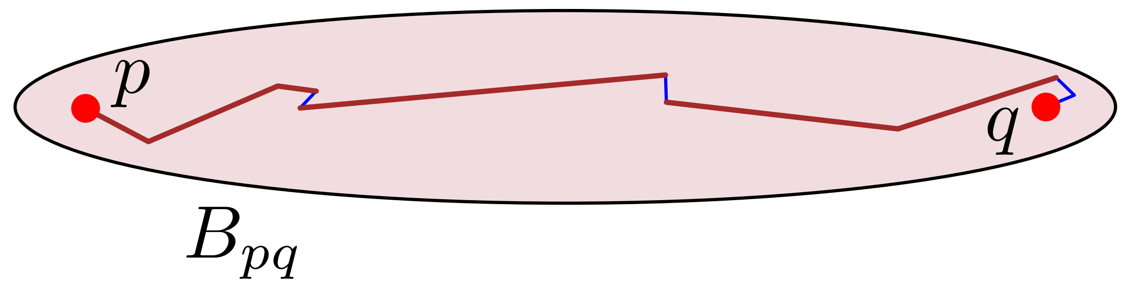

Lemma (B&T, STACS 2021). In a pq-path of weight $\leq (1+\varepsilon)\|pq\|$, the edges e such that $\angle(pq, e) \leq \sqrt{\varepsilon}$ have total weight at least $\frac{1}{2}\|pq\|$.

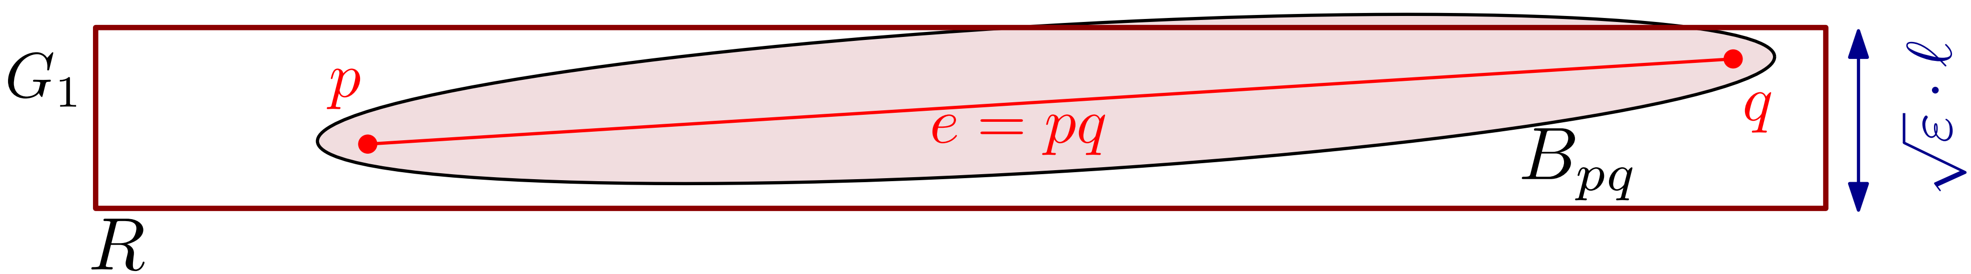

- The ellipse B_{pq} lies in a small neighborhood of the cyclinder.
- OPT contains edges of weight $\geq \frac{1}{2}\|pq\|$ of direction $\angle(e', pq) \leq \sqrt{\varepsilon}$) in a small neighborhood of the cylinder.
- The ratio $\frac{\text{ALG}}{\text{OPT}} \leq O\left(\varepsilon^{\frac{1-d}{2}}\right)$ holds for each cylinder & each direction.

Where do we stand ...

Without Steiner points.

Family	Stretch	# of edges	Lightness
General metrics	$(2k-1)(1+\varepsilon)$	$O(\varepsilon^{-1}\log(\frac{1}{\varepsilon}))n^{1+\frac{1}{k}}$	$O(n^{\frac{1}{k}}\varepsilon^{-1}\log^2 n)$
Euclidean d-space	$1+\varepsilon$	$\tilde{O}_d(\varepsilon^{1-d})n$	$O(\varepsilon^{-d}\log n)$
Real line	$1+\varepsilon$	$n-1$	$\tilde{\Theta}(\varepsilon^{-1}\log n)$
Doubling[GGN'06	$1+\varepsilon$	$\varepsilon^{-O(d)}n$	$\varepsilon^{-O(d)}\log n$
Family	**Stretch**	**# of edges**	**Comp. ratio**
General metrics	$(2k-1)$	$-$	$\Omega(\frac{1}{k}\cdot n^{\frac{1}{k}})$
Euclidean plane	$1+\varepsilon$	$\tilde{O}(\varepsilon^{-1})n$	$\tilde{O}(\varepsilon^{-3/2}\log n)$
$\mathbb{R}^d$ with L_1-norm	$1+\varepsilon$	$-$	$\Omega(\varepsilon^{-d})$

Where do we stand ...

With Steiner points.

Points in $\mathbb{R}^d$ with Steiner points.

- Upper Bound - $\Omega(\varepsilon^{(1-d)/2} \log n)$.
- Lower Bound - $\Omega(f(n))$ for some function $f(n)$, $\lim_{n \to \infty} f(n) = \infty$.

Under L_1 norm.

- Lower bounds - $\Omega(\varepsilon^{-2} / \log \varepsilon^{-1})$ in $\mathbb{R}^2$ and is $\Omega(\varepsilon^{-d})$ in $\mathbb{R}^d$ for $d \geq 3$.

Where do we stand ...

With Steiner points.

Points in $\mathbb{R}^d$ with Steiner points.

- Upper Bound - $\Omega(\varepsilon^{(1-d)/2} \log n)$.
- Lower Bound - $\Omega(f(n))$ for some function $f(n)$, $\lim_{n\to\infty} f(n) = \infty$.

Under L_1 norm.

- Lower bounds - $\Omega(\varepsilon^{-2}/\log \varepsilon^{-1})$ in $\mathbb{R}^2$ and is $\Omega(\varepsilon^{-d})$ in $\mathbb{R}^d$ for $d \geq 3$.

Future directions.

Without Steiner points, Log-dependence is unavoidable, due to LB in $\mathbb{R}$.

Q: Does the competitive ratio of an online $(1+\varepsilon)$-spanner algorithm for n points in $\mathbb{R}^d$ necessarily grow proportionally with $\varepsilon^{-f(d)} \cdot \log n$, where $\lim_{d\to\infty} f(d) = \infty$?

Thank you for your attention!