On Notions of Distortion and an Almost Minimum Spanning Tree with Constant Average Distortion

Yair Bartal[1] **Arnold Filtser**[2] Ofer Neiman[2]

[1] Hebrew University of Jerusalem
[2] Ben-Gurion University of the Negev

January 11, 2016

Warning

The talk is about an improved version of the paper.

Warning

The talk is about an improved version of the paper.

For details: www.cs.bgu.ac.il/~arnoldf/

Embedding

Embedding

$(X, d_X), (Y, d_Y)$ metric spaces.
$f : (X, d_X) \to (Y, d_Y)$ is called an **embedding**.

Embedding

Embedding

$(X, d_X), (Y, d_Y)$ metric spaces.
$f : (X, d_X) \to (Y, d_Y)$ is called an **embedding**.

Distortion

f has **distortion** t if for every $x, y \in X$,

$$d_X(x, y) \leq d_Y(f(x), f(y)) \leq t \cdot d_X(x, y) \ .$$

Embedding

Embedding

$(X, d_X), (Y, d_Y)$ metric spaces.
$f : (X, d_X) \to (Y, d_Y)$ is called an **embedding**.

Distortion

f has **distortion** t if for every $x, y \in X$,

$$d_X(x, y) \le d_Y(f(x), f(y)) \le t \cdot d_X(x, y) \ .$$

Average distortion

$$\frac{1}{\binom{|X|}{2}} \cdot \sum_{v,u \in X} \frac{d_Y(f(v), f(u))}{d_X(v, u)}$$

Spanner

Graph spanner

Given a weighted graph $G = (V, E, w)$, a **subgraph** $H = (V, E_H, w)$ of G is a **spanner** of G with **distortion** t if

$$\forall u, v \in V, \quad d_H(u, v) \leq t \cdot d_G(u, v)$$

Spanner

Graph spanner

Given a weighted graph $G = (V, E, w)$, a **subgraph** $H = (V, E_H, w)$ of G is a **spanner** of G with **distortion** t if

$$\forall u, v \in V, \quad d_H(u, v) \leq t \cdot d_G(u, v)$$

The **lightness** of a H is $$\Psi(H) = \frac{\sum_{e \in E_H} w(e)}{w(MST)} .$$

Spanner

Graph spanner

Given a weighted graph $G = (V, E, w)$, a **subgraph** $H = (V, E_H, w)$ of G is a **spanner** of G with **distortion** t if

$$\forall u, v \in V, \quad d_H(u, v) \leq t \cdot d_G(u, v)$$

The **lightness** of a H is $$\Psi(H) = \frac{\sum_{e \in E_H} w(e)}{w(MST)} .$$

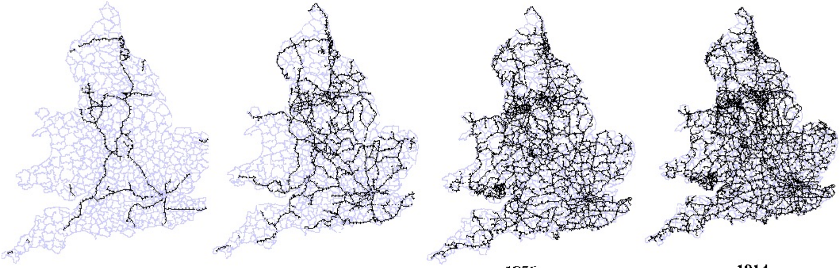

1845 1854 1876 1914

Lightness vs Average Distortion in Trees

$G = (V, E, w)$ is a weighted graph.

Lightness vs Average Distortion in Trees

$G = (V, E, w)$ is a weighted graph.

The **MST** has **lightness** 1!

Lightness vs Average Distortion in Trees

$G = (V, E, w)$ is a weighted graph.

The **MST** has **lightness** 1!
But unbounded **average distortion**...

Lightness vs Average Distortion in Trees

$G = (V, E, w)$ is a weighted graph.

The **MST** has **lightness** 1!
But unbounded **average distortion**...

Theorem (Abraham, Bartal and Neiman 2006)

Every weighted graph contains a **spanning tree** *with* $O(1)$ **average distortion**.

Lightness vs Average Distortion in Trees

$G = (V, E, w)$ is a weighted graph.

The **MST** has **lightness** 1!
But unbounded **average distortion**...

Theorem (Abraham, Bartal and Neiman 2006)

Every weighted graph contains a **spanning tree** *with* $O(1)$ **average distortion**.
But unbounded **lightness**...

Main result

Theorem (Constant lightness and average distortion (This work))

For any parameter $0 < \rho < 1$,
every weighted graph contains a **spanning tree** *with*

Main result

Theorem (Constant lightness and average distortion (This work))

For any parameter $0 < \rho < 1$,
every weighted graph contains a **spanning tree** *with*

- $1 + \rho$ *lightness.*

Main result

Theorem (Constant lightness and average distortion (This work))

For any parameter $0 < \rho < 1$*,*
every weighted graph contains a **spanning tree** *with*

- $1 + \rho$ *lightness.*
- $O(1/\rho)$ **average distortion**.

Main result

Theorem (Constant lightness and average distortion (This work))

For any parameter $0 < \rho < 1$,
every weighted graph contains a **spanning tree** *with*

- $1 + \rho$ *lightness.*
- $O(1/\rho)$ **average distortion**.

Tight!

Prioritized Distortion

$f:(X,d_X)\to(Y,d_Y)$. The **Distortion** of f is $\max_{x,y}\frac{d_Y(f(x),f(y))}{d_X(x,y)}$.

Prioritized Distortion

$f : (X, d_X) \to (Y, d_Y)$. The **Distortion** of f is $\max_{x,y} \frac{d_Y(f(x),f(y))}{d_X(x,y)}$.

Given a monotone increasing function $\alpha : \mathbb{N} \to \mathbb{R}_+$

Prioritized Distortion

$f : (X, d_X) \to (Y, d_Y)$. The **Distortion** of f is $\max_{x,y} \frac{d_Y(f(x),f(y))}{d_X(x,y)}$.

Given a monotone increasing function $\alpha : \mathbb{N} \to \mathbb{R}_+$
Priority $\pi = (x_1, \ldots, x_n)$.

Prioritized Distortion

$f : (X, d_X) \to (Y, d_Y)$. The **Distortion** of f is $\max_{x,y} \frac{d_Y(f(x),f(y))}{d_X(x,y)}$.

Given a monotone increasing function $\alpha : \mathbb{N} \to \mathbb{R}_+$
Priority $\pi = (x_1, \ldots, x_n)$.

Priority Distortion

$f : X \to Y$ has **priority distortion** α w.r.t. π if

$$\forall x_j, y \in X \quad d_X(x_j, y) \leq d_Y(f(x_j), f(y)) \leq \alpha(j) \cdot d_X(x_j, y)$$

Prioritized Distortion

$f : (X, d_X) \to (Y, d_Y)$. The **Distortion** of f is $\max_{x,y} \frac{d_Y(f(x),f(y))}{d_X(x,y)}$.

Given a monotone increasing function $\alpha : \mathbb{N} \to \mathbb{R}_+$
Priority $\pi = (x_1, \ldots, x_n)$.

Priority Distortion

$f : X \to Y$ has **priority distortion** α w.r.t. π if

$$\forall x_j, y \in X \quad d_X(x_j, y) \leq d_Y(f(x_j), f(y)) \leq \alpha(j) \cdot d_X(x_j, y)$$

Theorem (Prioritized Spanner (This work))

Given a graph $G = (V, E)$, parameter $0 < \rho < 1$ and any priority **ranking** *π of V, there exists a* **spanner** *H with* **lightness** $1 + \rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$.

Scaling Distortion

Scaling Distortion

Embedding $f : X \to Y$ has **scaling distortion** $\beta : (0,1) \to \mathbb{R}_+$ if

$\forall \epsilon \in (0,1)$ **at least** $(1-\epsilon)$-fraction of the pairs suffer distortion **at most** $\beta(\epsilon)$.

Scaling Distortion

Scaling Distortion

Embedding $f : X \to Y$ has **scaling distortion** $\beta : (0,1) \to \mathbb{R}_+$ if

$$\forall \epsilon \in (0,1) \textbf{ at least } (1-\epsilon)\text{-fraction of the pairs suffer distortion } \textbf{at most } \beta(\epsilon) .$$

Theorem (Abraham, Bartal and Neiman 2006)

Every weighted graph contains a **spanning tree** *with* **scaling distortion** $O(1/\sqrt{\epsilon})$.

Scaling Distortion

Scaling Distortion

Embedding $f : X \to Y$ has **scaling distortion** $\beta : (0,1) \to \mathbb{R}_+$ if

$\forall \epsilon \in (0,1)$ **at least** $(1-\epsilon)$-fraction of the pairs suffer distortion **at most** $\beta(\epsilon)$.

Theorem (Abraham, Bartal and Neiman 2006)

Every weighted graph contains a **spanning tree** *with* **scaling distortion** $O(1/\sqrt{\epsilon})$.

Scaling Distortion implies constant average distortion

If f has **scaling distortion** $O(\frac{1}{\epsilon^{1-\delta}})$ for $\delta > 0$ then

$$\textbf{Average Distortion} = O(1) \ .$$

Priority implies Scaling - main technical contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X),

Priority implies Scaling - main technical contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X)*,*
there exists a **priority ranking** $\pi = (x_1, \ldots, x_n)$

Priority implies Scaling - main technical contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X)*,*
there exists a **priority ranking** $\pi = (x_1, \ldots, x_n)$
s.t. every embedding with **priority distortion** α **w.r.t** π *into* (Y, d_Y)

Priority implies Scaling - main technical contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X)*,*
there exists a **priority ranking** $\pi = (x_1, \ldots, x_n)$
s.t. every embedding with **priority distortion** α **w.r.t** π *into* (Y, d_Y)
has **scaling distortion** $O(\alpha(4/\epsilon))$.

Priority implies Scaling - main technical contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X)*,*
there exists a **priority ranking** $\pi = (x_1, \ldots, x_n)$
s.t. every embedding with **priority distortion** α **w.r.t** π *into* (Y, d_Y)
has **scaling distortion** $O(\alpha(4/\epsilon))$*.*

Scaling also implies priority!

Scaling Light Spanner

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}\left(\log j\right)/\rho$.

Scaling Light Spanner

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}\left(\log j\right)/\rho$.

Theorem (Priority implies Scaling)

Priority distortion α *w.r.t* π *implies* **scaling distortion** $O(\alpha(4/\epsilon))$.

Scaling Light Spanner

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$.

Theorem (Priority implies Scaling)

Priority distortion α *w.r.t* π *implies* **scaling distortion** $O(\alpha(4/\epsilon))$.

⇓ ⇓ ⇓

Theorem (Scaling Spanner)

Spanner *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}(\log(1/\epsilon))/\rho$.

Theorem (Scaling Spanner)

Spanner *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}\left(\log\left(1/\epsilon\right)\right)/\rho$.

Theorem (Scaling Spanner)

Spanner *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}\left(\log\left(1/\epsilon\right)\right)/\rho$.

Theorem (Abraham, Bartal and Neiman 2006)

Any w.graph contains a **spanning tree** *with* **scaling distortion** $O(1/\sqrt{\epsilon})$.

Theorem (Scaling Spanner)

Spanner *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}\left(\log\left(1/\epsilon\right)\right)/\rho$.

Theorem (Abraham, Bartal and Neiman 2006)

Any w.graph contains a **spanning tree** *with* **scaling distortion** $O(1/\sqrt{\epsilon})$.

Lemma (Scaling embeddings Composition)

If $\mathbf{f}:(X,d_X)\rightarrow(Y,d_Y)$ *(respectively,* $\mathbf{g}:(Y,d_Y)\rightarrow(Z,d_Z)$*) has scaling distortion* α *(resp.,* β*). Then* $\mathbf{f}\circ\mathbf{g}$ *has scaling distortion* $\alpha(\epsilon/2)\cdot\beta(\epsilon/2)$.

Theorem (Scaling Spanner)

Spanner *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}\left(\log\left(1/\epsilon\right)\right)/\rho$.

Theorem (Abraham, Bartal and Neiman 2006)

Any w.graph contains a **spanning tree** *with* **scaling distortion** $O(1/\sqrt{\epsilon})$.

Lemma (Scaling embeddings Composition)

If $\mathbf{f}:(X,d_X)\rightarrow(Y,d_Y)$ *(respectively,* $\mathbf{g}:(Y,d_Y)\rightarrow(Z,d_Z)$*) has scaling distortion* α *(resp.,* β*). Then* $\mathbf{f}\circ\mathbf{g}$ *has scaling distortion* $\alpha(\epsilon/2)\cdot\beta(\epsilon/2)$.

⇩ ⇩ ⇩

Theorem

Spanning tree *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}(\sqrt{1/\epsilon})/\rho$.

Light Tree with Constant Average Distortion

Theorem

Spanning tree *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}(\sqrt{1/\epsilon})/\rho$.

Light Tree with Constant Average Distortion

Theorem

Spanning tree *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}(\sqrt{1/\epsilon})/\rho$.

Scaling distortion $O(\frac{1}{\epsilon^{1-\delta}})$ implies $O(1)$ **average distortion**.

Light Tree with Constant Average Distortion

Theorem

Spanning tree *with* **lightness** $1+\rho$ *and* **scaling distortion** $\tilde{O}(\sqrt{1/\epsilon})/\rho$.

Scaling distortion $O(\frac{1}{\epsilon^{1-\delta}})$ implies $O(1)$ **average distortion**.

⇩ ⇩ ⇩

Corollary (Main result)

Spanning tree *with* **lightness** $1+\rho$ *and* **average distortion** $O(1/\rho)$.

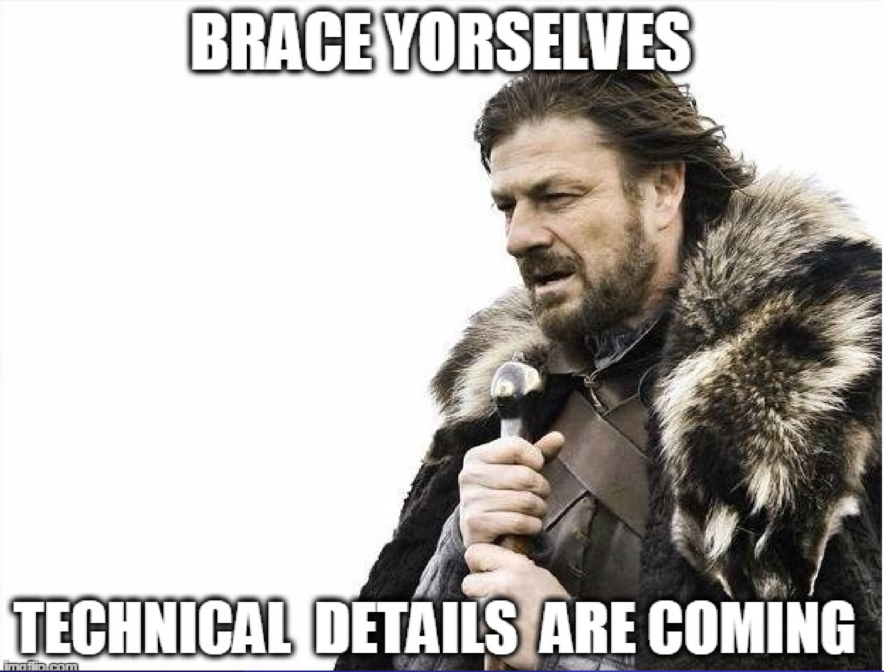
BRACE YORSELVES
TECHNICAL DETAILS ARE COMING
imgflip.com

Priority implies Scaling - main tequniqal contribution

Theorem (Priority implies scaling)

Given a metric space (X, d_X)*,*
there exists a **priority ranking** $\pi = (x_1, \dots, x_n)$
s.t. every embedding with **priority distortion** α **w.r.t** π *into* (Y, d_Y)
has **scaling distortion** $O(\alpha(4/\epsilon))$*.*

ϵ-Density Net

Given $x \in X$ and $\epsilon \in (0,1)$, $R(x,\epsilon)$ is the **minimal** radius r such that

$$|B_X(x,r)| \geq \epsilon \cdot n$$

ϵ-Density Net

Given $x \in X$ and $\epsilon \in (0,1)$, $R(x,\epsilon)$ is the **minimal** radius r such that

$$|B_X(x,r)| \geq \epsilon \cdot n$$

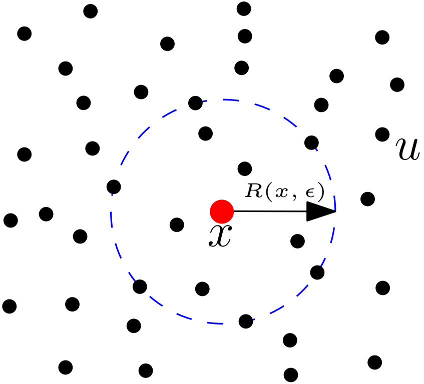

ϵ-Density Net

ϵ**-Density Net** is a subset $N \subseteq X$ such that:

- $\forall x \in X$ there exists $y \in N$ such that $d_X(x, y) \leq 2R(x, \epsilon)$.
- $|N| \leq 1/\epsilon$.

ϵ-Density Net is a subset $N \subseteq X$ such that:

- $\forall x \in X$ there exists $y \in N$ such that $d_X(x, y) \leq 2R(x, \epsilon)$.
- $|N| \leq 1/\epsilon$.

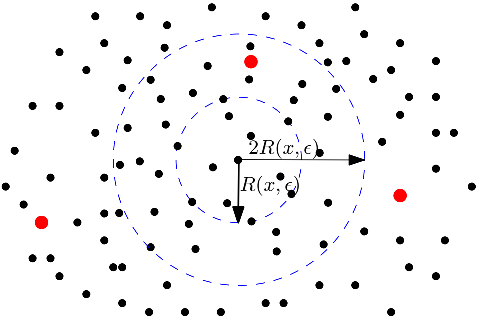

ϵ-Density Net

ϵ-Density Net is a subset $N \subseteq X$ such that:

- $\forall x \in X$ there exists $y \in N$ such that $d_X(x, y) \leq 2R(x, \epsilon)$.
- $|N| \leq 1/\epsilon$.

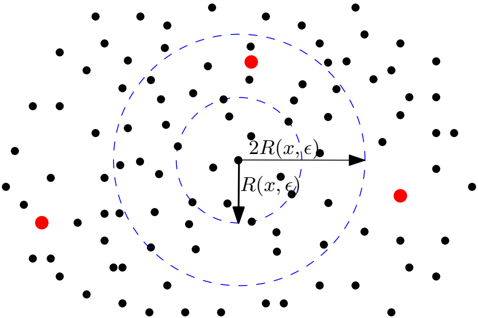

Theorem (H.Chan, M.Dinitz and A.Gupta 2006)

For every metric space and $\epsilon \in (0, 1)$ there **exists** *an ϵ-density-net.*

Priority implies Scaling - Proof

For $1 \leq i \leq \lceil \log n \rceil$ set $\epsilon_i = 2^{-i}$. Let N_i be an ϵ_i-density net.

Priority implies Scaling - Proof

For $1 \leq i \leq \lceil \log n \rceil$ set $\epsilon_i = 2^{-i}$. Let N_i be an ϵ_i-density net.

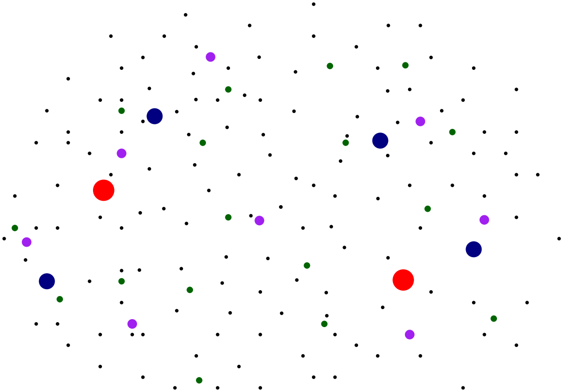

Priority implies Scaling - Proof

For $1 \leq i \leq \lceil \log n \rceil$ set $\epsilon_i = 2^{-i}$. Let N_i be an ϵ_i-density net.

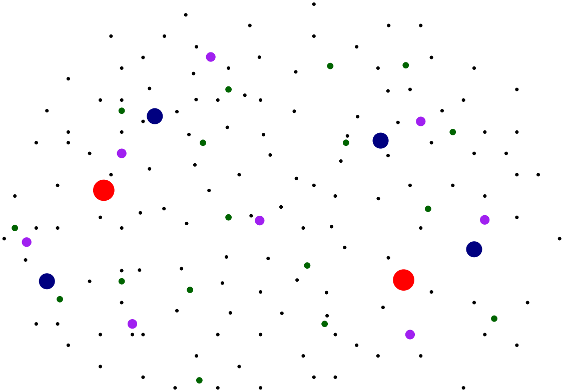

Permutation selection:

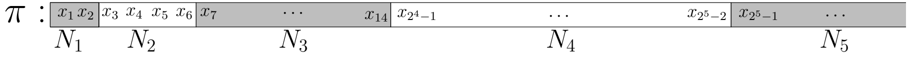

Prioritized Light Spanner

Theorem (Prioritized Spanner)

Given a graph $G = (V, E)$, parameter $0 < \rho < 1$ and any priority **ranking** *π of V, there exists a* **spanner** *H with* **lightness** $1 + \rho$ *and* **prioritized distortion** $\tilde{O}(\log j) / \rho$.

Prioritized Light Spanner

Theorem (Prioritized Spanner)

Given a graph $G = (V, E)$, parameter $0 < \rho < 1$ and any priority **ranking** π *of* V, *there exists a* **spanner** H *with* **lightness** $1 + \rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$.

Lemma (Terminal light spanner)

Given a graph $G = (V, E)$, a subset $K \subseteq V$ of **terminals** *of size k, and a parameter $0 < \delta < 1$, there exists a spanner H that:*

Prioritized Light Spanner

Theorem (Prioritized Spanner)

Given a graph $G = (V, E)$*, parameter* $0 < \rho < 1$ *and any priority* **ranking** π *of* V*, there exists a* **spanner** H *with* **lightness** $1 + \rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$*.*

Lemma (Terminal light spanner)

Given a graph $G = (V, E)$*, a subset* $K \subseteq V$ *of* **terminals** *of size* k*, and a parameter* $0 < \delta < 1$*, there exists a spanner* H *that:*

1) *Contains the* **MST** *of* G*.*

Prioritized Light Spanner

Theorem (Prioritized Spanner)

Given a graph $G = (V, E)$, parameter $0 < \rho < 1$ and any priority **ranking** *π of V, there exists a* **spanner** *H with* **lightness** $1 + \rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$.

Lemma (Terminal light spanner)

Given a graph $G = (V, E)$, a subset $K \subseteq V$ of **terminals** *of size k, and a parameter $0 < \delta < 1$, there exists a spanner H that:*

1) *Contains the* **MST** *of G.*

2) *Has* **lightness** $1 + \delta$.

Prioritized Light Spanner

Theorem (Prioritized Spanner)

Given a graph $G = (V, E)$, parameter $0 < \rho < 1$ and any priority **ranking** *π of V, there exists a* **spanner** *H with* **lightness** *$1 + \rho$ and* **prioritized distortion** *$\tilde{O}(\log j)/\rho$.*

Lemma (Terminal light spanner)

Given a graph $G = (V, E)$, a subset $K \subseteq V$ of **terminals** *of size k, and a parameter $0 < \delta < 1$, there exists a spanner H that:*

1) *Contains the* **MST** *of G.*

2) *Has* **lightness** *$1 + \delta$.*

3) *Every pair in $K \times V$ has* **distortion** *$O\left(\frac{\log k}{\delta}\right)$.*

Previus results and terminal spanner

Theorem (Chechik and Wulff-Nilsen (SODA 16), following Chandra et.al and Elkin et.al.)

For every weighted n-vertex graph G and parameters $t > 1, \epsilon > 0$ there exist a $(2t-1)(1+\epsilon)$ **spanner** *of* **lightness** $O_\epsilon(n^{1/t})$.

Previus results and terminal spanner

Theorem (Chechik and Wulff-Nilsen (SODA 16), following Chandra et.al and Elkin et.al.)

For every weighted n-vertex graph G and parameters $t > 1, \epsilon > 0$ *there exist a* $(2t-1)(1+\epsilon)$ **spanner** *of* **lightness** $O_\epsilon(n^{1/t})$.

For $t = \log n$ and $\epsilon = 1$, they get $O(\log n)$-**spanner** with **lightness** $O(1)$.

General reduction to light spanners

Theorem (Light spanners reduction)

Suppose that for every n vertex graph G there is a spanner H that:

1. *Has* **lightness** ℓ.
2. *Has* **distortion** t.

General reduction to light spanners

Theorem (Light spanners reduction)

Suppose that for every n vertex graph G there is a spanner H that:

1. *Has* **lightness** ℓ.
2. *Has* **distortion** t.

Then, for every n vertex graph G and **parameter** $0 < \delta < 1$,
there is a **spanner** *H that:*

General reduction to light spanners

Theorem (Light spanners reduction)

Suppose that for every n vertex graph G there is a spanner H that:

1. *Has* **lightness** ℓ.
2. *Has* **distortion** t.

Then, for every n vertex graph G and **parameter** $0 < \delta < 1$, *there is a* **spanner** *H that:*

1. *Has* **lightness** $1 + \delta\ell$.
2. *Has* **distortion** t/δ.
3. *Contains the* **MST**.

General reduction to light spanners

Theorem (Light spanners reduction)

Suppose that for every n vertex graph G there is a spanner H that:

1. *Has* **lightness** ℓ. $O(1)$
2. *Has* **distortion** t. $O(\log n)$

Then, for every n vertex graph G and **parameter** $0 < \delta < 1$, *there is a* **spanner** H *that:*

1. *Has* **lightness** $1 + \delta\ell$.
2. *Has* **distortion** t/δ.
3. *Contains the* **MST**.

General reduction to light spanners

Theorem (Light spanners reduction)

Suppose that for every n vertex graph G there is a spanner H that:

1. *Has* **lightness** ℓ. $O(1)$
2. *Has* **distortion** t. $O(\log n)$

Then, for every n vertex graph G and **parameter** $0 < \delta < 1$, *there is a* **spanner** H *that:*

1. *Has* **lightness** $1 + \delta\ell$. $1 + \delta$
2. *Has* **distortion** t/δ. $O(\log n)/\delta$
3. *Contains the* **MST**.

Open Problems

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}(\log j)/\rho$.

Open Problems

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}\left(\log j\right)/\rho$.

Is it possible to get **prioritized distortion** $O\left(\log j\right)/\rho$?

Open Problems

Theorem (Prioritized Spanner)

Spanner with **lightness** $1+\rho$ *and* **prioritized distortion** $\tilde{O}\left(\log j\right)/\rho$.

Is it possible to get **prioritized distortion** $O\left(\log j\right)/\rho$?

Efficient implimintation

While the current implementation is polynomial, it is still far from practical.