European Symposium on Algorithms (ESA) 2026, L'Aquila, Italy

Hierarchical Spanners

Davide Bilò[1] Luciano Gualà[2] Stefano Leucci[1] Guido Proietti[1] Alessandro Straziota[2]





[1] University of L'Aquila, Italy
[2] *Tor Vergata* University of Rome, Italy

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \alpha \cdot d_G(u, v) + \beta, \quad \forall (u, v) \in P.$$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size $+$ low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Choice of P

$P = \{(s,t)\}$ $\longrightarrow$ single-pair

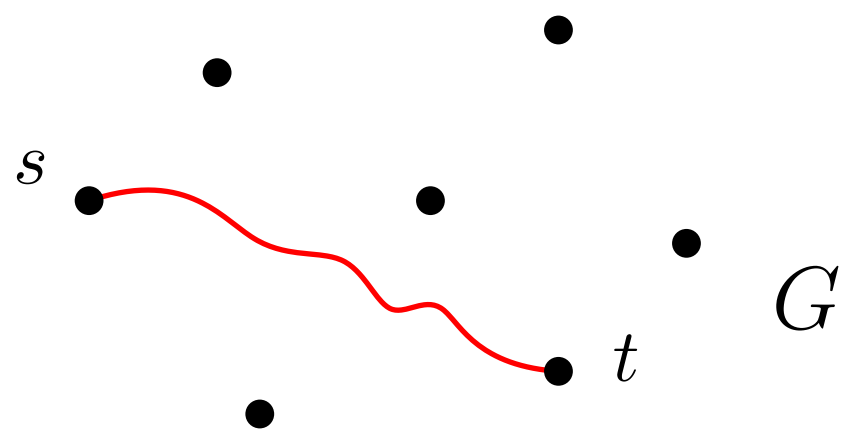

$P = \{(s,t)\}$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Choice of P

$P = \{(s,t)\}$ $\longrightarrow$ single-pair

$P = V \times V$ $\longrightarrow$ all-pair

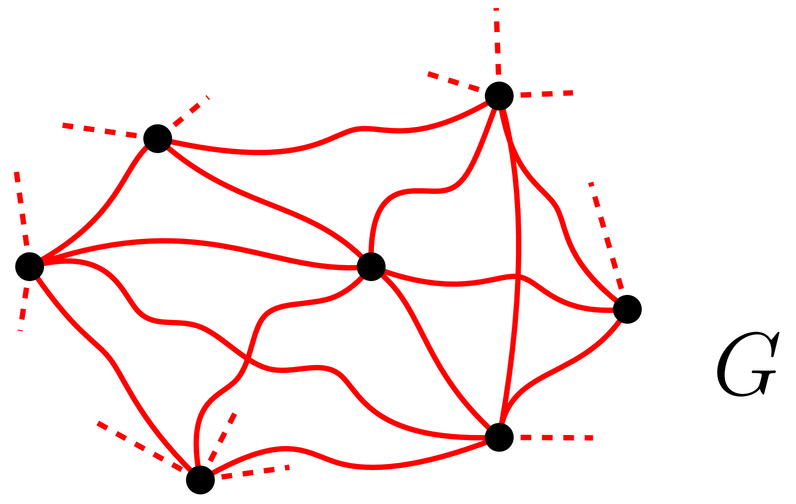

$P = V \times V$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Choice of P

$P = \{(s,t)\}$ $\longrightarrow$ single-pair

$P = V \times V$ $\longrightarrow$ all-pair

$P = S \times S$ $\longrightarrow$ subsetwise

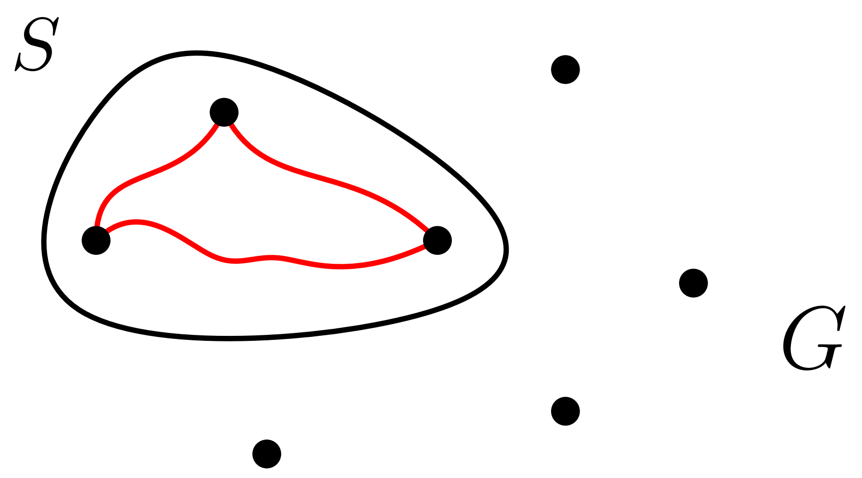

$P = S \times S$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Choice of P

$P = \{(s, t)\} \longrightarrow$ single-pair

$P = S \times T \longrightarrow S \times T$ spanner

$P = V \times V \longrightarrow$ all-pair

$P = S \times S \longrightarrow$ subsetwise

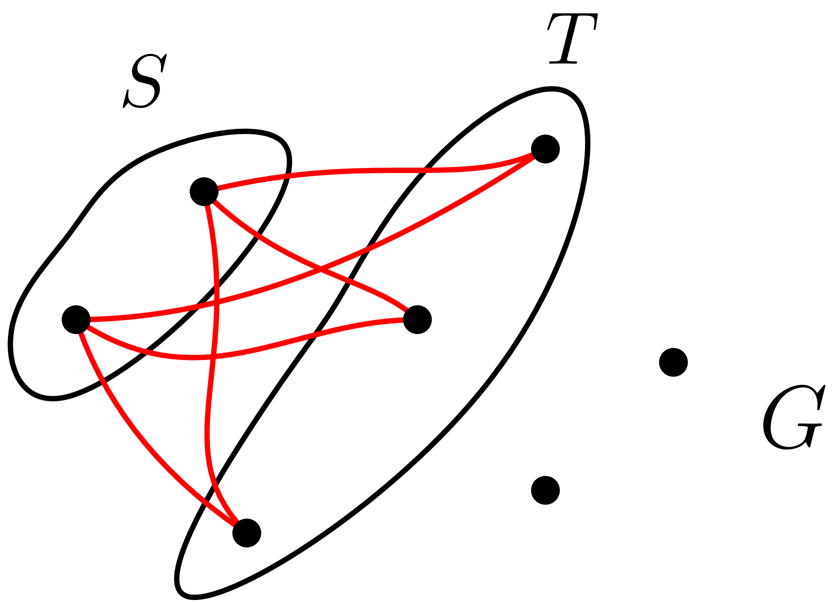

$P = S \times S$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

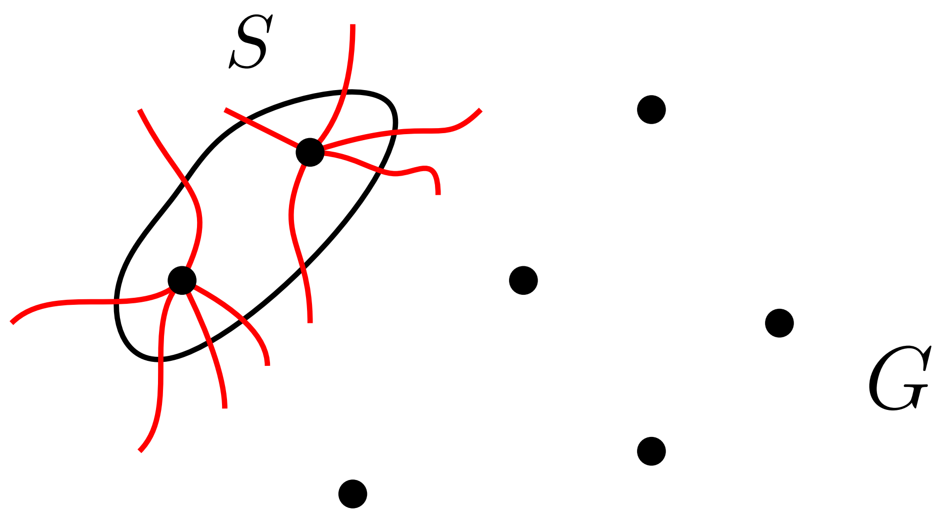

$P = S \times V$

Choice of P

$P = \{(s,t)\}$ $\longrightarrow$ single-pair

$P = V \times V$ $\longrightarrow$ all-pair

$P = S \times S$ $\longrightarrow$ subsetwise

$P = S \times T$ $\longrightarrow$ $S \times T$ spanner

$P \subseteq S \times V$ $\longrightarrow$ sourcewise

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

Choice of the stretch

$(\alpha, 0)$-spanner $\longrightarrow$ α-spanner

$(1, \beta)$-spanner $\longrightarrow$ β-additive spanner

1-spanner $\longrightarrow$ **preserver**

Choice of P

$P = \{(s, t)\}$ $\longrightarrow$ single-pair

$P = V \times V$ $\longrightarrow$ all-pair

$P = S \times S$ $\longrightarrow$ subsetwise

$P = S \times T$ $\longrightarrow$ $S \times T$ spanner

$P \subseteq S \times V$ $\longrightarrow$ sourcewise

$P \subseteq V \times V$ $\longrightarrow$ pairwise

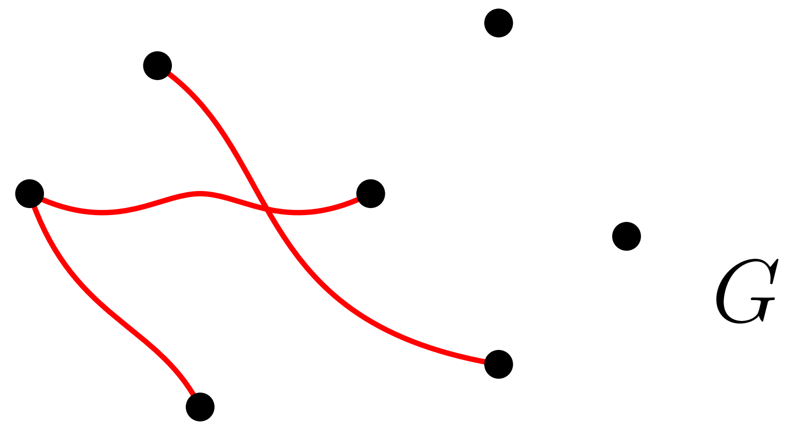

$P \subseteq V \times V$

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

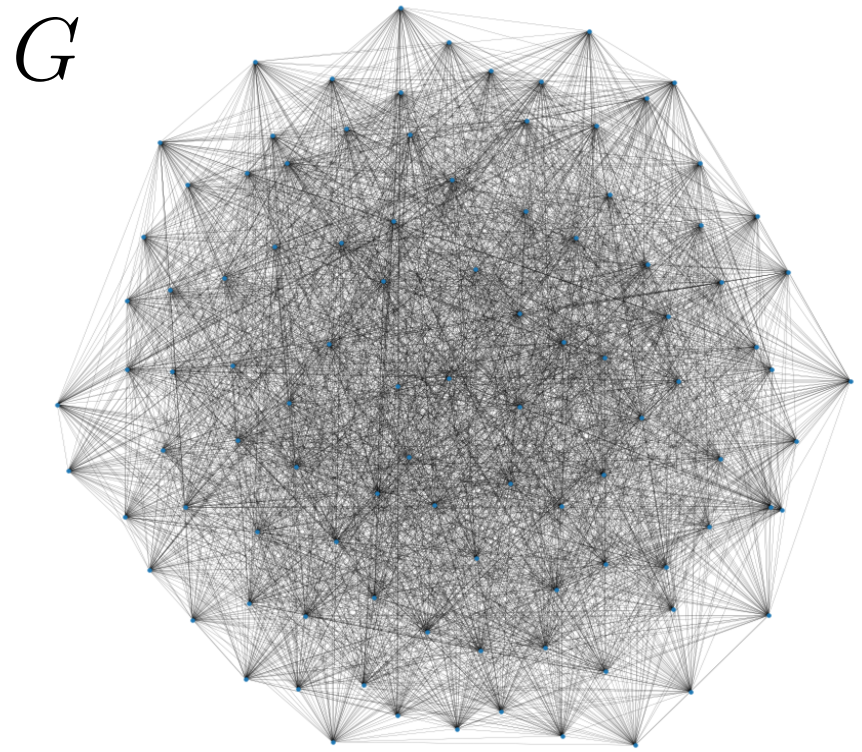

$n = 100$ vertices and $m = 3300$ edges

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u, v) \leq \boxed{\alpha} \cdot d_G(u, v) + \boxed{\beta}, \quad \forall (u, v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

G

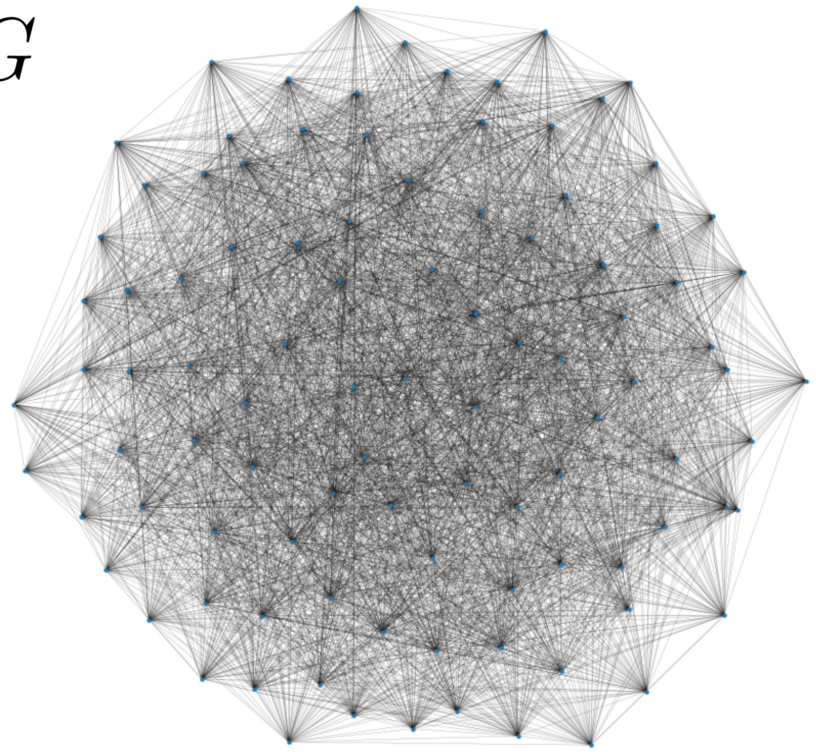

$\longrightarrow$

$H \subseteq G$

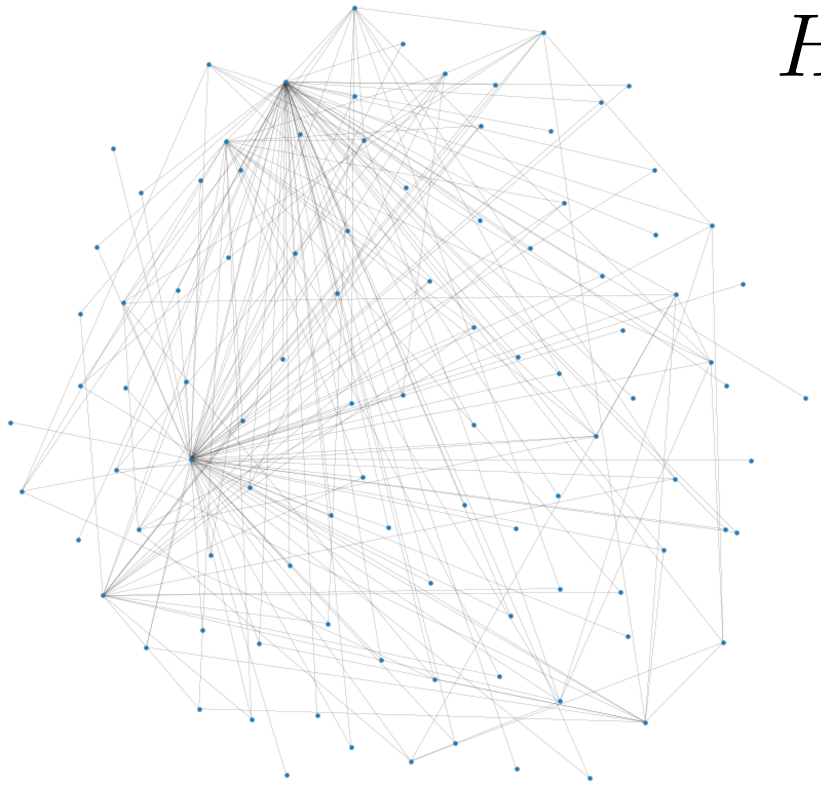

$n = 100$ vertices and $m = 3300$ edges

all-pairs 7-spanner with $m = 214$ edges

Graph Spanner

Graph Spanner. Giaven a graph $G = (V, E)$, and $\alpha \geq 1$ and $\beta \geq 0$, a subgraph $H \subseteq G$ is an (α, β)-spanner of G with respect to $P \subseteq V \times V$ if

$$d_H(u,v) \leq \boxed{\alpha} \cdot d_G(u,v) + \boxed{\beta}, \quad \forall (u,v) \in P.$$

multiplicative and *additive* *stretch* of H

Goal: small size + low stretch.

G

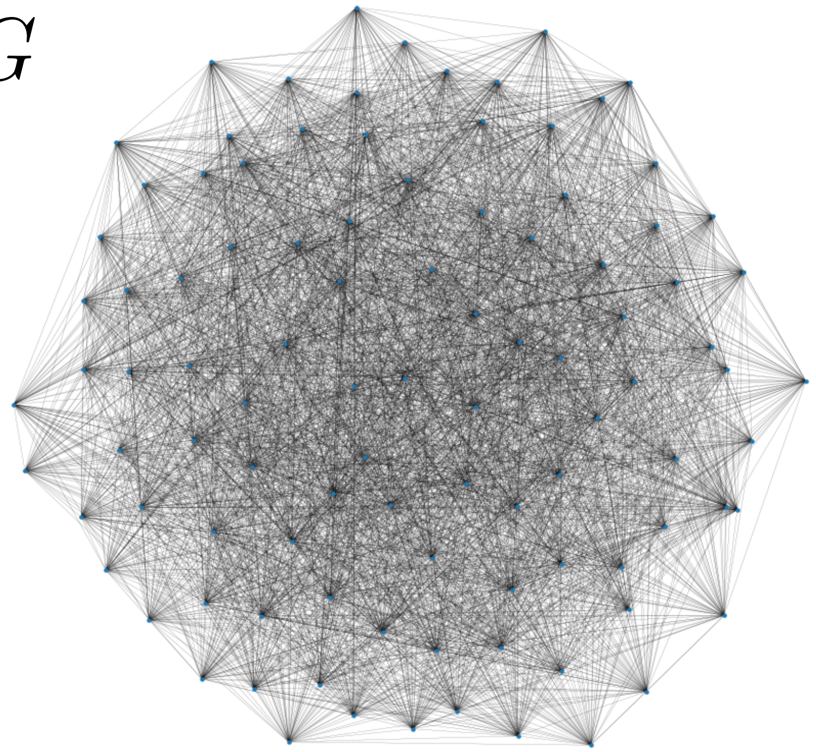

$\longrightarrow$

$H \subseteq G$

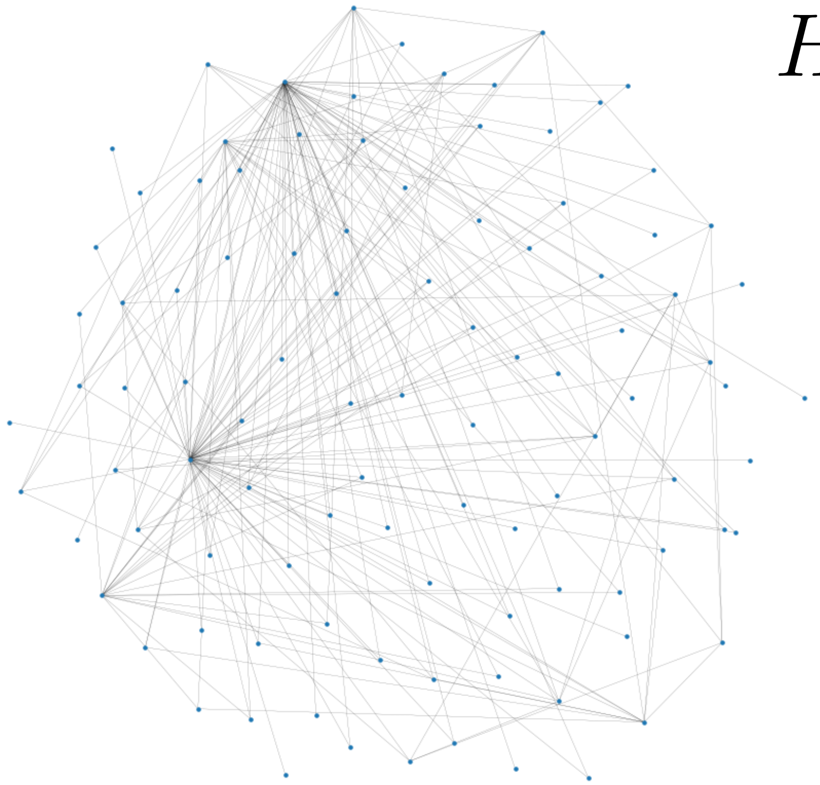

$n = 100$ vertices and $m = 3300$ edges

all-pairs 7-spanner with $m = 214$ edges
only 6.5% of edges of G!

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices

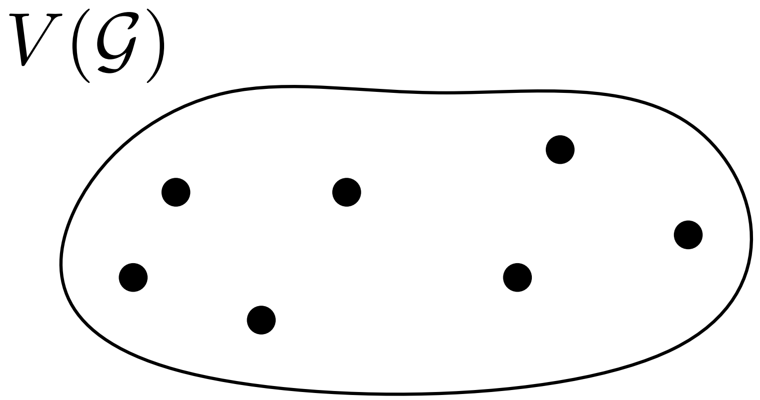

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices
- L pairwise disjoint subsets of edges $E_1, \ldots, E_L$

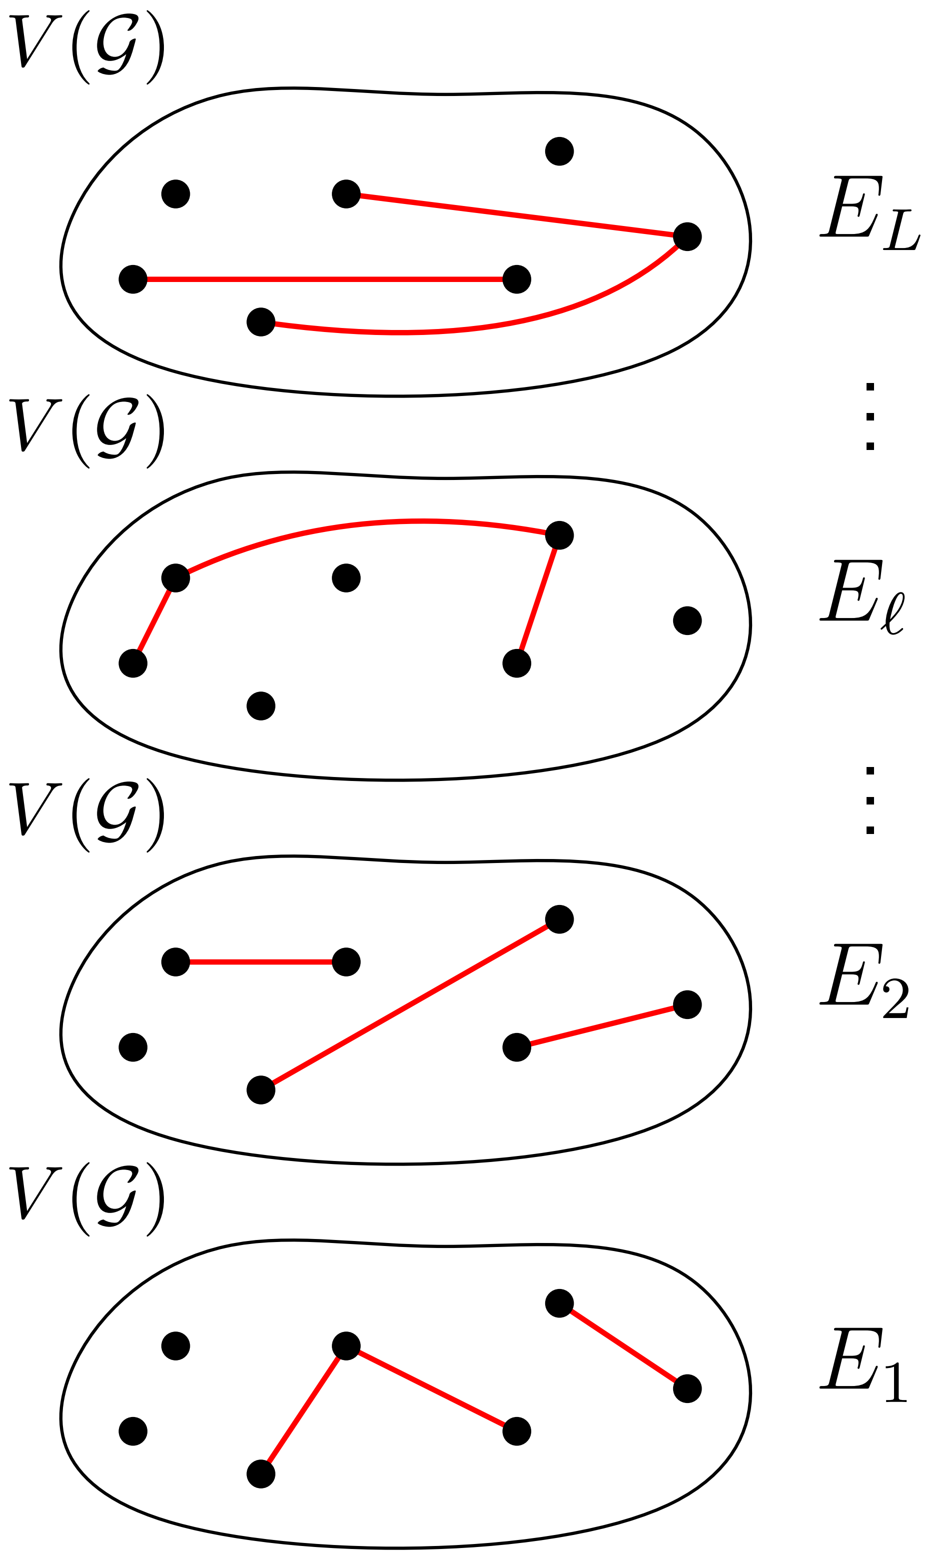

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices
- L pairwise disjoint subsets of edges $E_1, \ldots, E_L$
- *hierarchical graph* $\mathcal{G} = (V(\mathcal{G}), E_1, \ldots, E_L)$
- $E(\mathcal{G}) = \cup_{i=1}^{L} E_i$, with $m = |E(\mathcal{G})|$

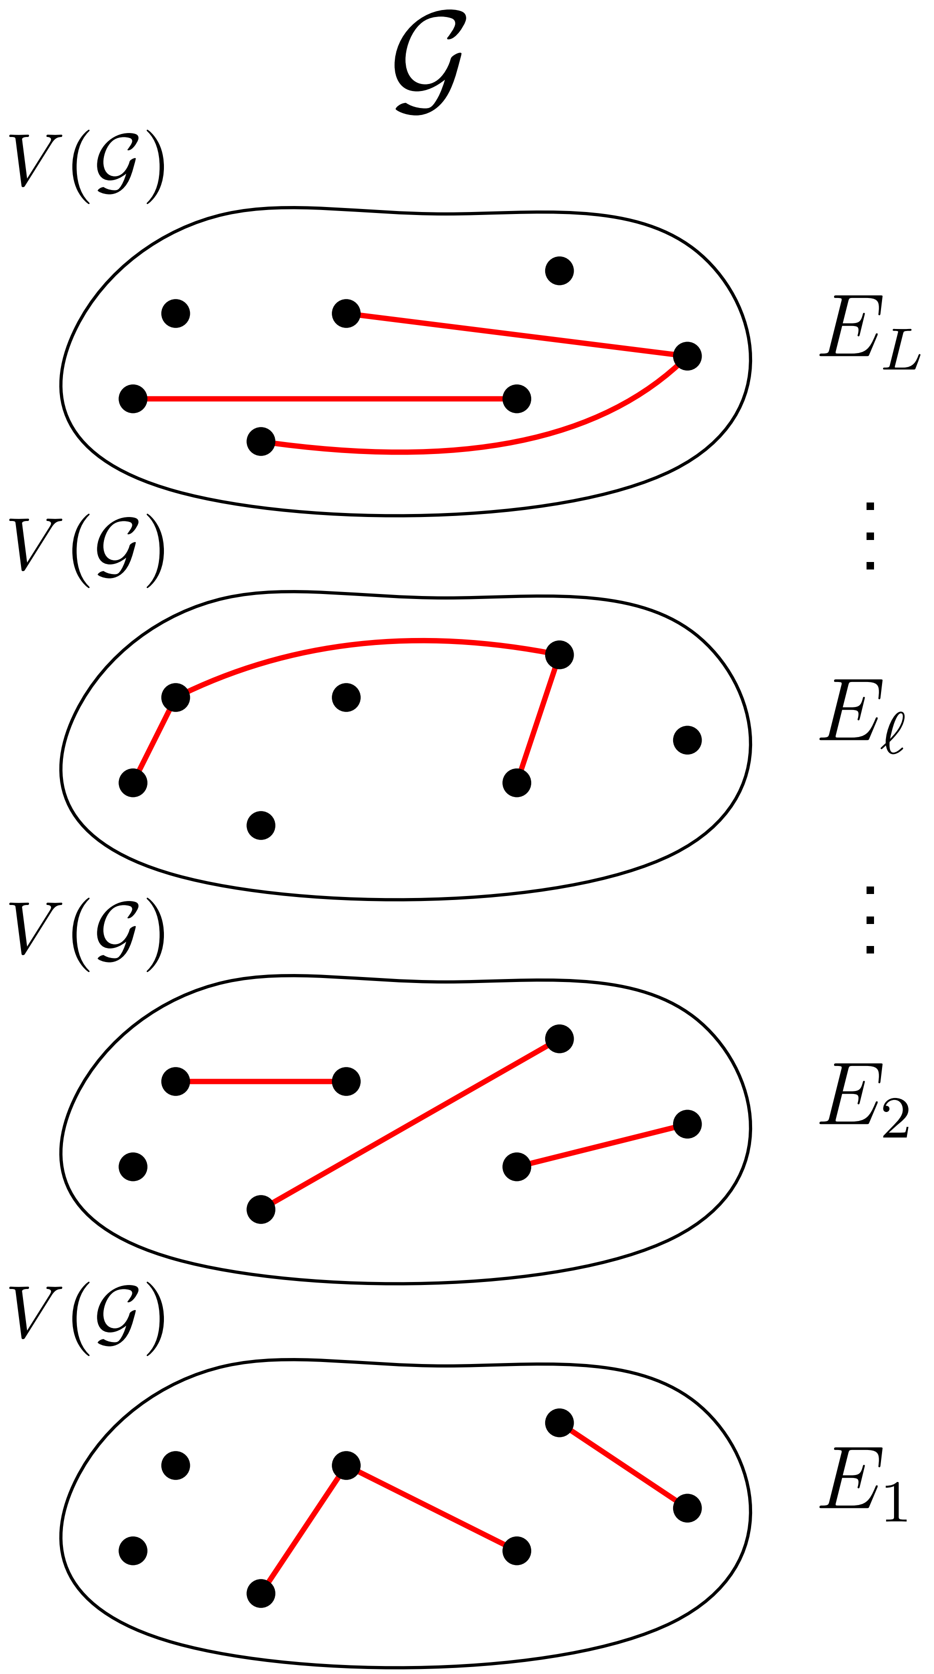

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices
- L pairwise disjoint subsets of edges $E_1, \ldots, E_L$
- *hierarchical graph* $\mathcal{G} = (V(\mathcal{G}), E_1, \ldots, E_L)$
- $E(\mathcal{G}) = \cup_{i=1}^{L} E_i$, with $m = |E(\mathcal{G})|$
- $G_\ell = (V(\mathcal{G}), \cup_{i=1}^{\ell} E_i)$, for $\ell = 0, 1, \ldots, L$

$$\emptyset = G_0 \subseteq G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L = G$$

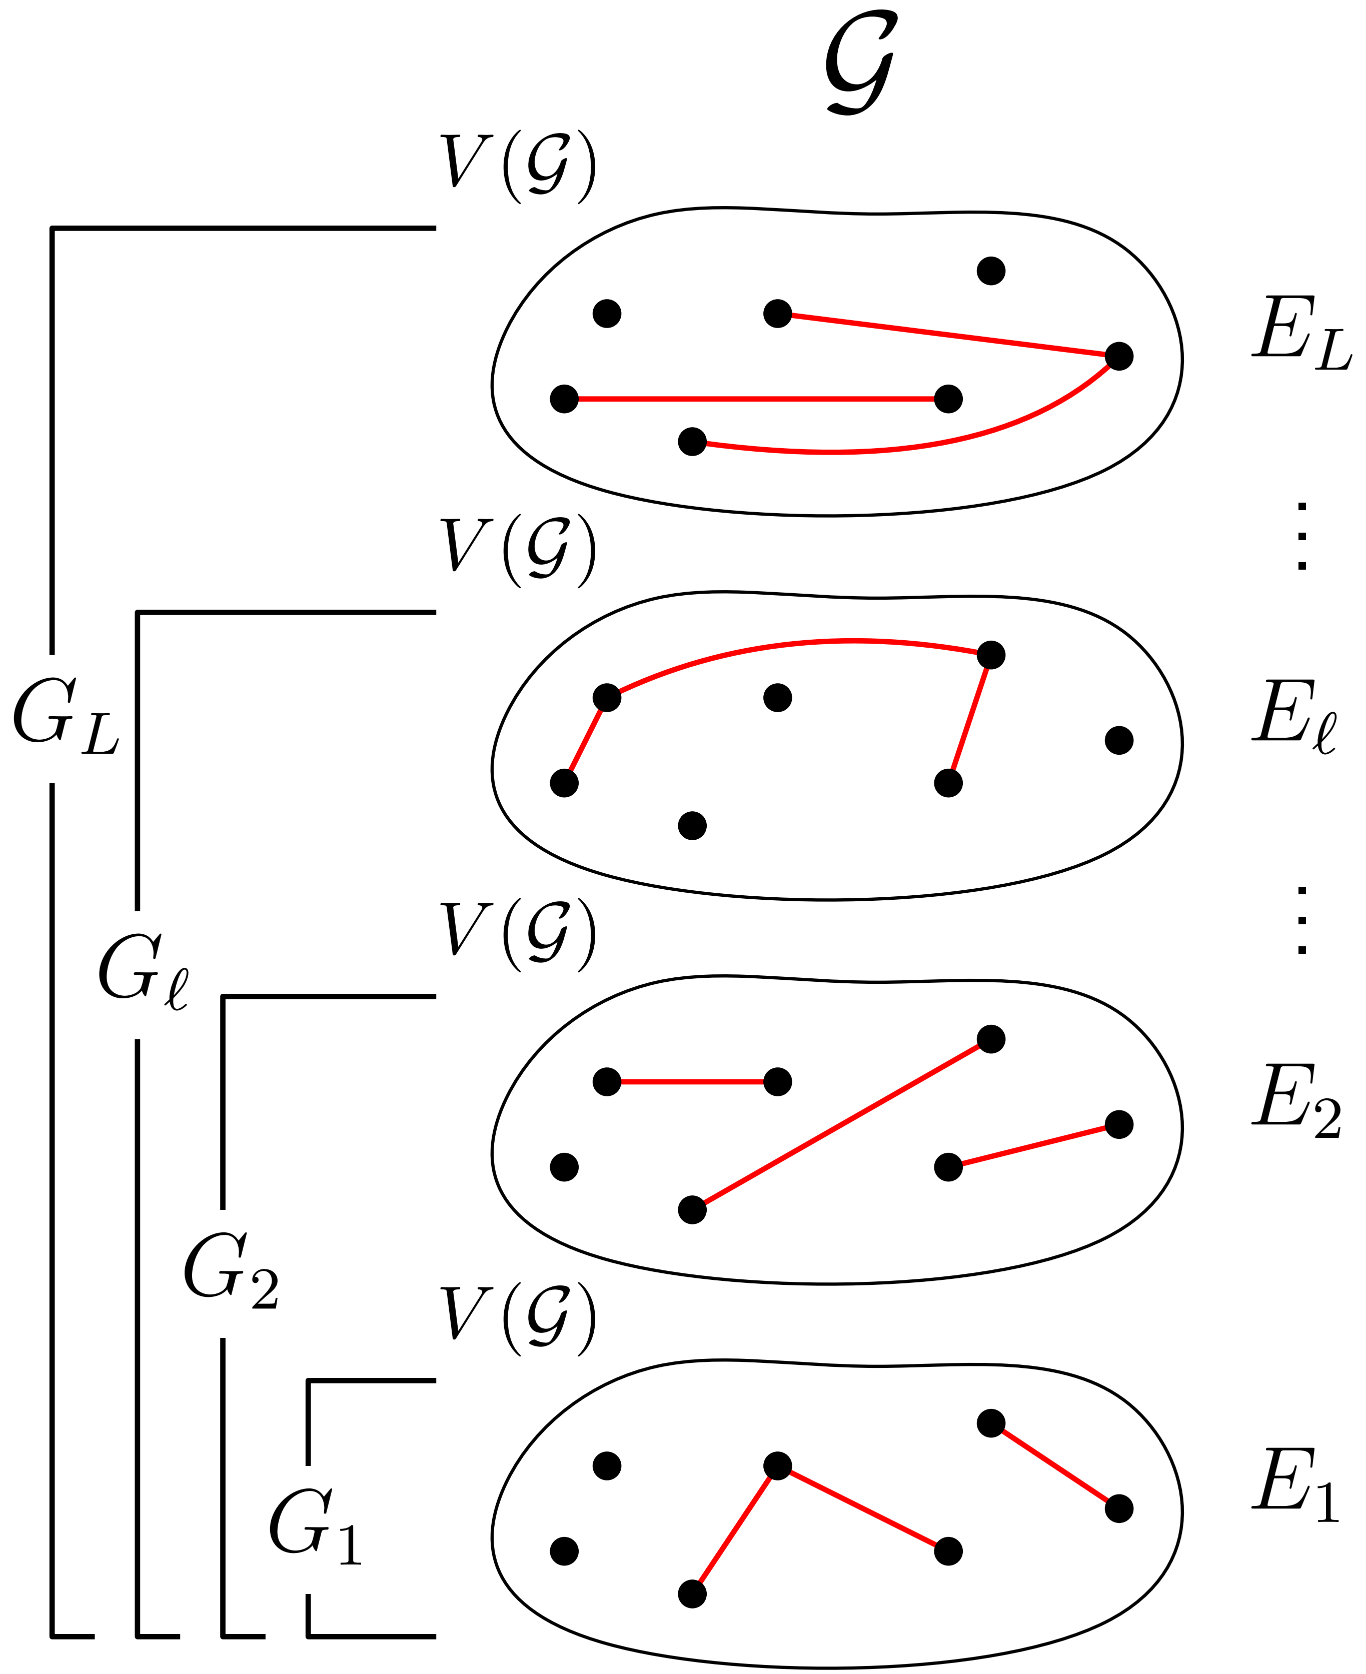

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices
- L pairwise disjoint subsets of edges $E_1, \ldots, E_L$
- *hierarchical graph* $\mathcal{G} = (V(\mathcal{G}), E_1, \ldots, E_L)$
- $E(\mathcal{G}) = \cup_{i=1}^{L} E_i$, with $m = |E(\mathcal{G})|$
- $G_\ell = (V(\mathcal{G}), \cup_{i=1}^{\ell} E_i)$, for $\ell = 0, 1, \ldots, L$

$$\emptyset = G_0 \subseteq G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L = G$$

Hierarchical Spanner. a hierarchical subgraph $\mathcal{H}$ is a *hierarchical* (α, β)*-spanner* of $\mathcal{G}$ w.r.t. $P \subseteq V(\mathcal{G}) \times V(\mathcal{G})$ if H_ℓ is an (α, β)-spanner of G_ℓ, for every layer $\ell = 1, \ldots, L$.

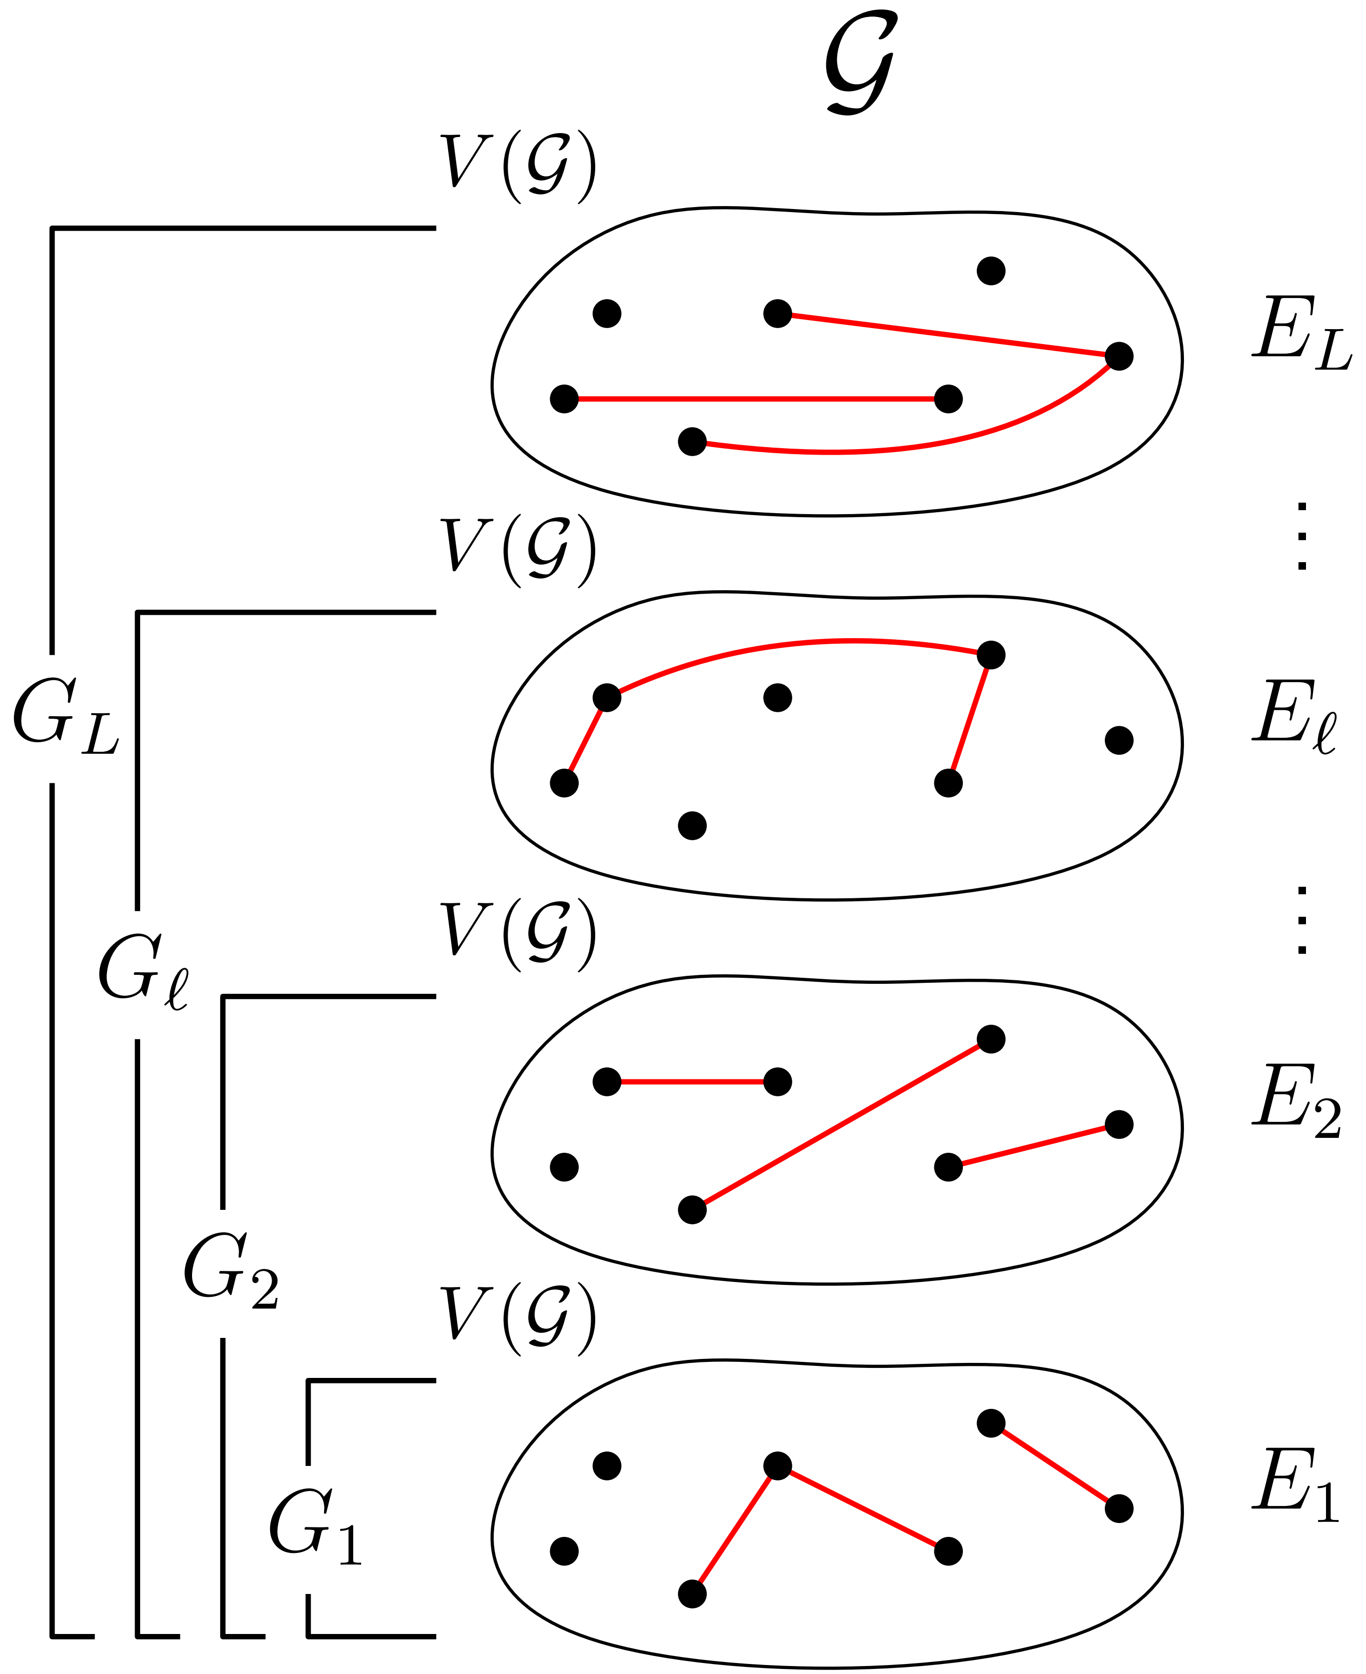

Hierarchical Graphs and Spanners

- a set $V(\mathcal{G})$ of n vertices
- L pairwise disjoint subsets of edges $E_1, \ldots, E_L$
- *hierarchical graph* $\mathcal{G} = (V(\mathcal{G}), E_1, \ldots, E_L)$
- $E(\mathcal{G}) = \cup_{i=1}^{L} E_i$, with $m = |E(\mathcal{G})|$
- $G_\ell = (V(\mathcal{G}), \cup_{i=1}^{\ell} E_i)$, for $\ell = 0, 1, \ldots, L$

$$\emptyset = G_0 \subseteq G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L = G$$

Hierarchical Spanner. a hierarchical subgraph $\mathcal{H}$ is a *hierarchical* (α, β)*-spanner* of $\mathcal{G}$ w.r.t. $P \subseteq V(\mathcal{G}) \times V(\mathcal{G})$ if H_ℓ is an (α, β)-spanner of G_ℓ, for every layer $\ell = 1, \ldots, L$.

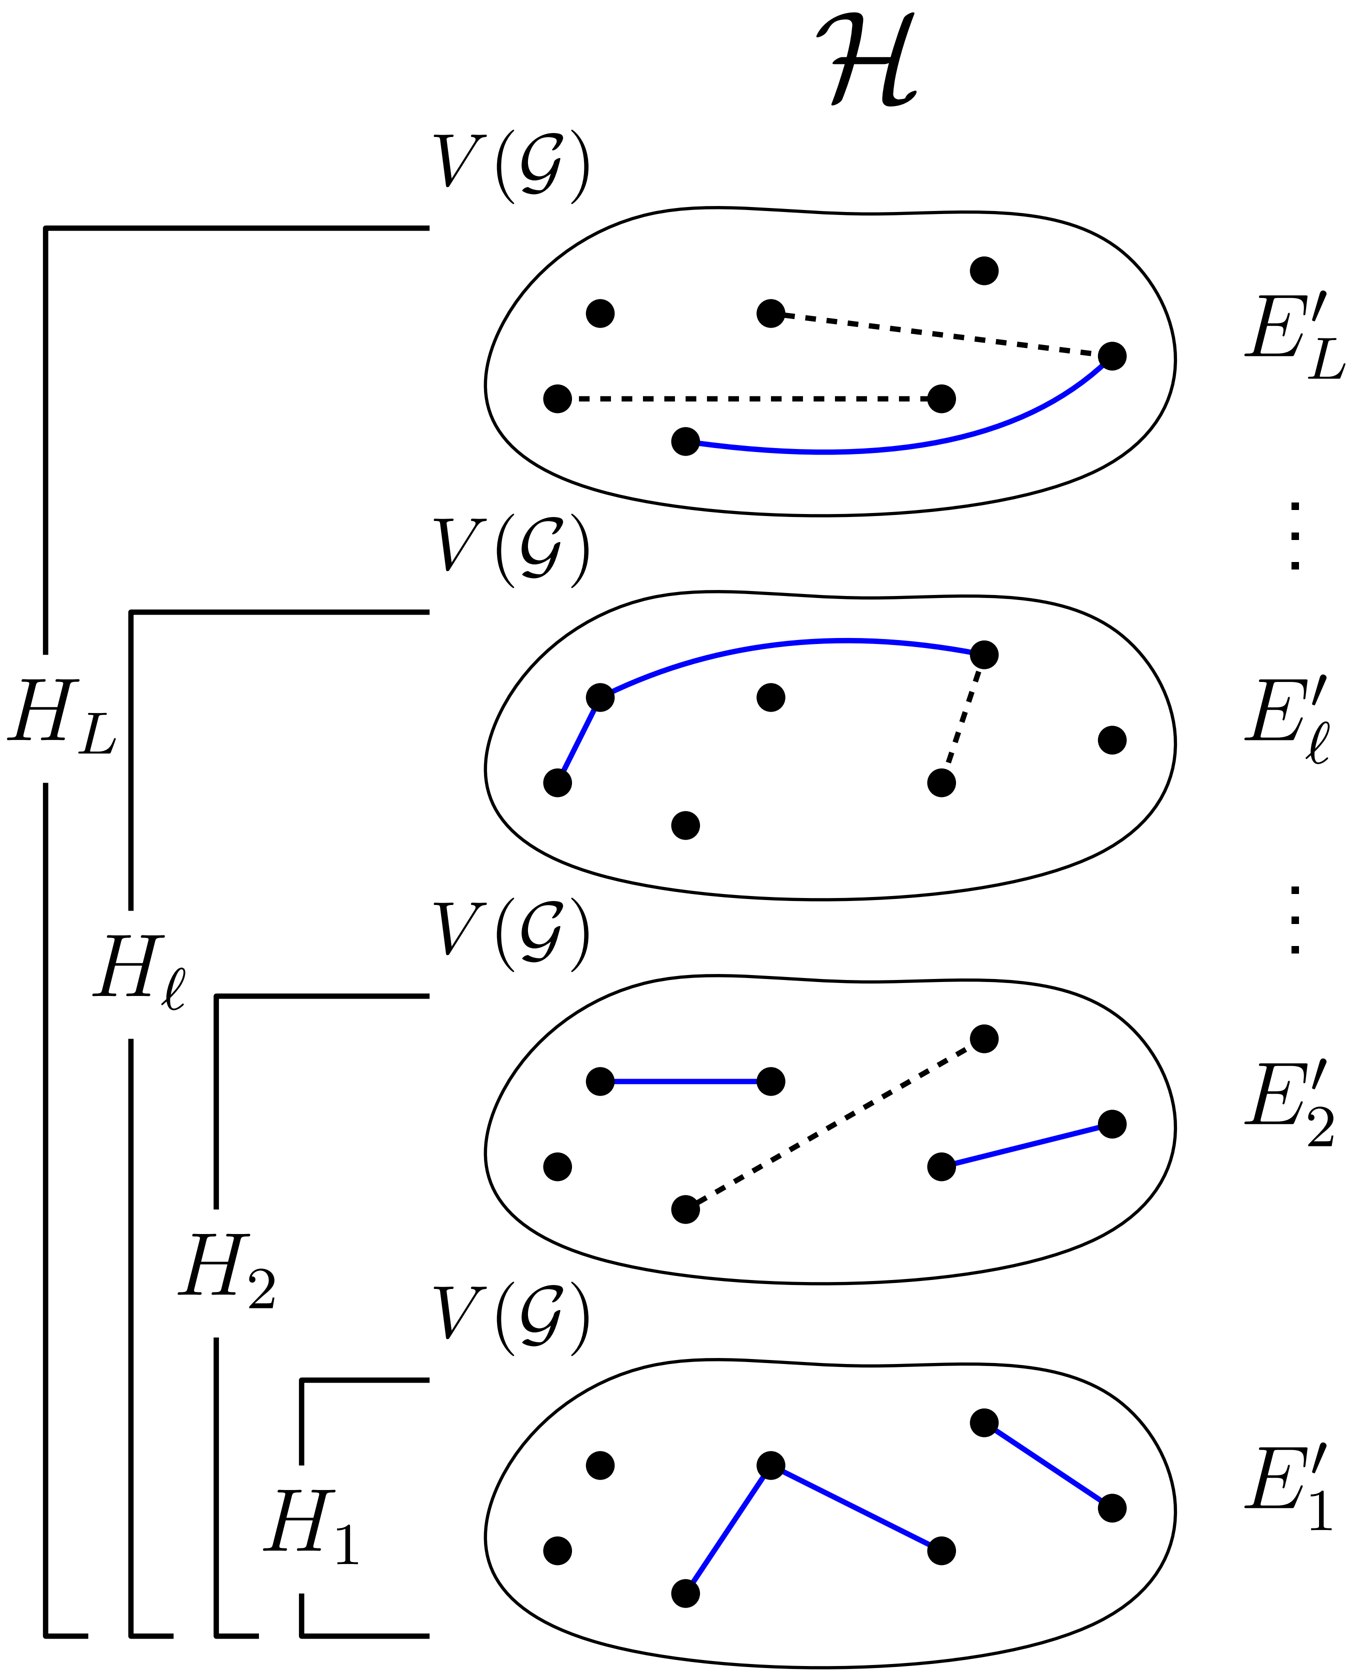

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Classical Graphs

Single-pair preserver

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Classical Graphs

Single-pair preserver $\longrightarrow O(n)$

just keep a shortest path

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Classical Graphs

Single-pair preserver $\longrightarrow O(n)$

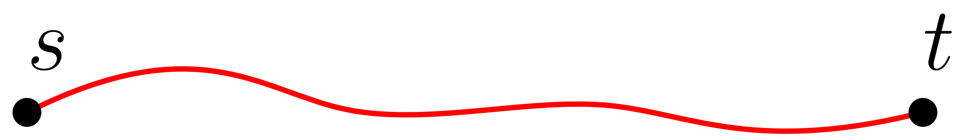

just keep a shortest path

Hierarchical Graphs

Single-pair preserver $\longrightarrow$?

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Classical Graphs

Single-pair preserver $\longrightarrow O(n)$

just keep a shortest path

Hierarchical Graphs

Single-pair preserver $\longrightarrow$?

spoiler: $\Theta(n\sqrt{n})$ tight bound

Q: How much harder does sparsification become when it has to preserve distances at every level of a hierarchy?

Classical Graphs

Single-pair preserver $\longrightarrow O(n)$

just keep a shortest path

Hierarchical Graphs

Single-pair preserver $\longrightarrow$?

spoiler: $\Theta(n\sqrt{n})$ tight bound

$\boxed{\text{Do classical size-stretch tradeoffs survive the hierarchy?}}$

Results roadmap

Results roadmap

Single-pair

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

Results roadmap

Single-pair

$$\beta\text{-additive} \longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$$

$$O(1)\text{-additive / preserver} \longrightarrow \Theta\left(n\sqrt{n}\right)$$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

	Classical	**Hierarchical**
	$(2k-1)$-stretch, $O(n^{1+1/k})$	$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

Results roadmap

Single-pair

$$\beta\text{-additive} \longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$$

$$O(1)\text{-additive / preserver} \longrightarrow \Theta\left(n\sqrt{n}\right)$$

All-pairs

Classical	Hierarchical
$(2k-1)$-stretch, $O(n^{1+1/k})$	$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)
2-additive, $O(n\sqrt{n})$	2-additive, ?

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

Classical	**Hierarchical**
$(2k-1)$-stretch, $O(n^{1+1/k})$	$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)
2-additive, $O(n\sqrt{n})$	2-additive, ? (unlikely $O(n\sqrt{n})$ size)

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta(n\sqrt{n})$

All-pairs

Classical	Hierarchical
$(2k-1)$-stretch, $O(n^{1+1/k})$	$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)
2-additive, $O(n\sqrt{n})$	2-additive, ? (unlikely $O(n\sqrt{n})$ size)
	(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

Classical

$(2k-1)$-stretch, $O(n^{1+1/k})$

2-additive, $O(n\sqrt{n})$

asymptotically similar for small constant ε

Hierarchical

$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

2-additive, ? (unlikely $O(n\sqrt{n})$ size)

(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

Classical

$(2k-1)$-stretch, $O(n^{1+1/k})$

2-additive, $O(n\sqrt{n})$

asymptotically similar for small constant ε

Hierarchical

$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

2-additive, ? (unlikely $O(n\sqrt{n})$ size)

(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

4-additive, $\widetilde{O}(n^{5/3})$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

Classical

$(2k-1)$-stretch, $O(n^{1+1/k})$

2-additive, $O(n\sqrt{n})$

asymptotically similar for small constant ε

Hierarchical

$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

2-additive, ? (unlikely $O(n\sqrt{n})$ size)

(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

4-additive, $\widetilde{O}(n^{5/3})$

$s \times T$

preserver $\longrightarrow O(n\sqrt{n|T|}), \Omega(n|T|)$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

	Classical	**Hierarchical**
	$(2k-1)$-stretch, $O(n^{1+1/k})$	$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)
	2-additive, $O(n\sqrt{n})$	2-additive, ? (unlikely $O(n\sqrt{n})$ size)
	asymptotically similar for small constant ε	(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$
		4-additive, $\widetilde{O}(n^{5/3})$

$s \times T$

preserver $\longrightarrow O(n\sqrt{n|T|}), \Omega(n|T|)$ $\Longrightarrow \Theta(n^2)$ for single-source $s \times V$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta\left(n\sqrt{n}\right)$

All-pairs

Classical

$(2k-1)$-stretch, $O(n^{1+1/k})$

2-additive, $O(n\sqrt{n})$

asymptotically similar for small constant ε

Hierarchical

$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

2-additive, ? (unlikely $O(n\sqrt{n})$ size)

(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

4-additive, $\widetilde{O}(n^{5/3})$

$s \times T$

preserver $\longrightarrow O(n\sqrt{n|T|}), \Omega(n|T|)$ $\implies \Theta(n^2)$ for single-source $s \times V$

$(1+\varepsilon)$-stretch $\longrightarrow \widetilde{O}(n/\varepsilon)$

Results roadmap

Single-pair

β-additive $\longrightarrow \Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$

$O(1)$-additive / preserver $\longrightarrow \Theta(n\sqrt{n})$

All-pairs

Classical

$(2k-1)$-stretch, $O(n^{1+1/k})$

2-additive, $O(n\sqrt{n})$

asymptotically similar for small constant ε

Hierarchical

$(2k-1)$-stretch, $O(n^{1+1/k})$ (no asymptotic loss)

2-additive, ? (unlikely $O(n\sqrt{n})$ size)

(nearly additive) $(1+\varepsilon, 2)$, $\widetilde{O}(n\sqrt{\frac{n}{\varepsilon}})$

4-additive, $\widetilde{O}(n^{5/3})$

$s \times T$

preserver $\longrightarrow O(n\sqrt{n|T|}), \Omega(n|T|)$ $\implies \Theta(n^2)$ for single-source $s \times V$

$(1+\varepsilon)$-stretch $\longrightarrow \widetilde{O}(n/\varepsilon)$

other upperbounds for subsetwise $(S \times S)$ and pairwise

A simple greedy construction: only keep critical shortest paths

A simple greedy construction: only keep critical shortest paths

Fix any pair (s, t), and let $d_\ell = d_{G_\ell}(s, t)$, for every hierarchy $\ell = 1, \ldots, L$.

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1, \ldots, L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

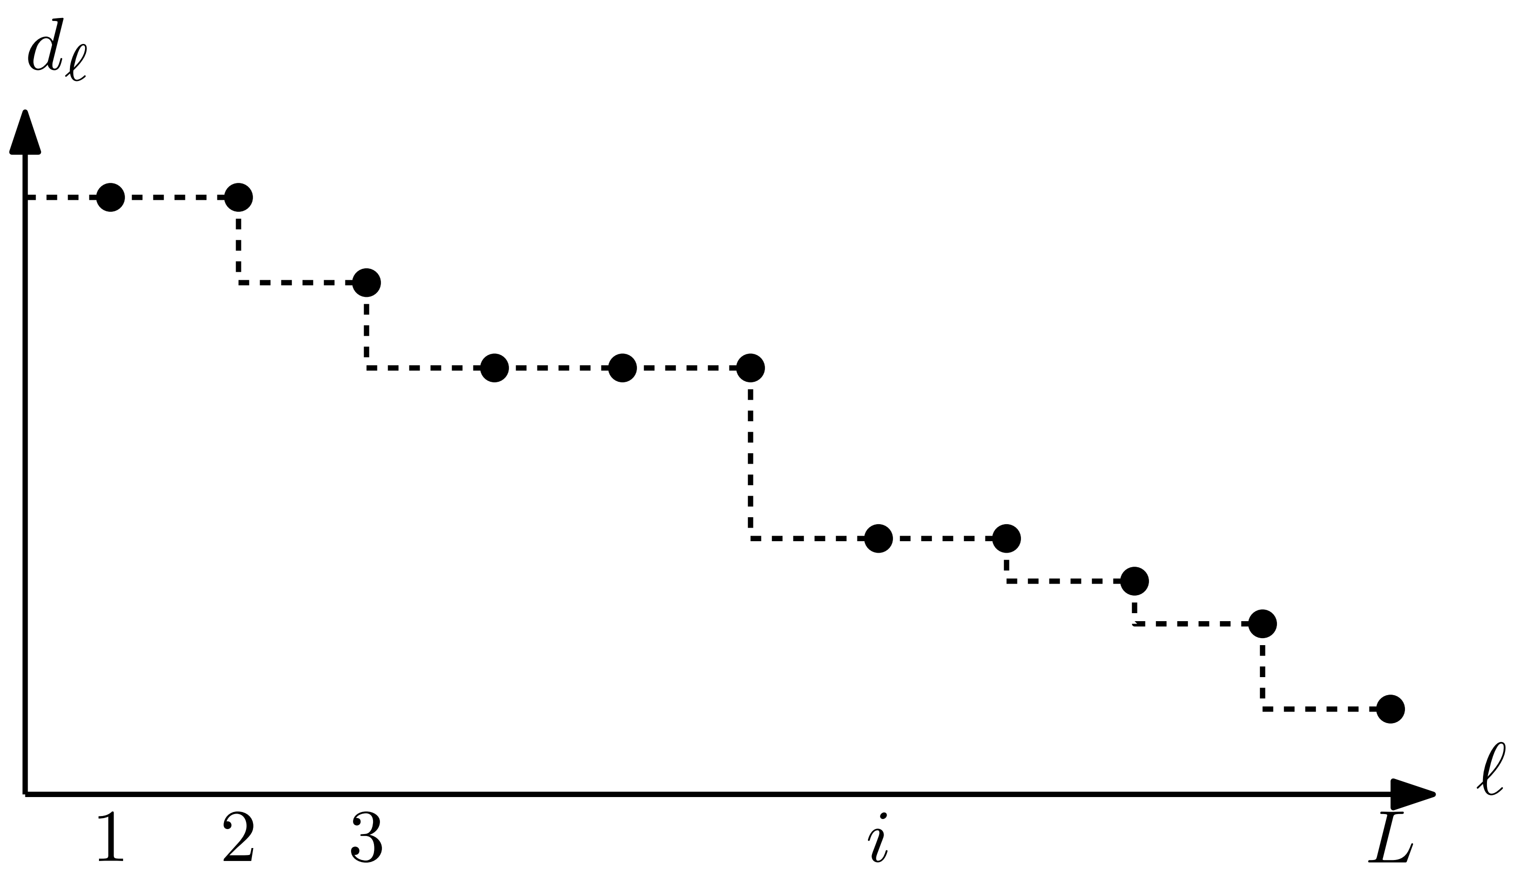

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

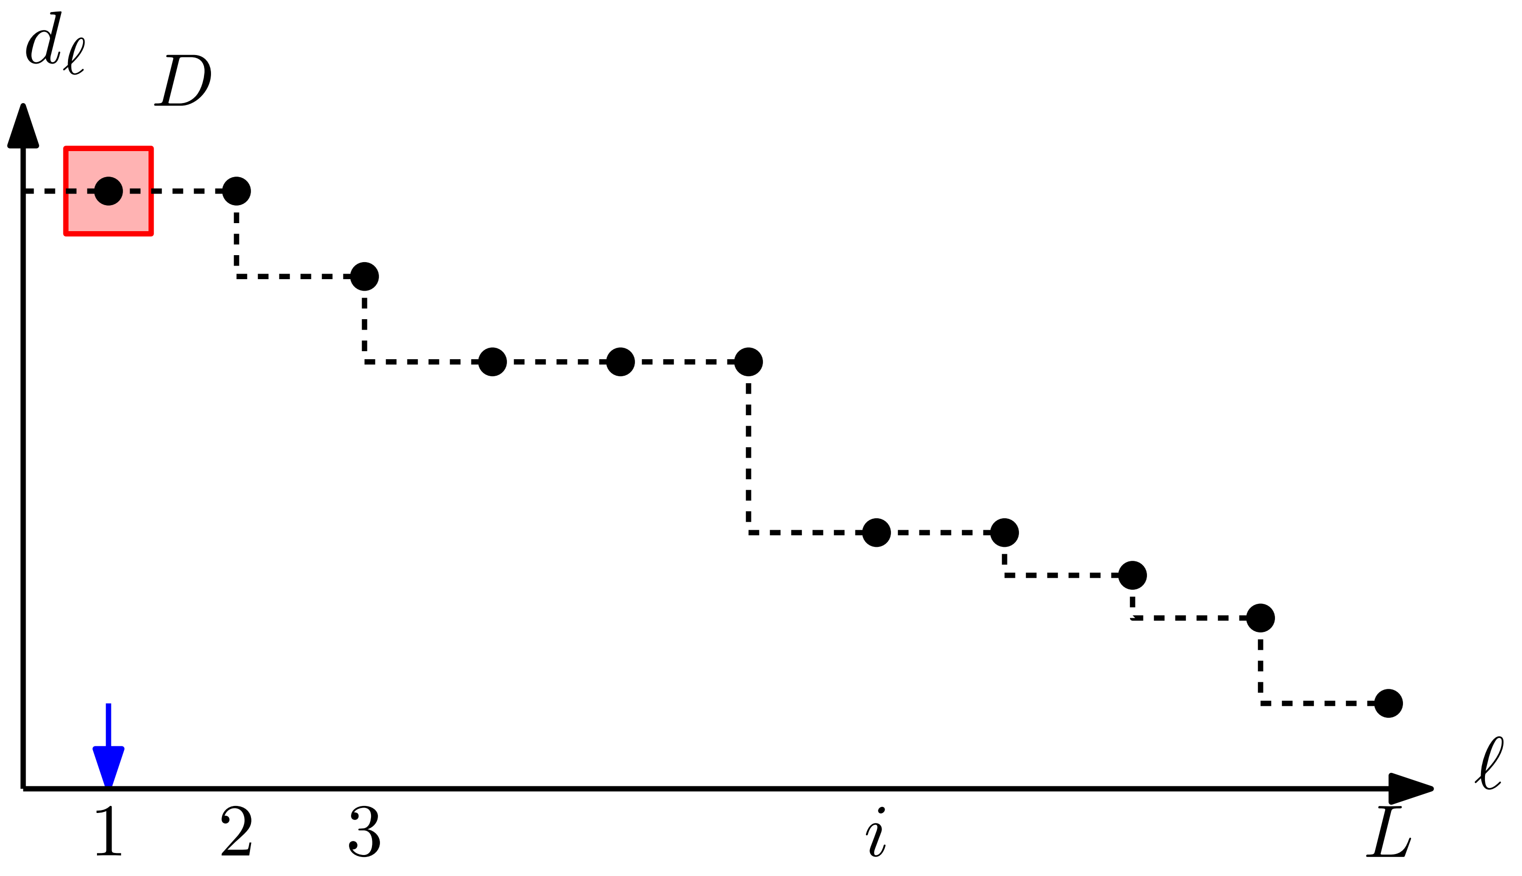

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

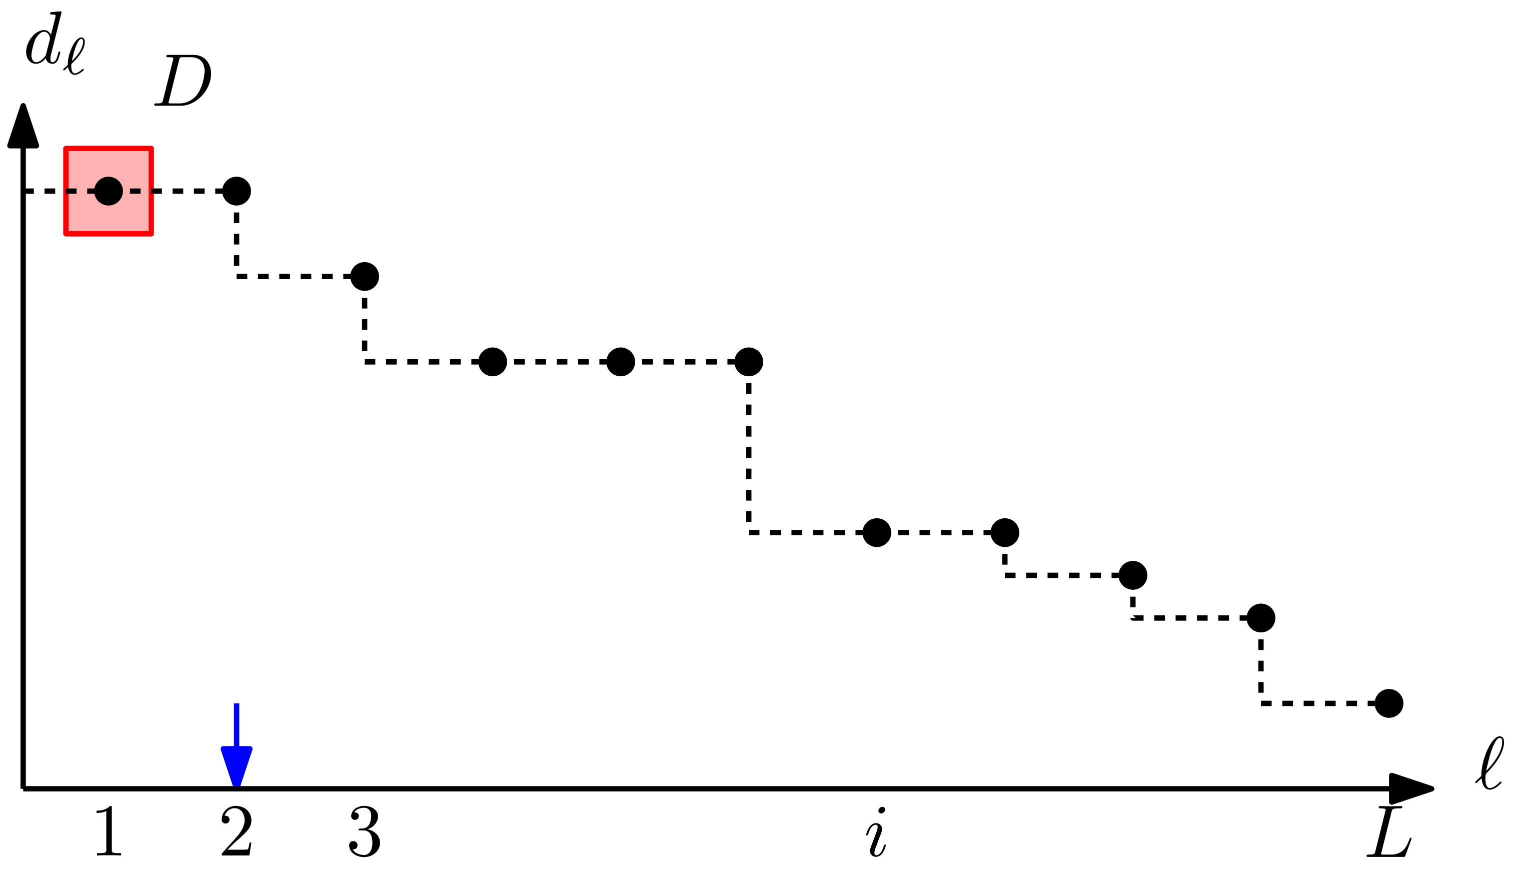

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

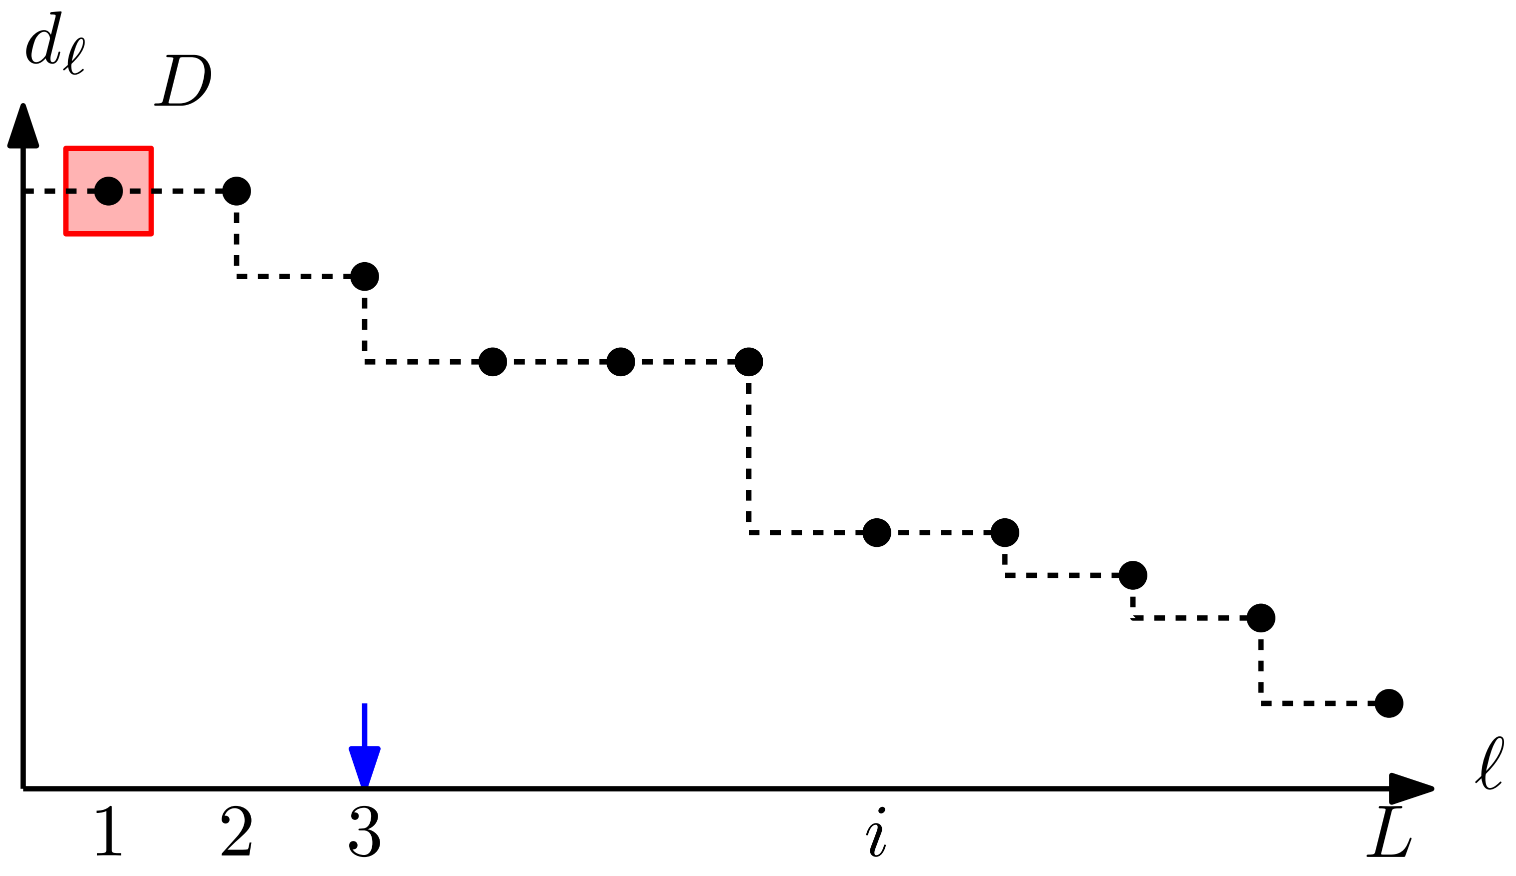

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1, \dots, L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1, \dots, L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

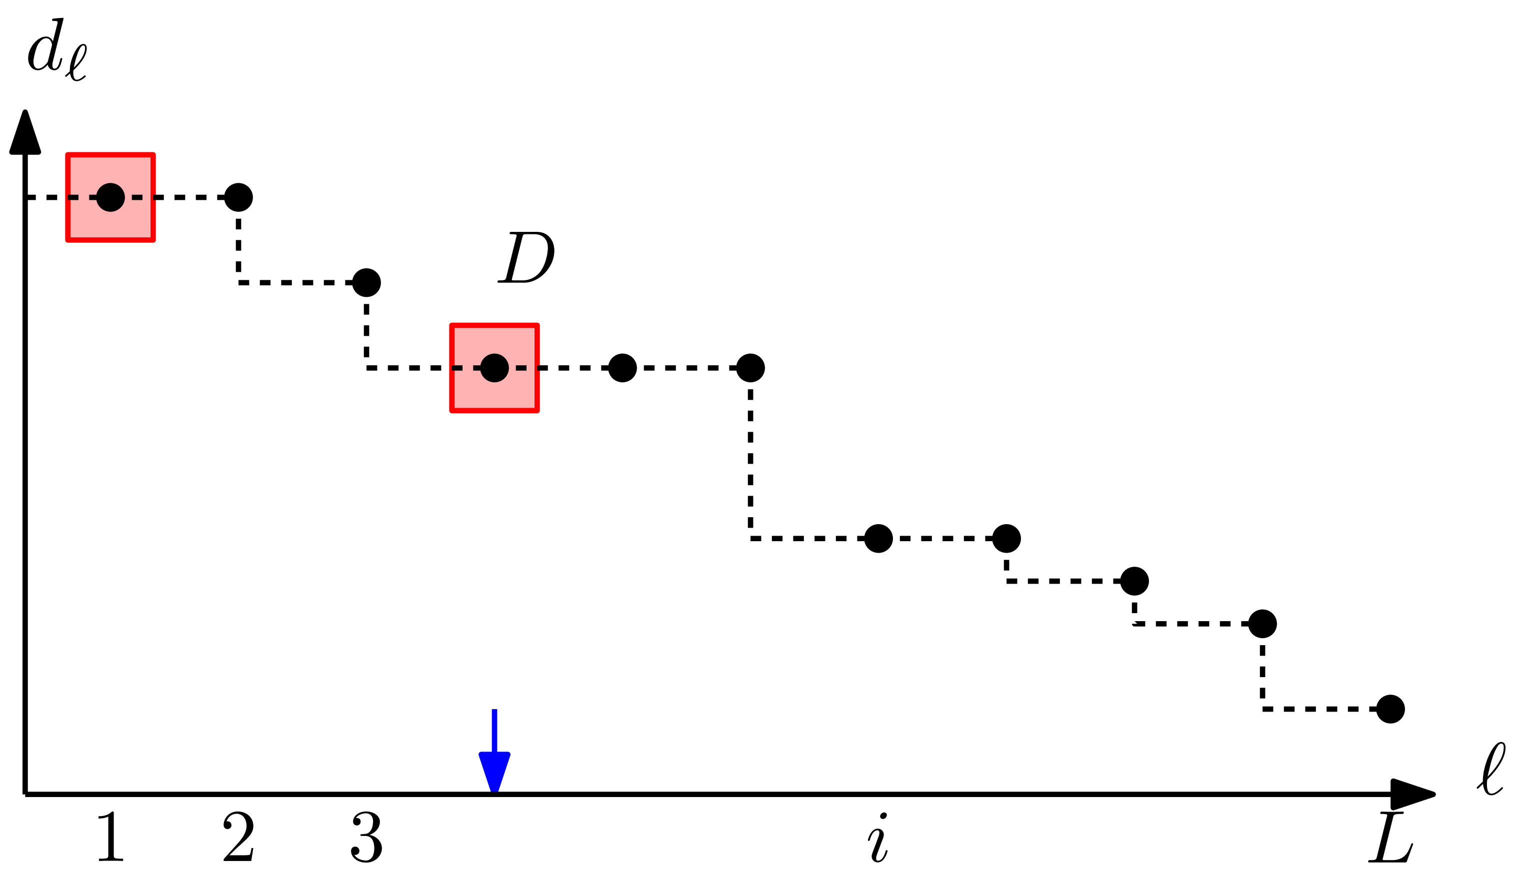

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

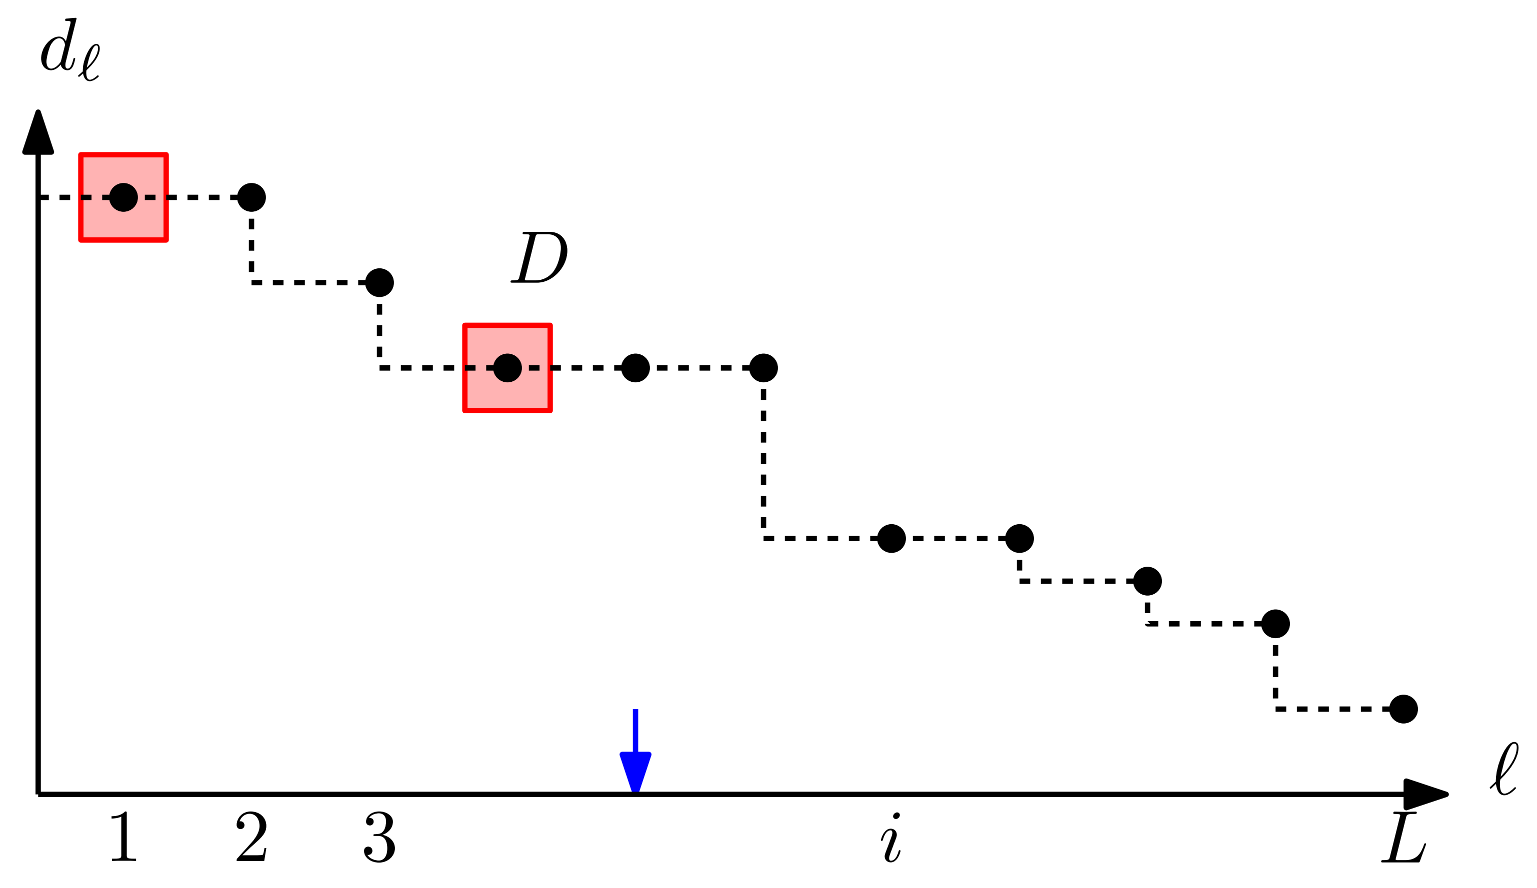

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

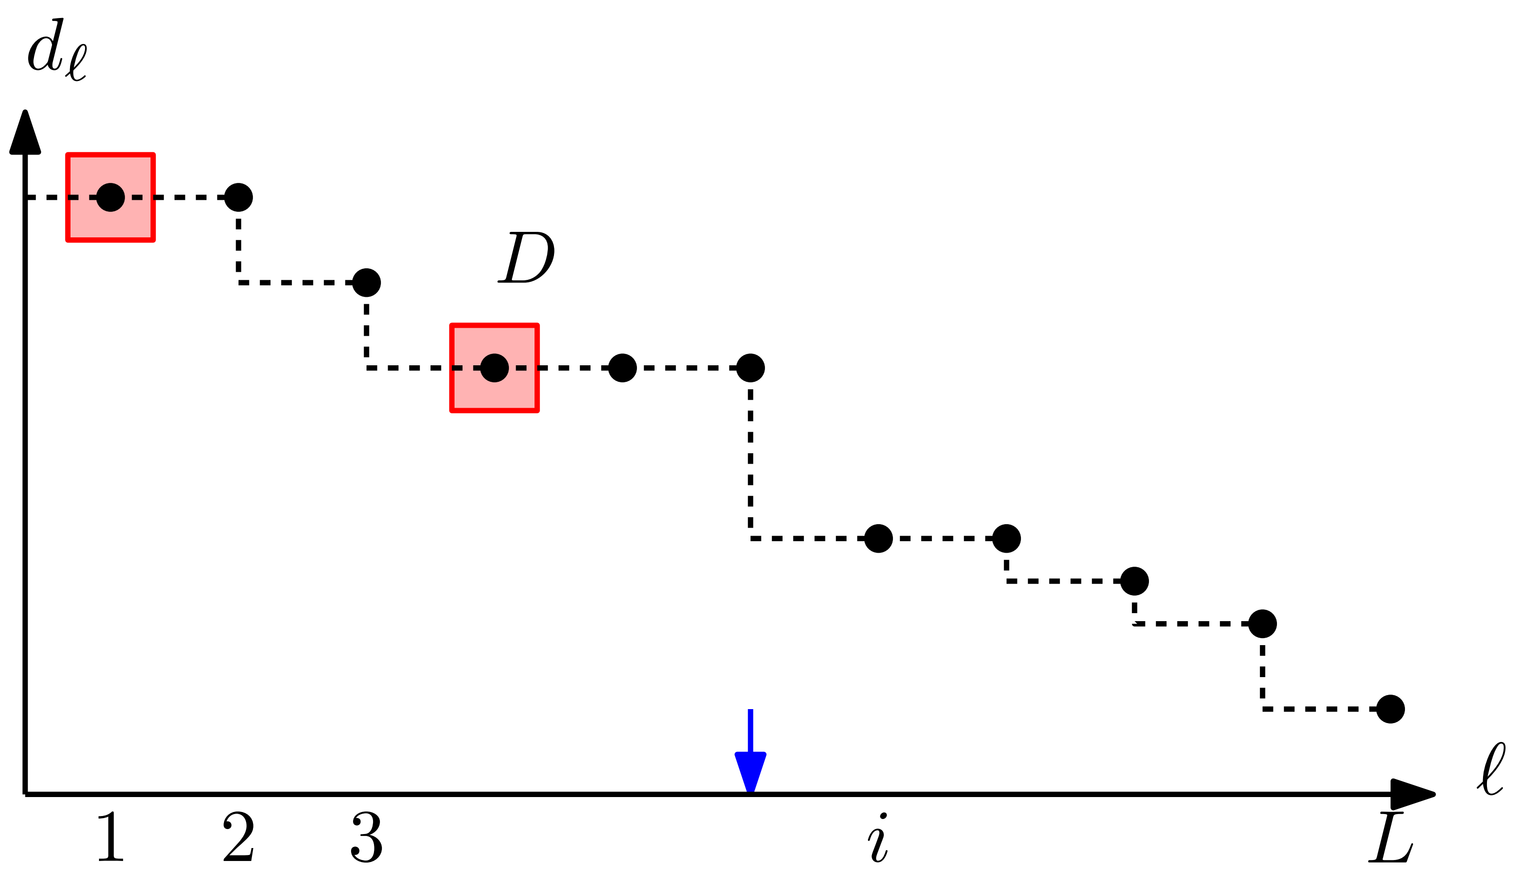

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

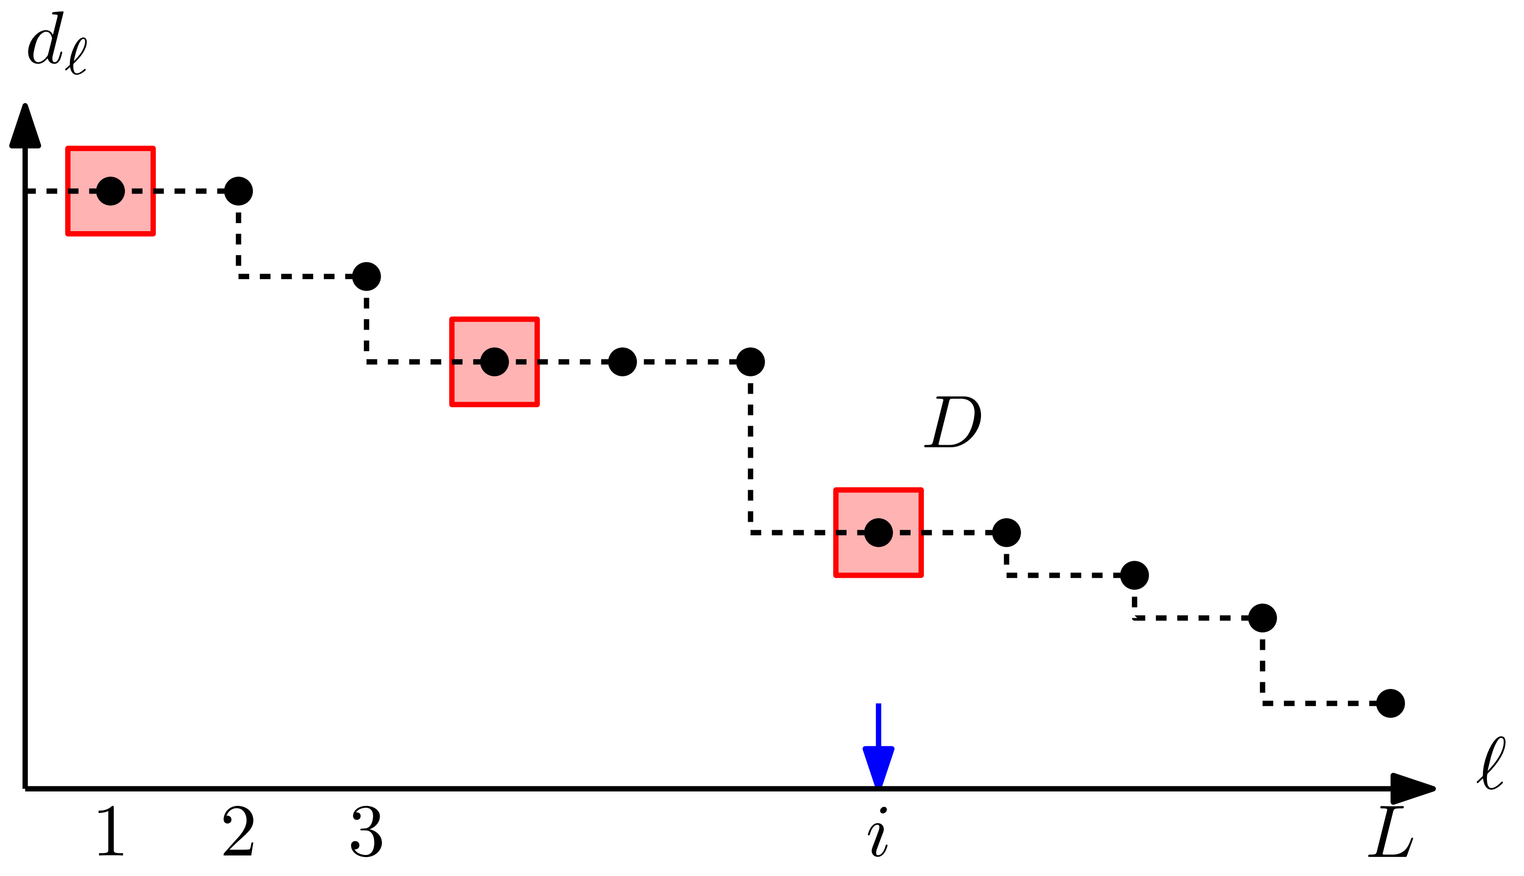

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

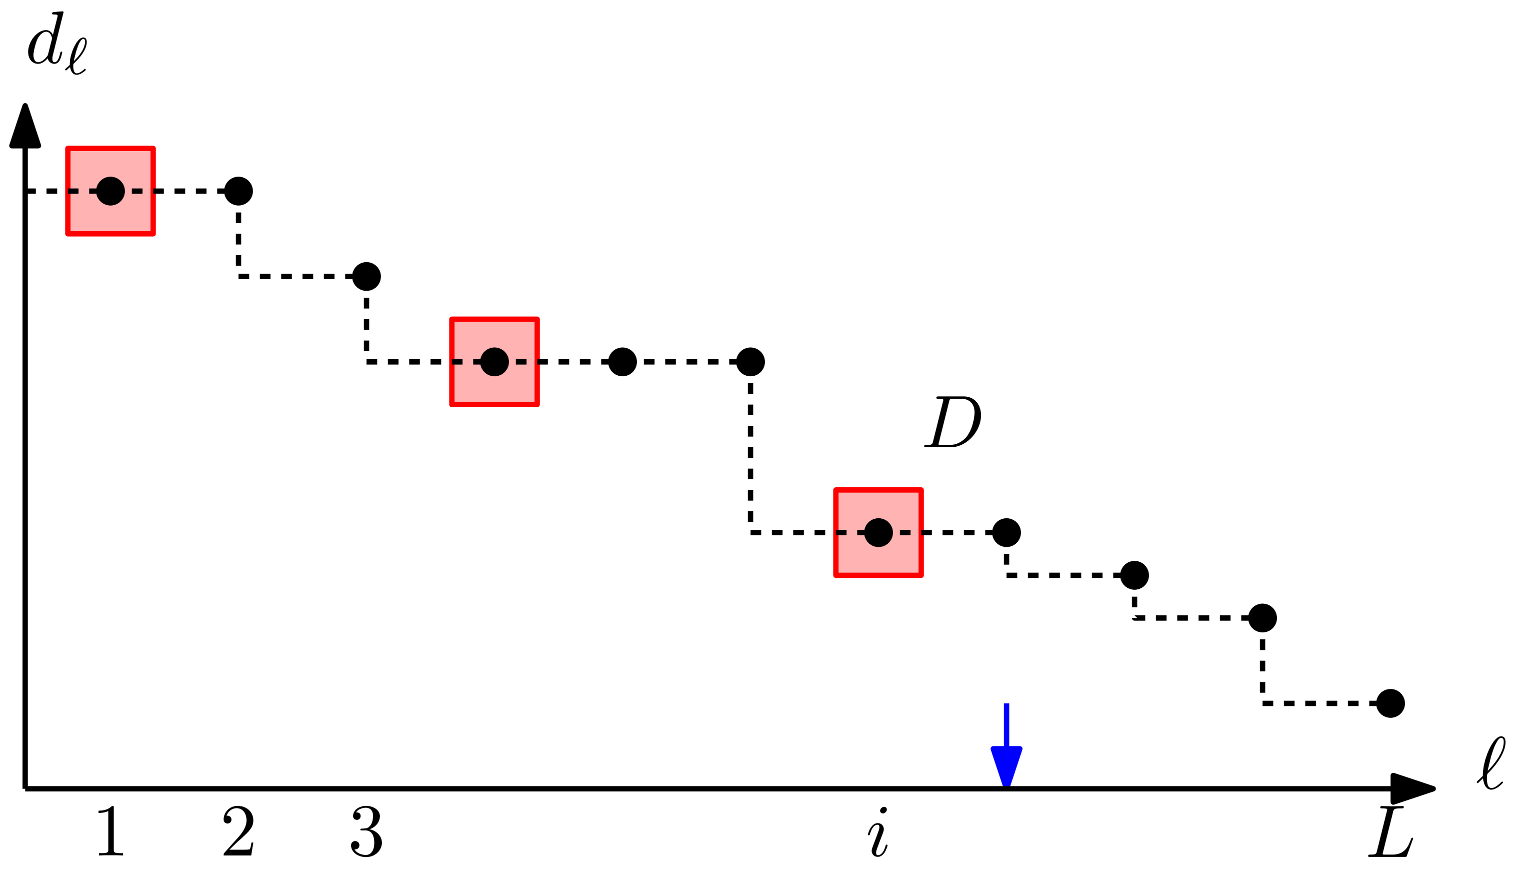

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

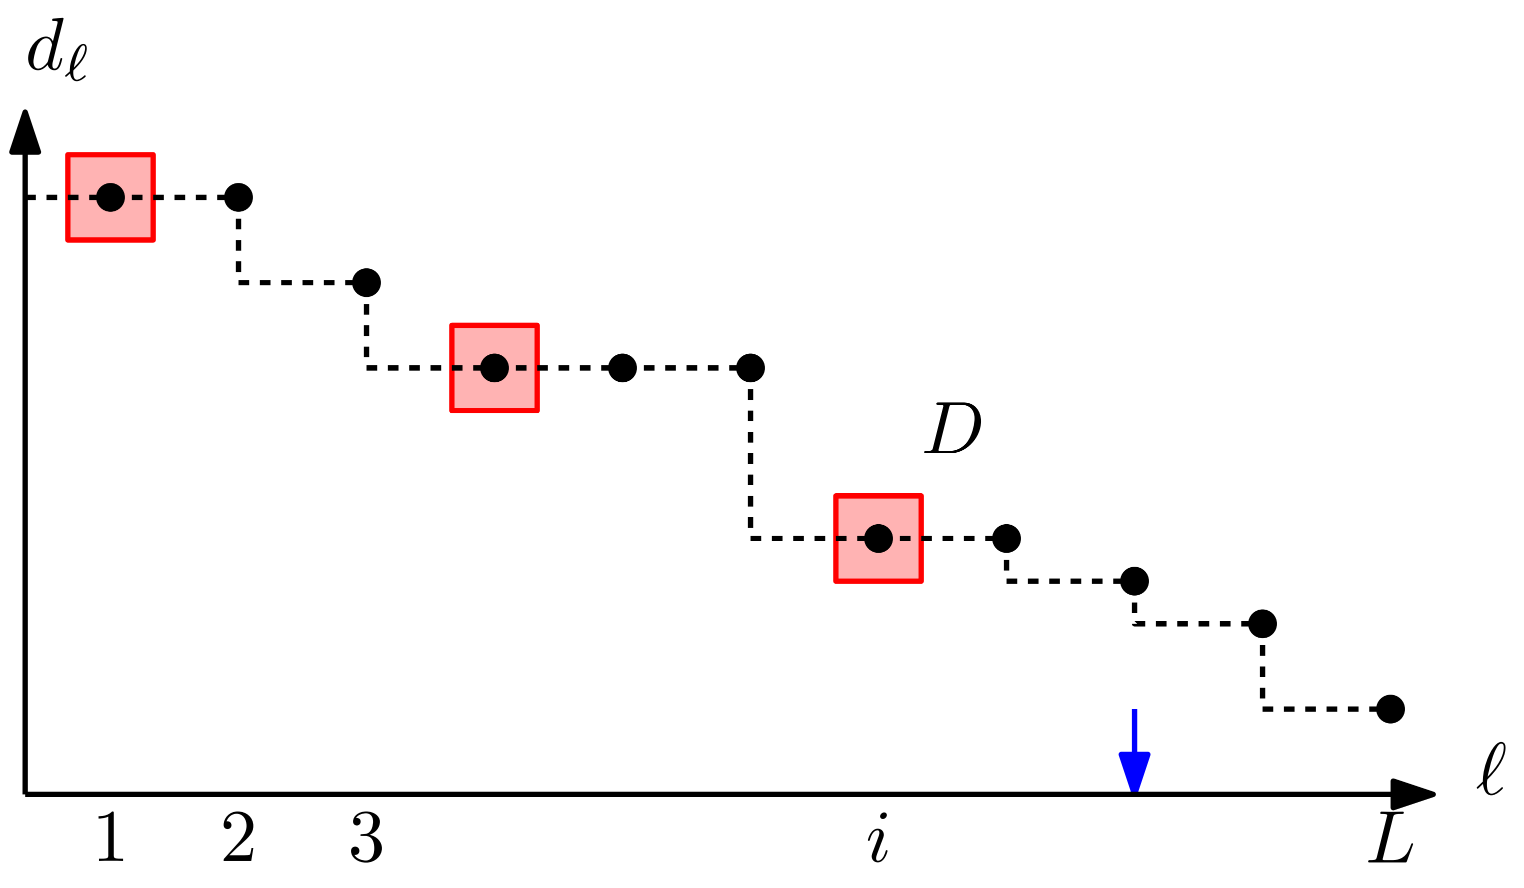

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1, \dots, L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1, \dots, L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

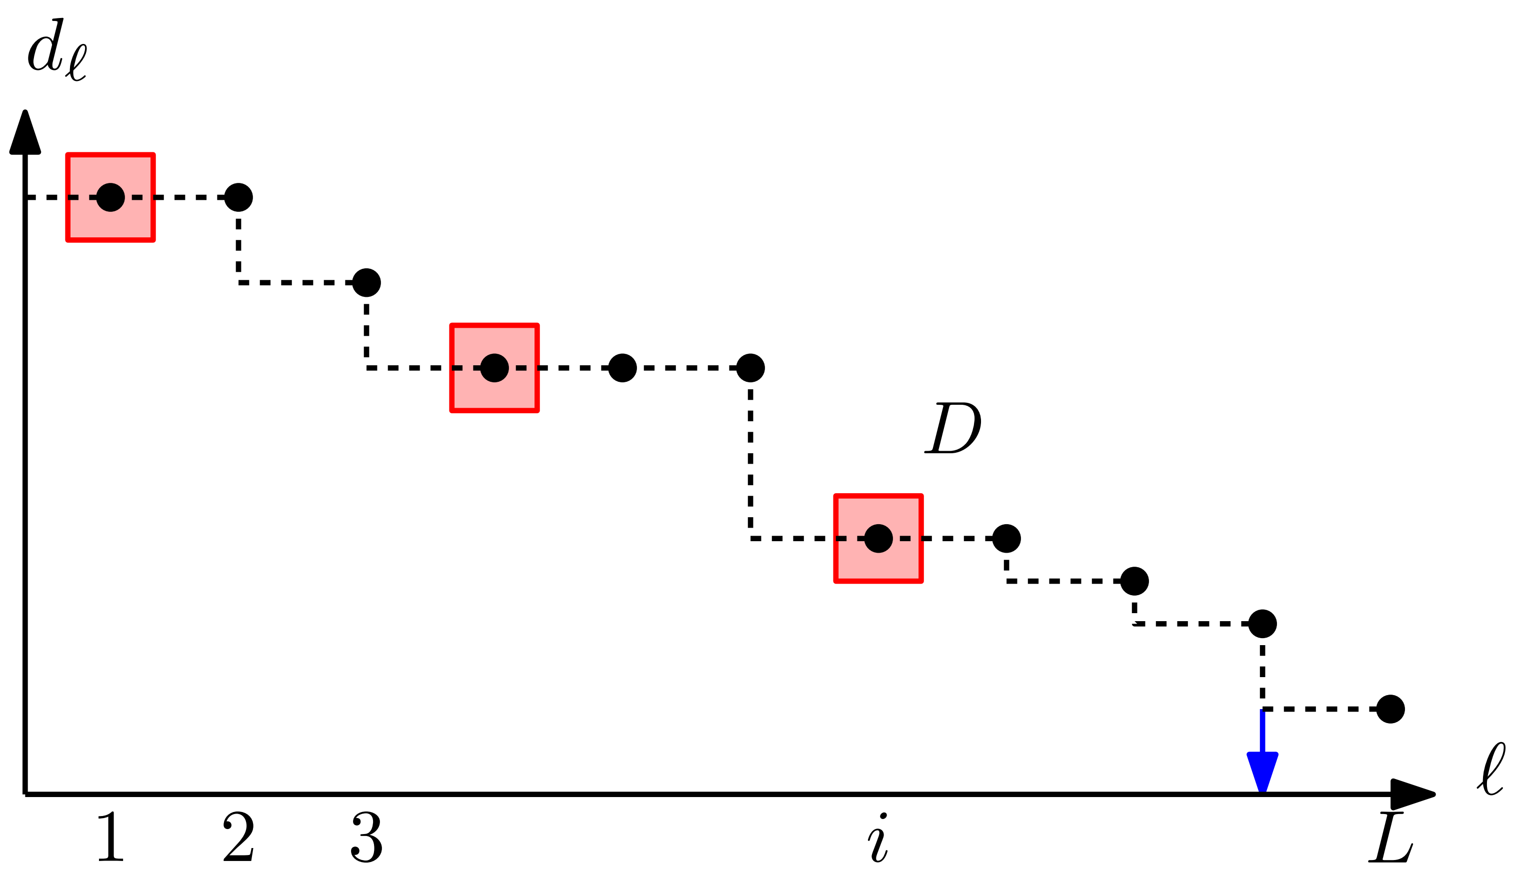

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

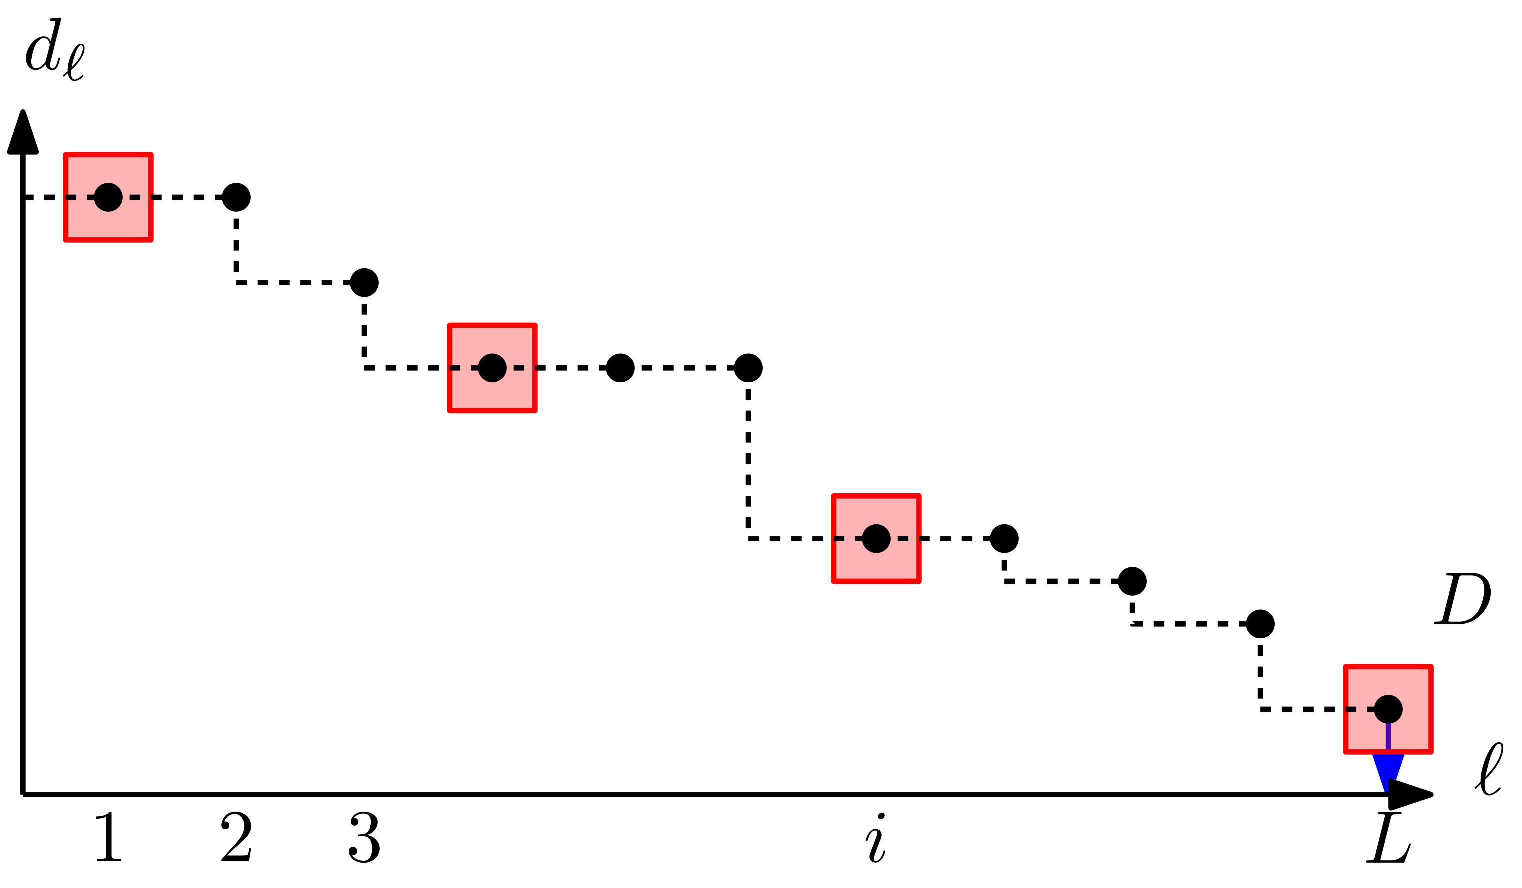

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

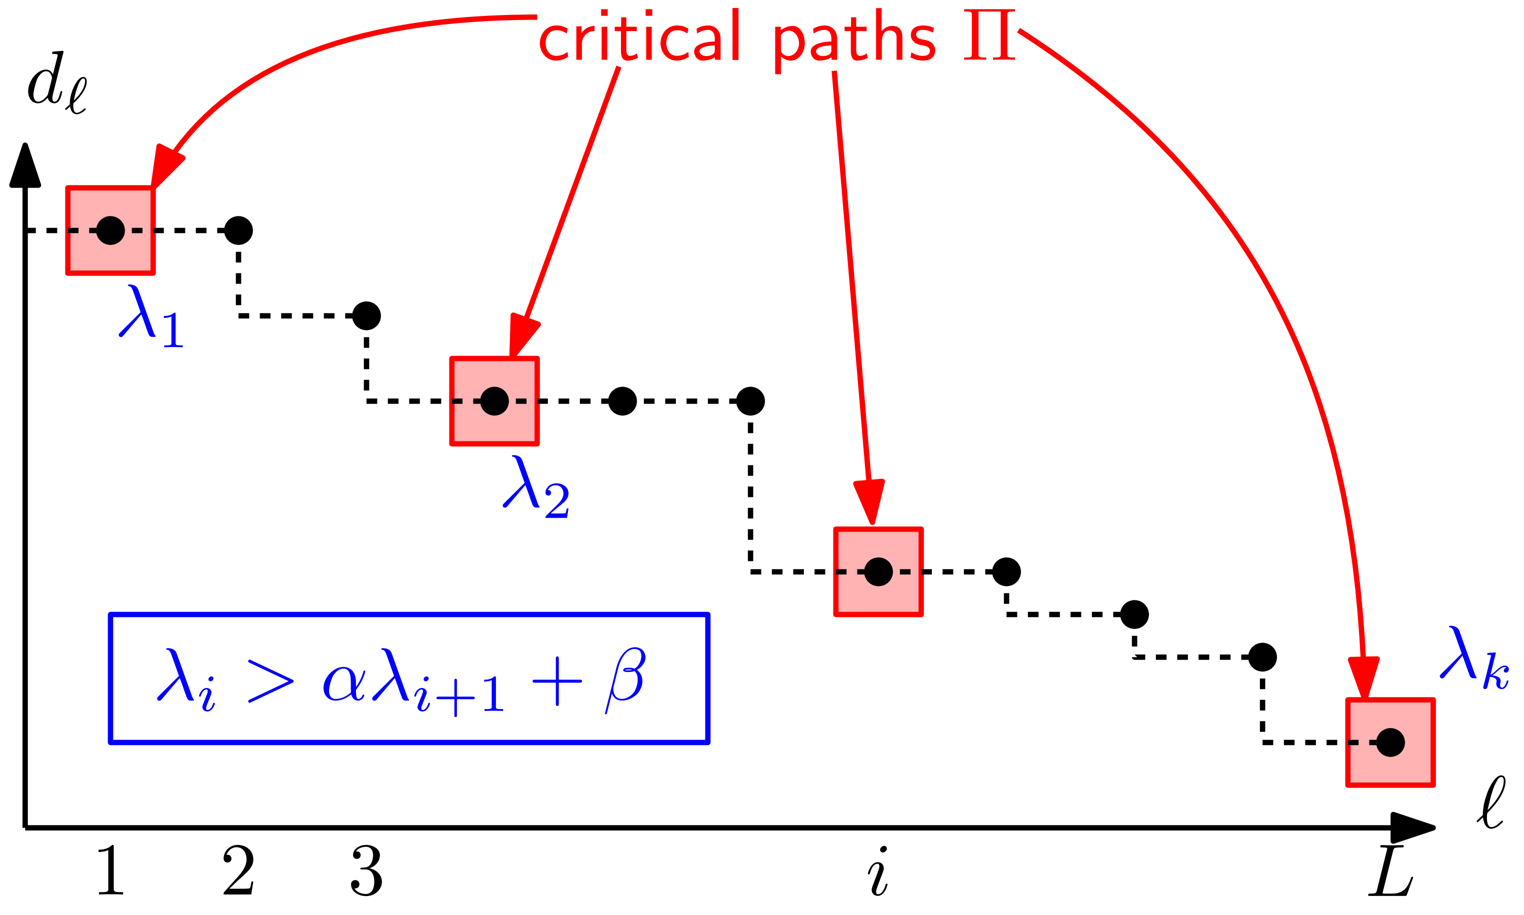

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\ldots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\ldots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

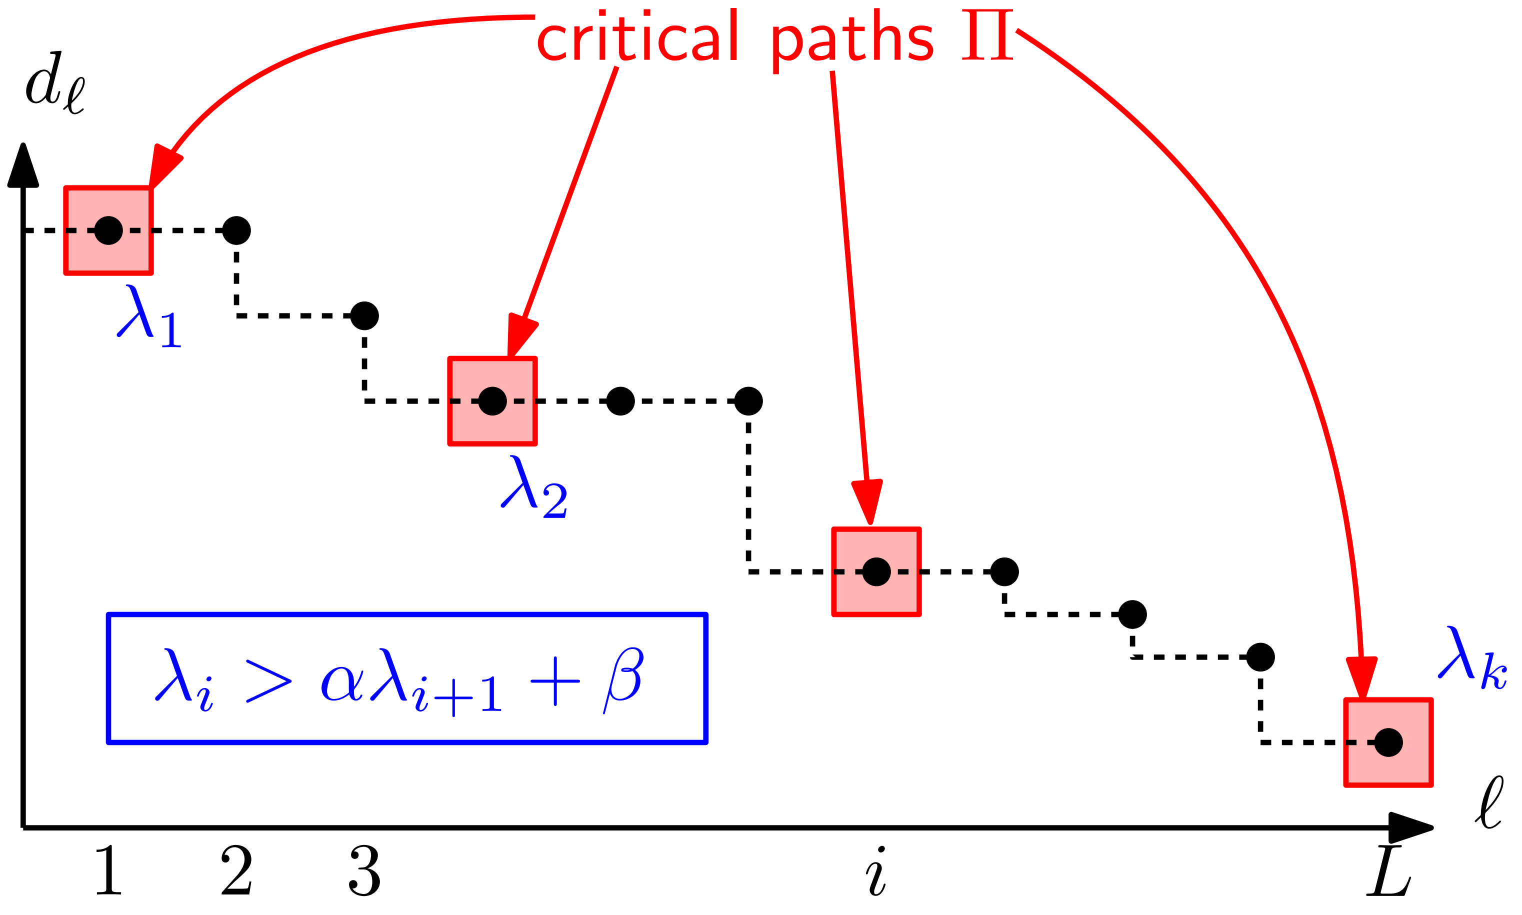

Pure additive case ($\alpha = 1, \beta > 0$):

A possible sequence of $\lambda_1,\ldots,\lambda_k$ could be

$$\begin{array}{ccccc} \lambda_1 & \lambda_2 & \ldots & \lambda_{k-1} & \lambda_k \\ n & n-\beta & \ldots & \beta+1 & 1 \end{array} \implies |\Pi| = O\left(\frac{n}{\beta+1}\right)$$

A simple greedy construction: only keep critical shortest paths

Fix any pair (s,t), and let $d_\ell = d_{G_\ell}(s,t)$, for every hierarchy $\ell = 1,\dots,L$.

Monotoniticy. $G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L \implies d_1 \geq d_2 \geq \cdots \geq d_L$.

Greedy Alg: let D be the lenght of the last path added to $\mathcal{H}$. For every hierarchy $\ell = 1,\dots,L$, store the (s,t)-shortest path π_ℓ (with lenght d_ℓ) iff $D > \alpha d_\ell + \beta$.

Proof: if we don't add π_ℓ to $\mathcal{H} \implies d_{H_\ell}(s,t) = D \leq \alpha d_\ell + \beta$

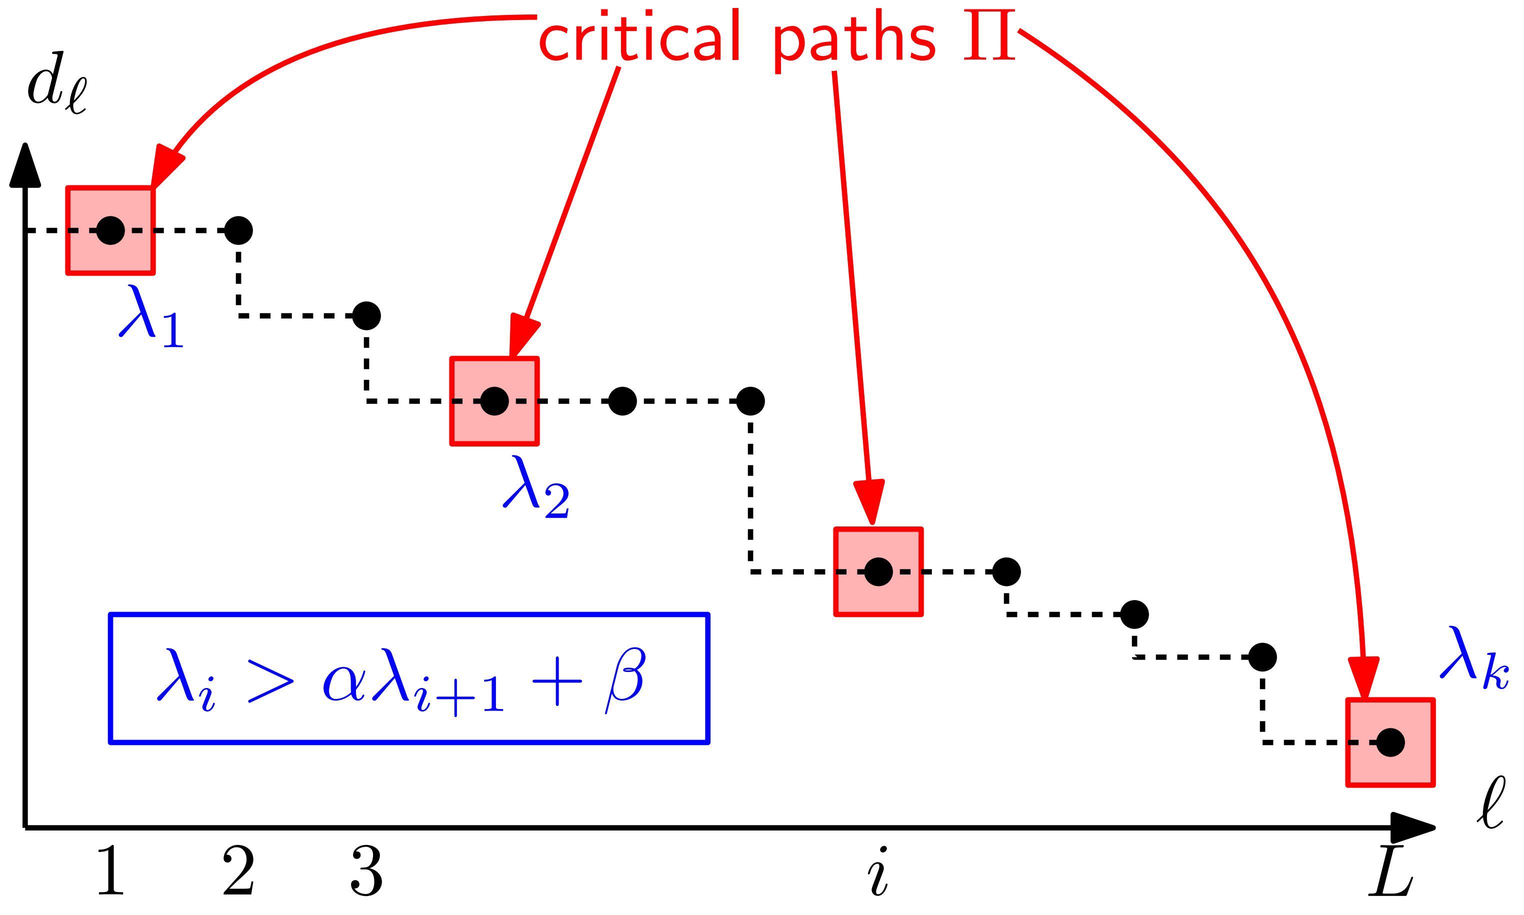

Pure additive case ($\alpha = 1, \beta > 0$):

A possible sequence of $\lambda_1,\dots,\lambda_k$ could be

$$\begin{array}{ccccc} \lambda_1 & \lambda_2 & \dots & \lambda_{k-1} & \lambda_k \\ n & n-\beta & \dots & \beta+1 & 1 \end{array} \implies |\Pi| = O\left(\tfrac{n}{\beta+1}\right)$$

Pure multiplicative case ($\alpha > 1, \beta = 0$):

A possible sequence of $\lambda_1,\dots,\lambda_k$ could be

$$\begin{array}{cccc} \lambda_1 & \lambda_2 & \lambda_3 & \dots \\ n & \frac{n}{\alpha} & \frac{n}{\alpha^2} & \dots \end{array} \implies |\Pi| = O(\log_\alpha n)$$

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

Trivial bound: $|E(\mathcal{H})| \leq n|\Pi|$ (with $|\Pi| \leq \min\{\frac{n}{\beta+1}, \log_\alpha n\}$).

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

Trivial bound: $|E(\mathcal{H})| \le n|\Pi|$ (with $|\Pi| \le \min\{\frac{n}{\beta+1}, \log_\alpha n\}$).

Idea: shortest paths cannot overlap arbitrarily badly.

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

Trivial bound: $|E(\mathcal{H})| \leq n|\Pi|$ (with $|\Pi| \leq \min\{\frac{n}{\beta+1}, \log_\alpha n\}$).

Idea: shortest paths cannot overlap arbitrarily badly.

Lemma 3. $|E(\mathcal{H})| = O(n\sqrt{|\Pi|})$.

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

Trivial bound: $|E(\mathcal{H})| \leq n|\Pi|$ (with $|\Pi| \leq \min\{\frac{n}{\beta+1}, \log_\alpha n\}$).

Idea: shortest paths cannot overlap arbitrarily badly.

Lemma 3. $|E(\mathcal{H})| = O(n\sqrt{|\Pi|})$.

i) single-pair preserver $\alpha = 1, \beta = 0 \implies O(n\sqrt{n})$

The union of critical shortest paths is sparse

Greedy alg.: for every hierarchy ℓ store all critical pats π_ℓ to Π

Trivial bound: $|E(\mathcal{H})| \leq n|\Pi|$ (with $|\Pi| \leq \min\{\frac{n}{\beta+1}, \log_\alpha n\}$).

Idea: shortest paths cannot overlap arbitrarily badly.

Lemma 3. $|E(\mathcal{H})| = O(n\sqrt{|\Pi|})$.

i) single-pair preserver $\alpha = 1, \beta = 0 \implies O(n\sqrt{n})$

ii) can be generalized to **pairwise**, with size $O(n\sqrt{|P||\Pi|})$

Even though a single pair may have $\Theta(n)$ critical shortest paths, their union always has only $O(n\sqrt{n})$ edges.
Can we do better?

Even though a single pair may have $\Theta(n)$ critical shortest paths, their union always has only $O(n\sqrt{n})$ edges.
Can we do better?

No, $\Omega(n\sqrt{n})$ edges can be necessary.
At many different levels, a different long path becomes unavoiable.

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

$\mathcal{X}$

$\ell = 1$

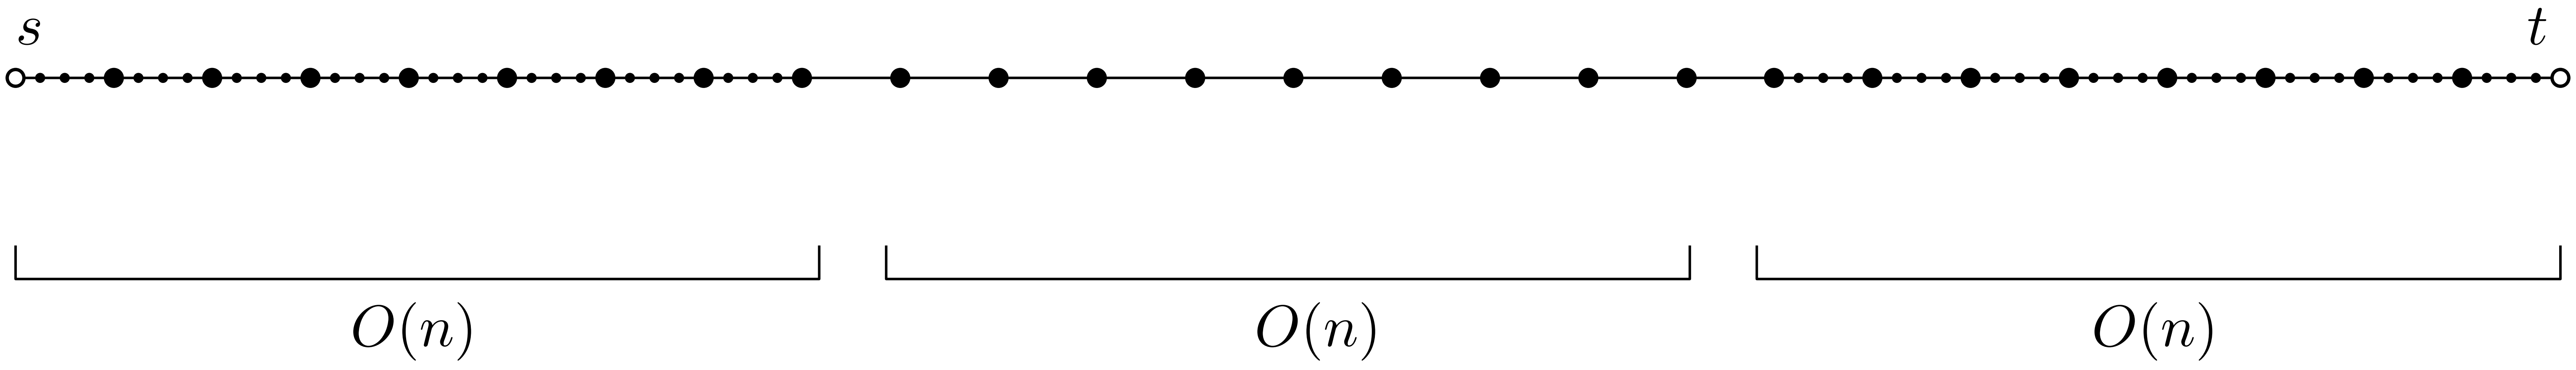

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

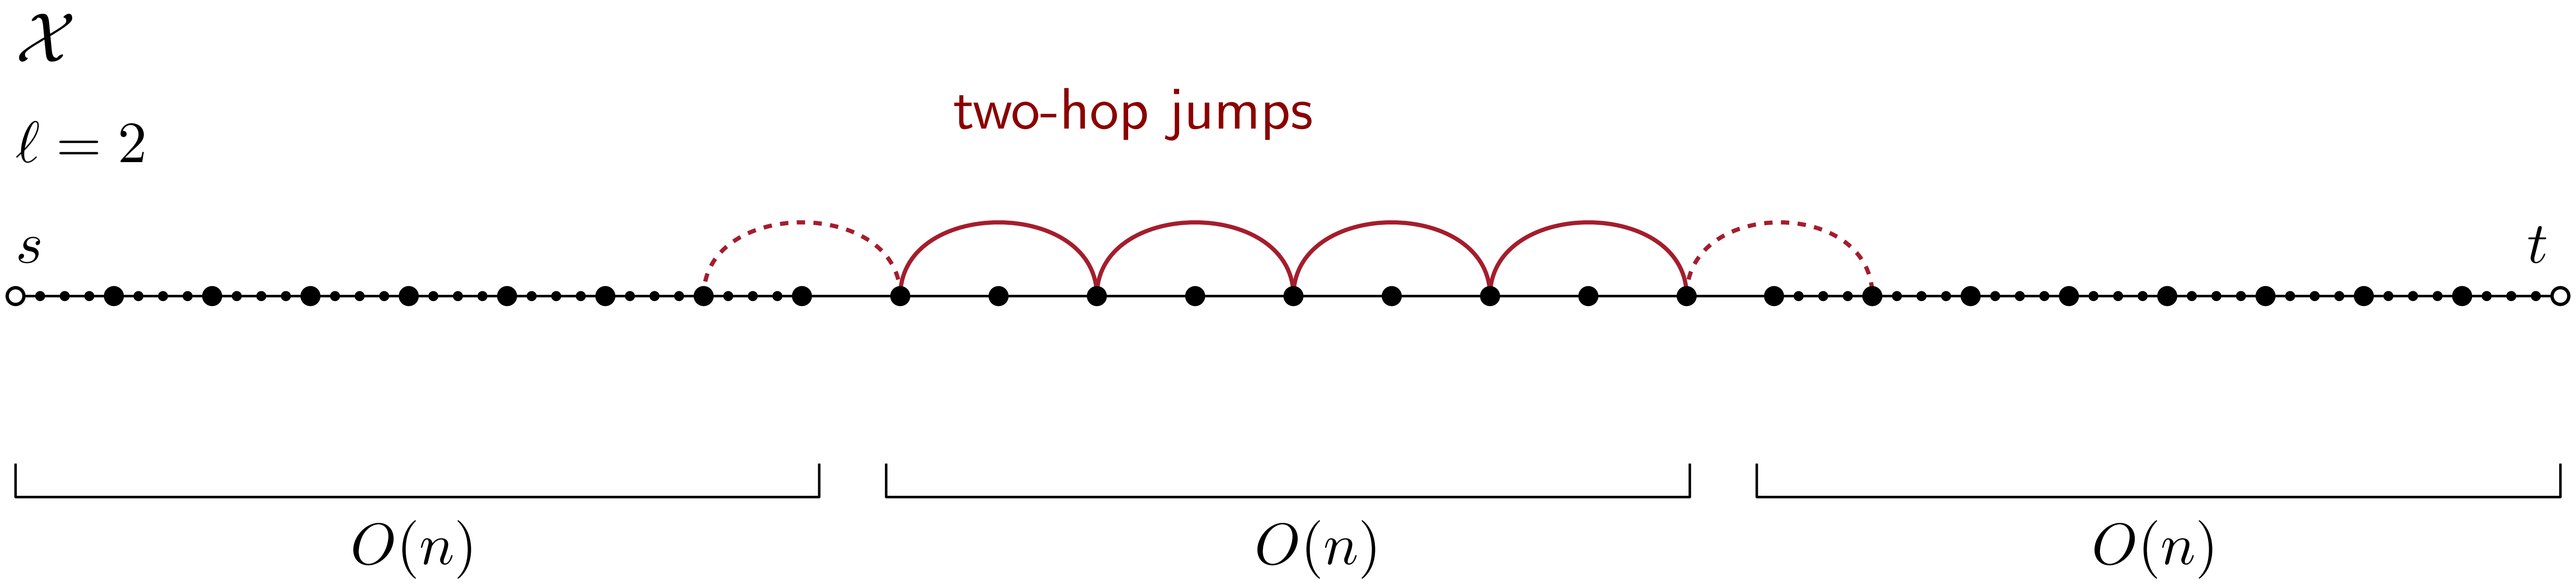

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

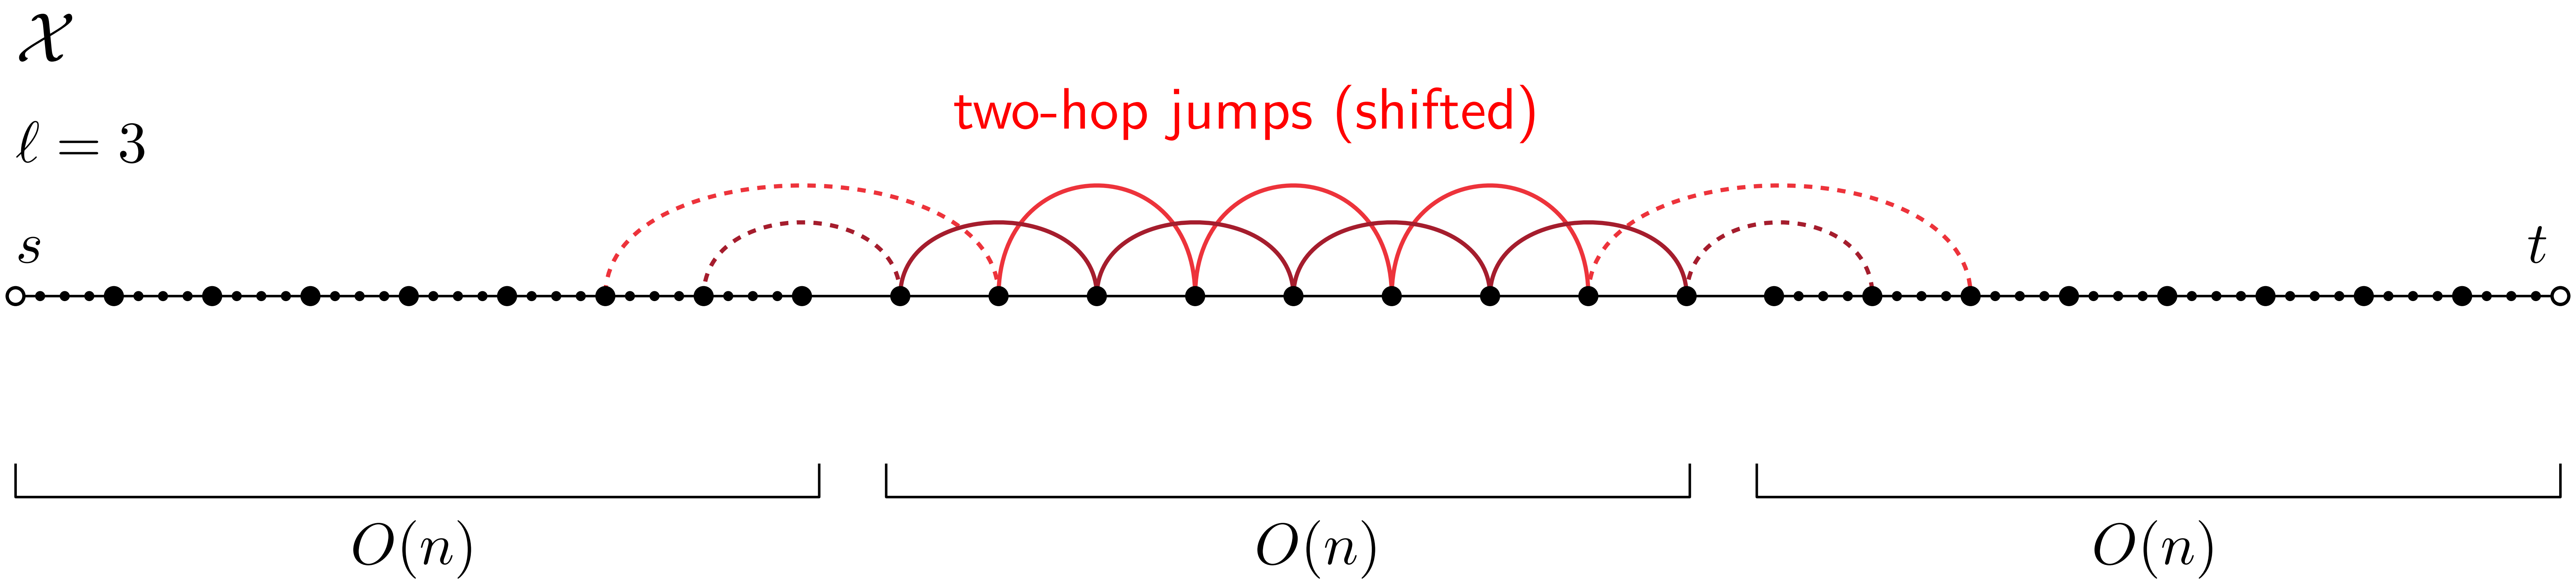

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

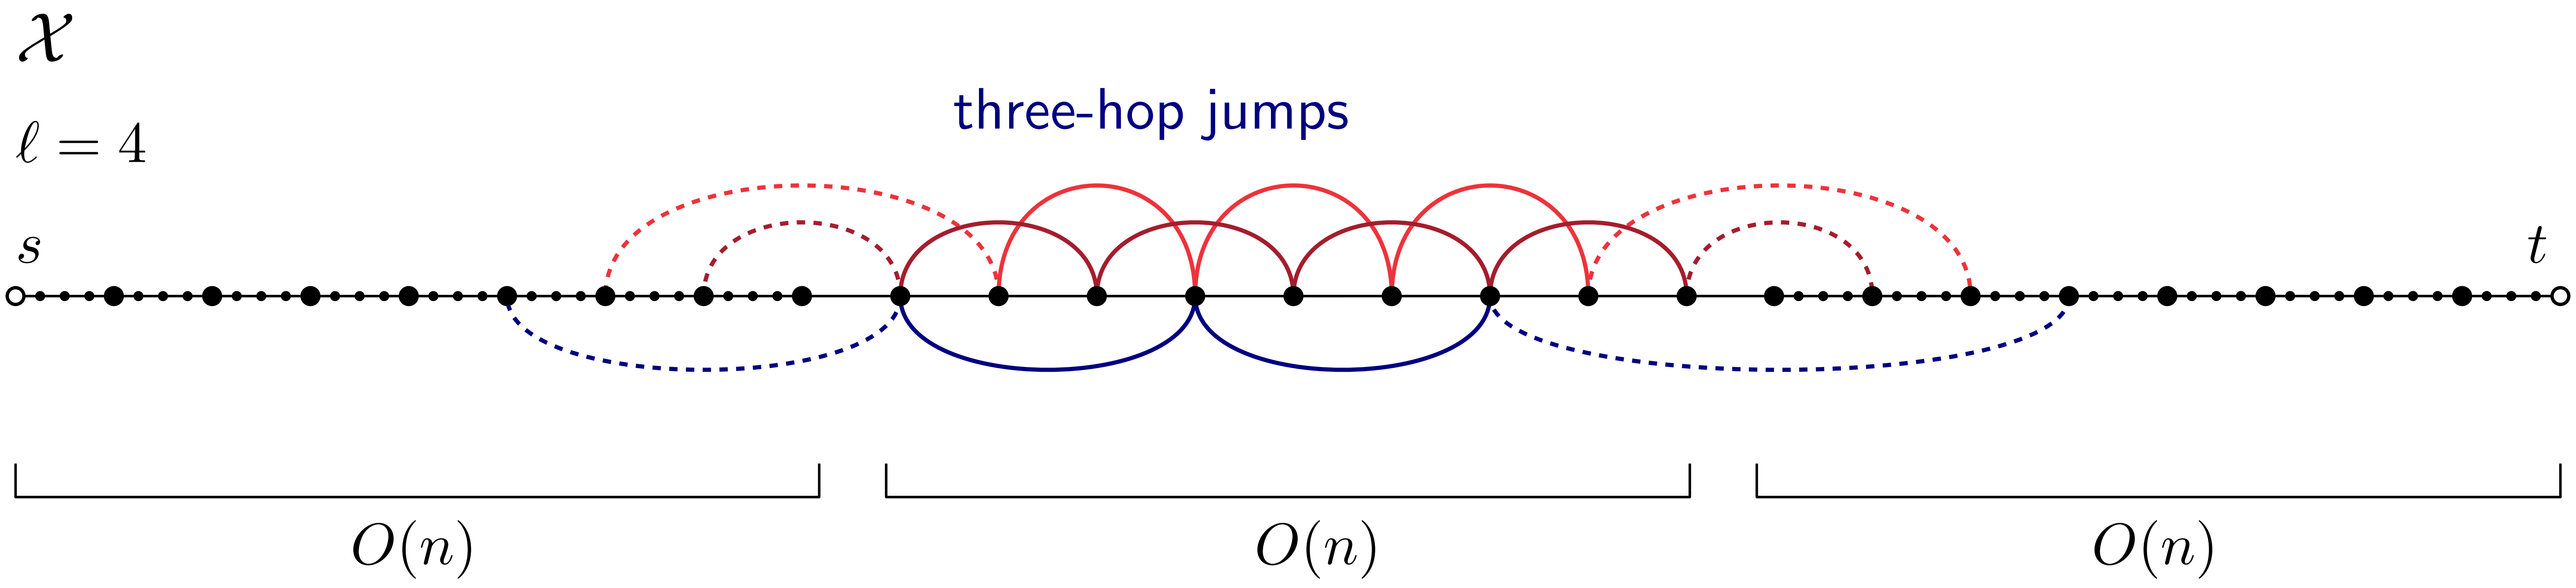

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

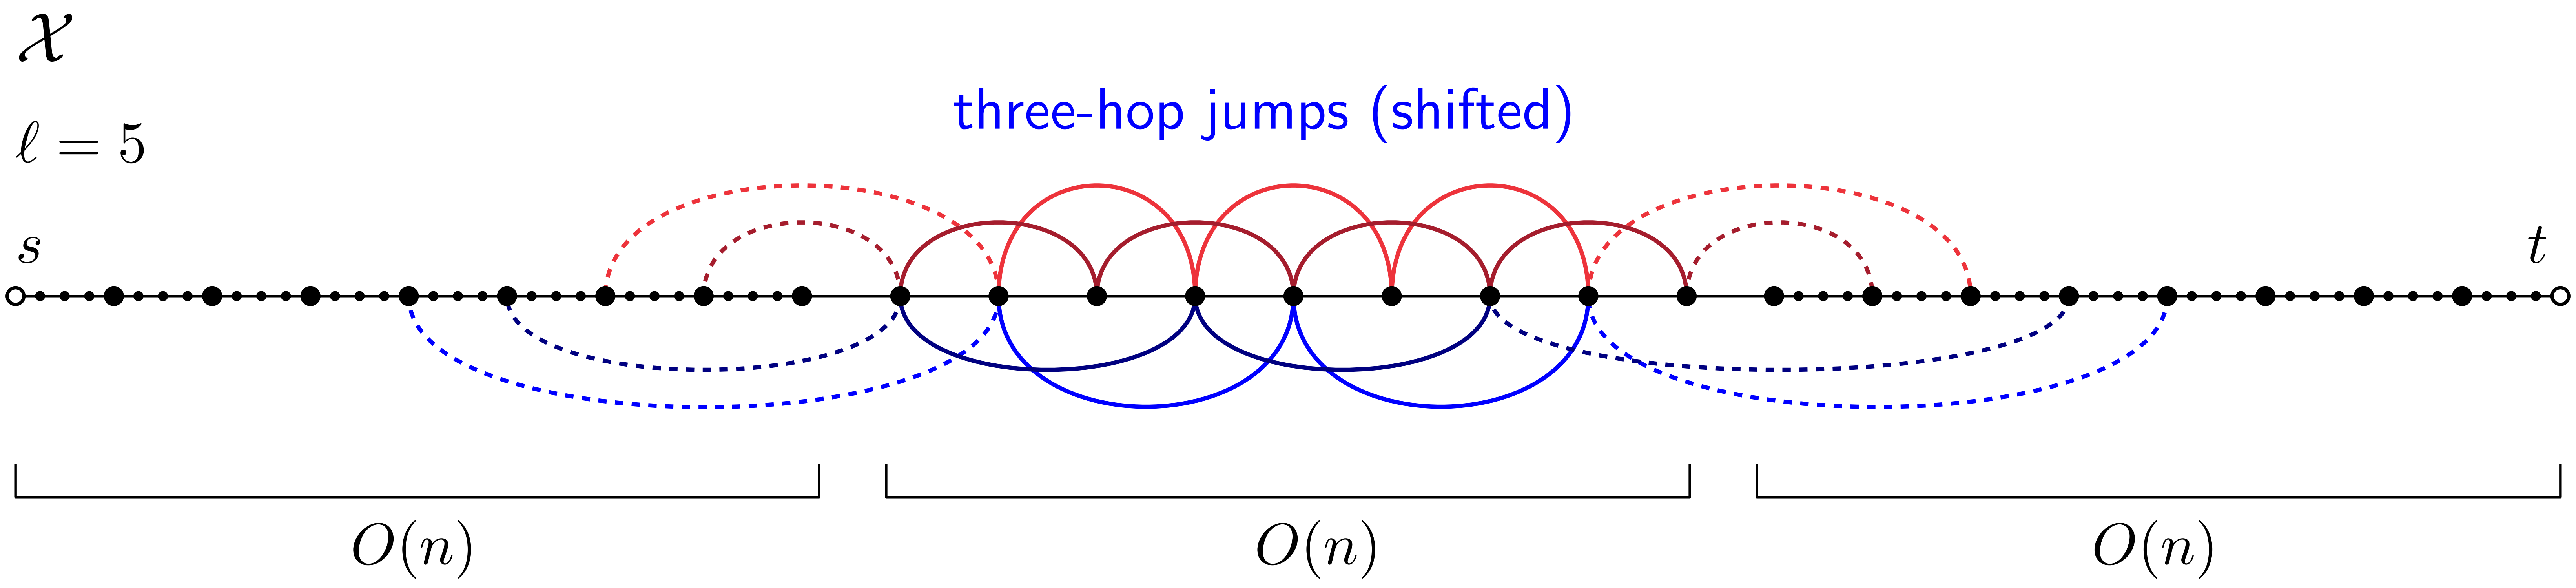

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

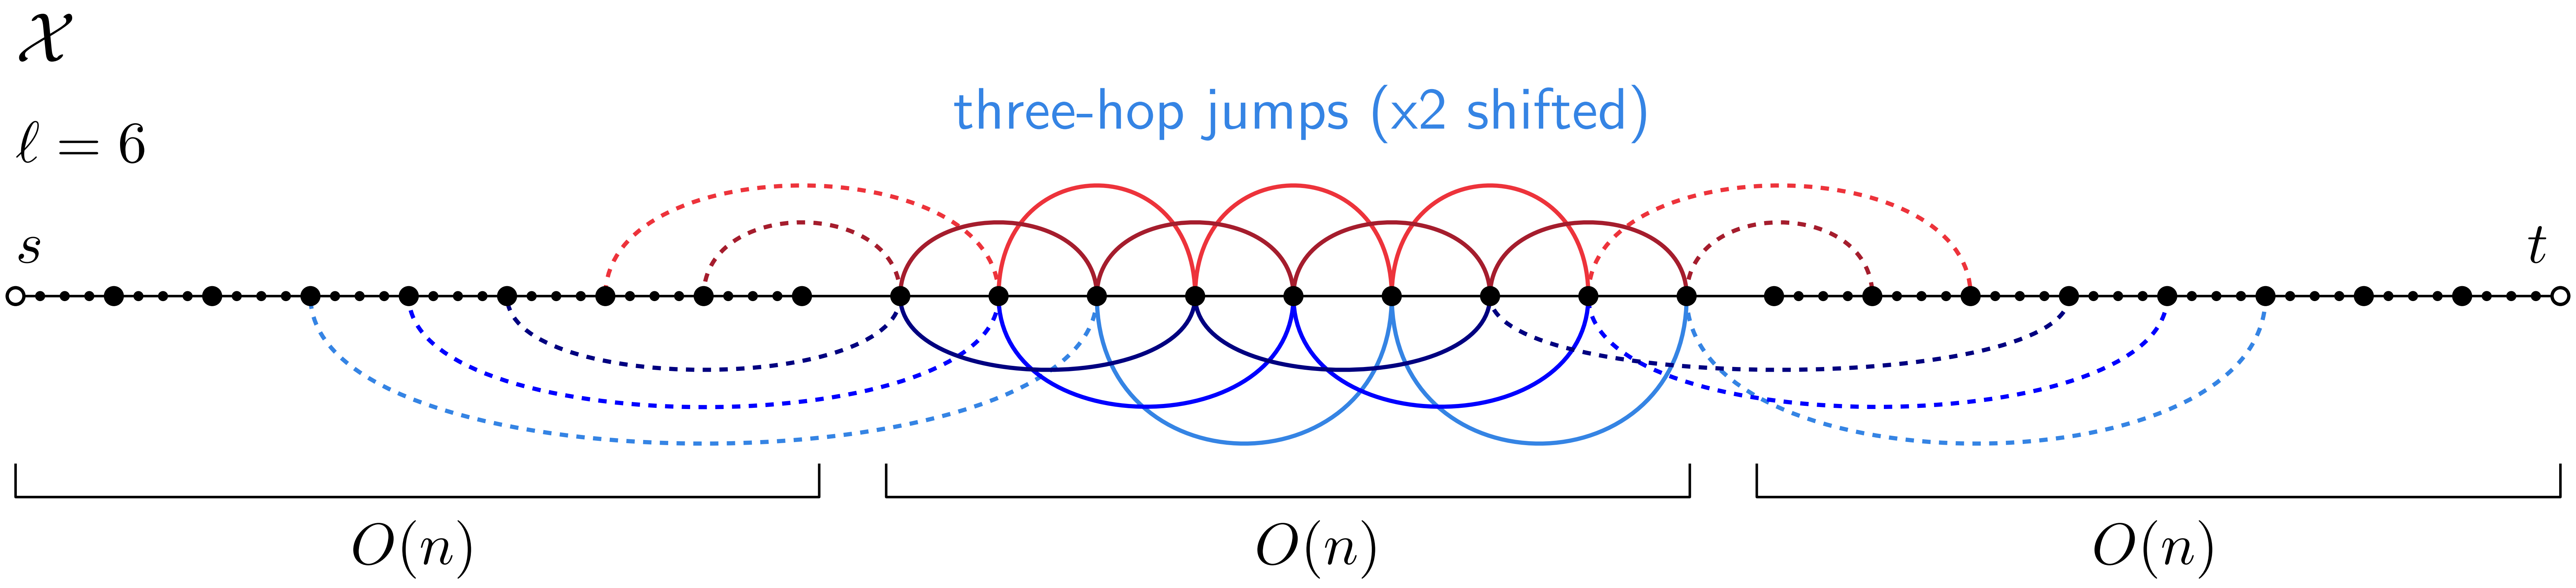

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

at most $\sqrt{n}$-hop jumps, $\sqrt{n}$ shift

$\mathcal{X}$

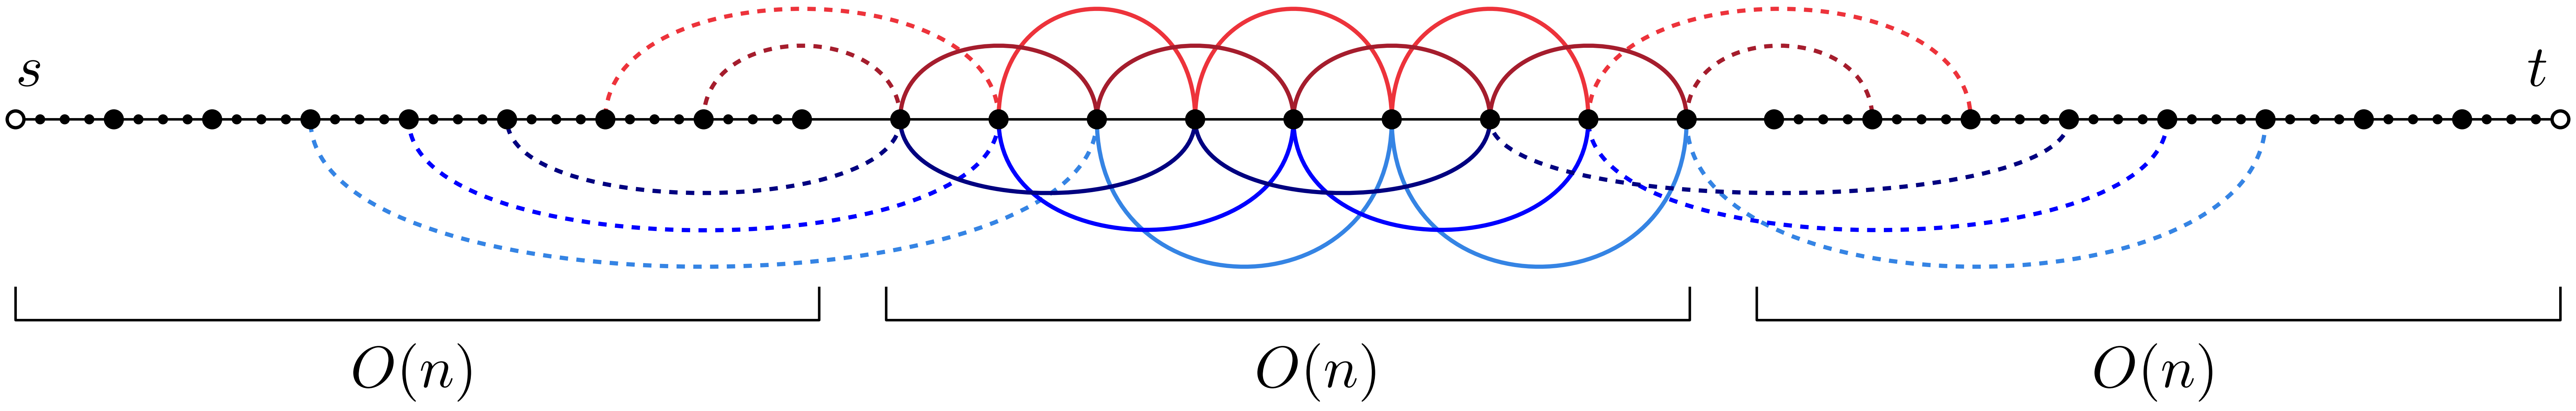

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

$L = \Theta(n)$ layers

at most $\sqrt{n}$-hop jumps, $\sqrt{n}$ shift $\quad \Longrightarrow$

$\mathcal{X}$

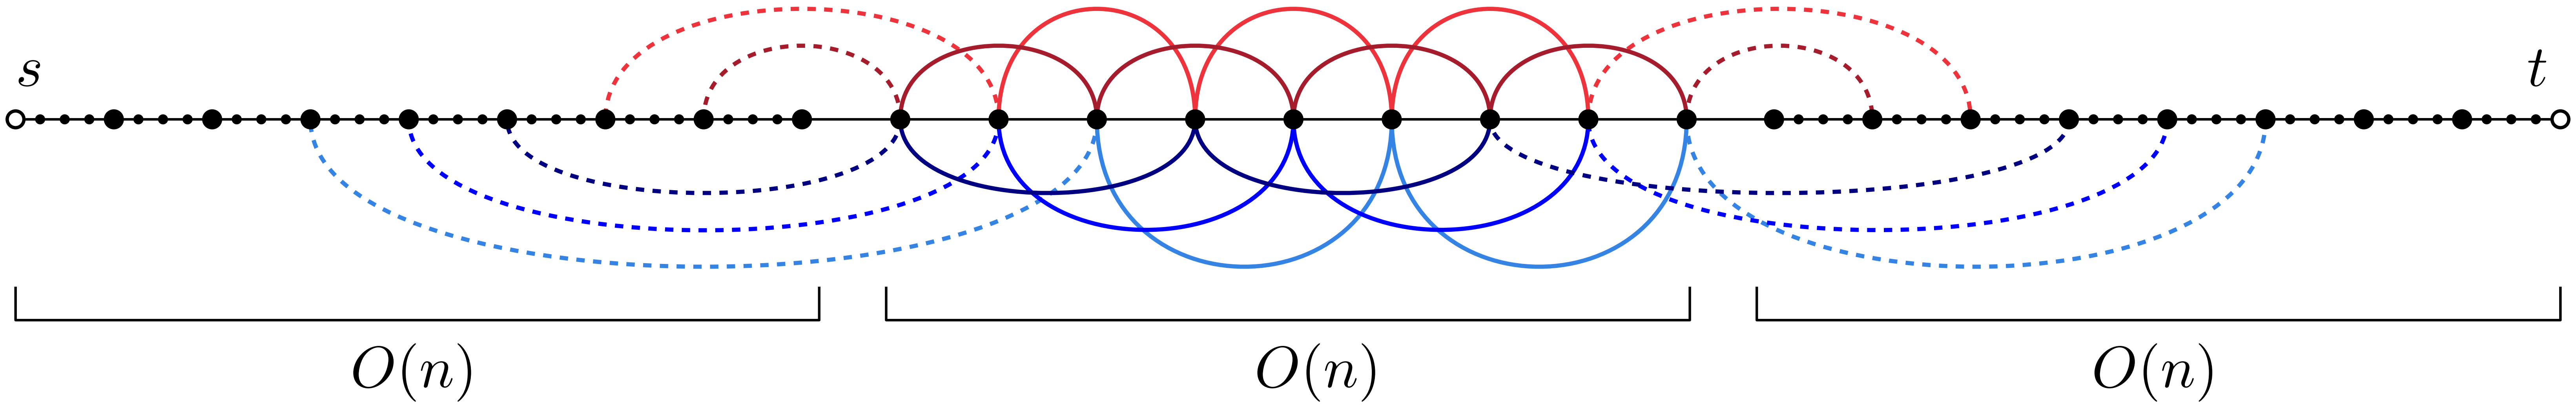

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

at most $\sqrt{n}$-hop jumps, $\sqrt{n}$ shift $\implies$ $L = \Theta(n)$ layers

$\Omega(\sqrt{n})$ per layer

$\mathcal{X}$

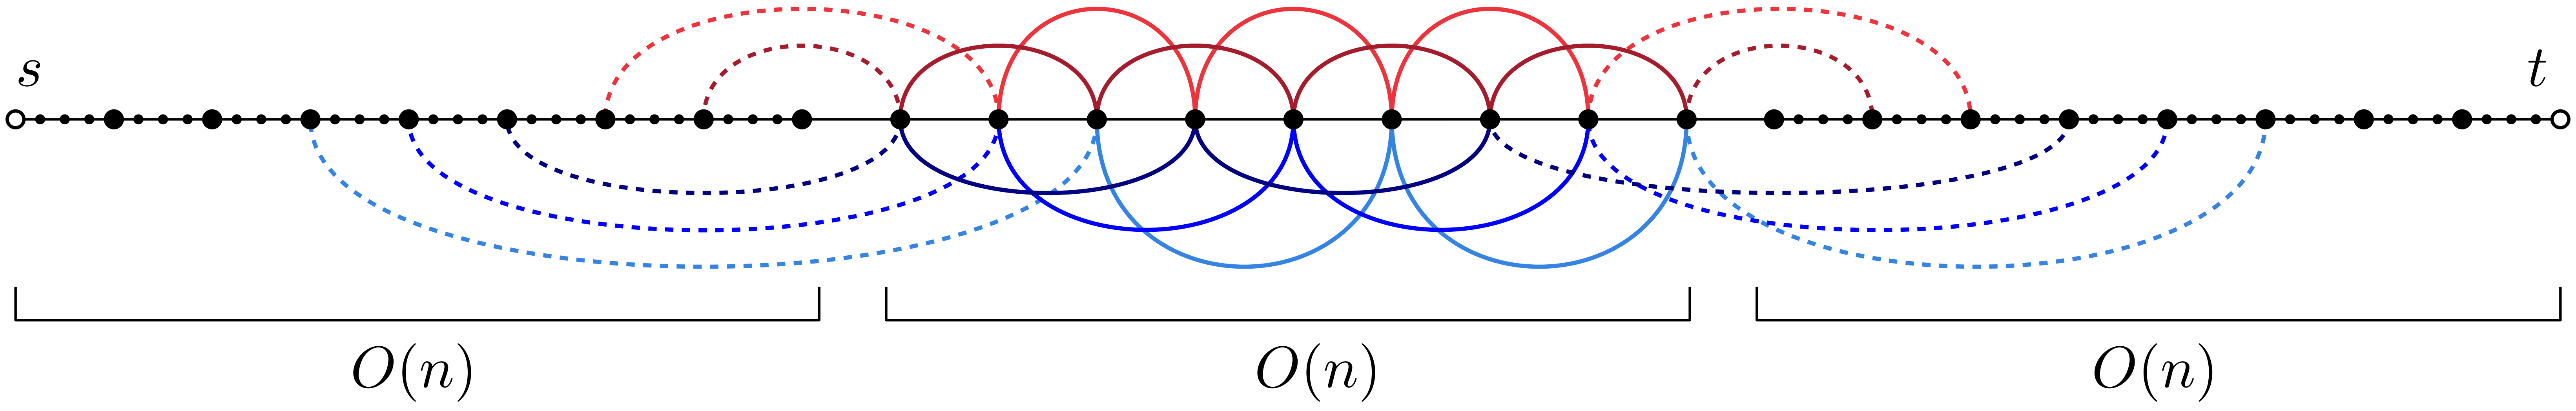

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

at most $\sqrt{n}$-hop jumps, $\sqrt{n}$ shift $\implies$ $L = \Theta(n)$ layers, $\Omega(\sqrt{n})$ per layer $\implies$ $\Omega(n\sqrt{n})$ edges

$\mathcal{X}$

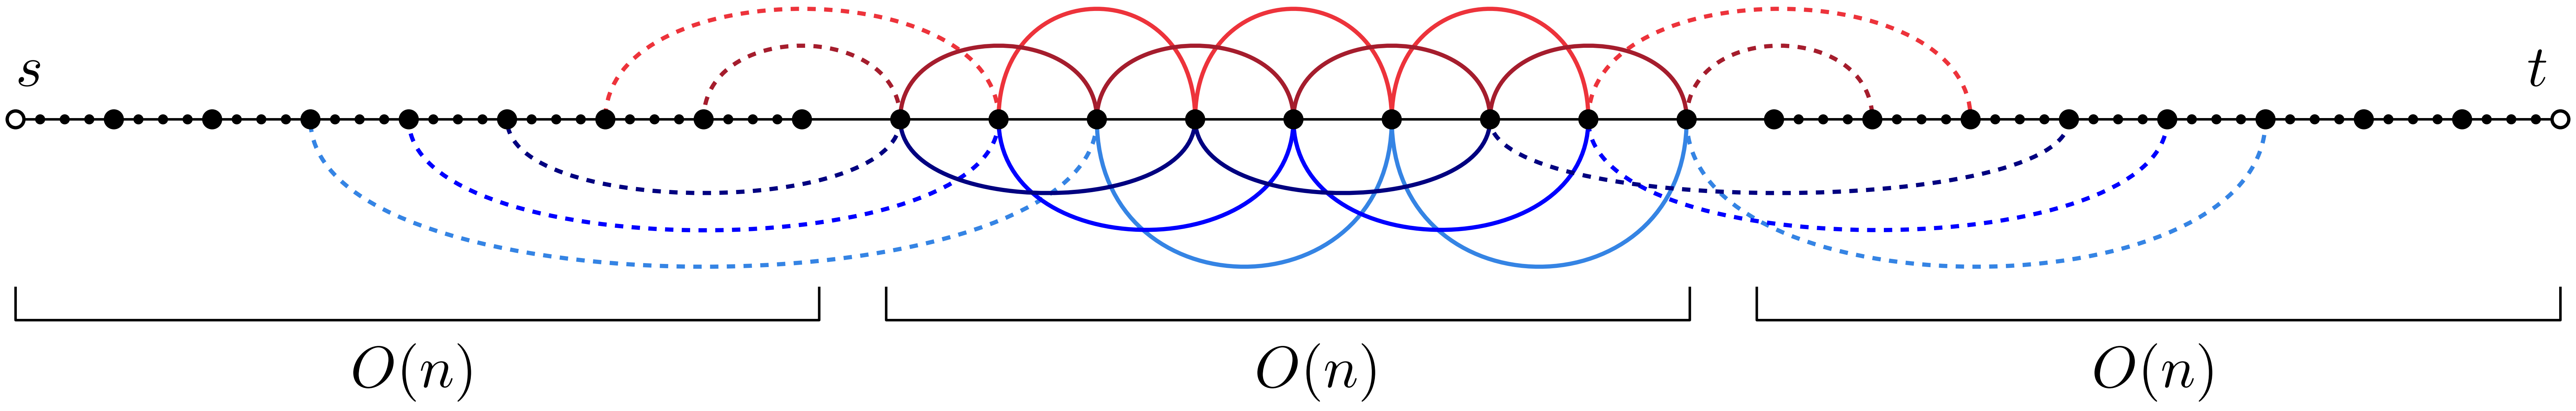

Lower bound

Single-pair preserver $O(n\sqrt{n})$ edges

at most $\sqrt{n}$-hop jumps, $\sqrt{n}$ shift $\implies$ $L = \Theta(n)$ layers, $\Omega(\sqrt{n})$ per layer $\implies$ $\Omega(n\sqrt{n})$ edges

$\mathcal{X}$

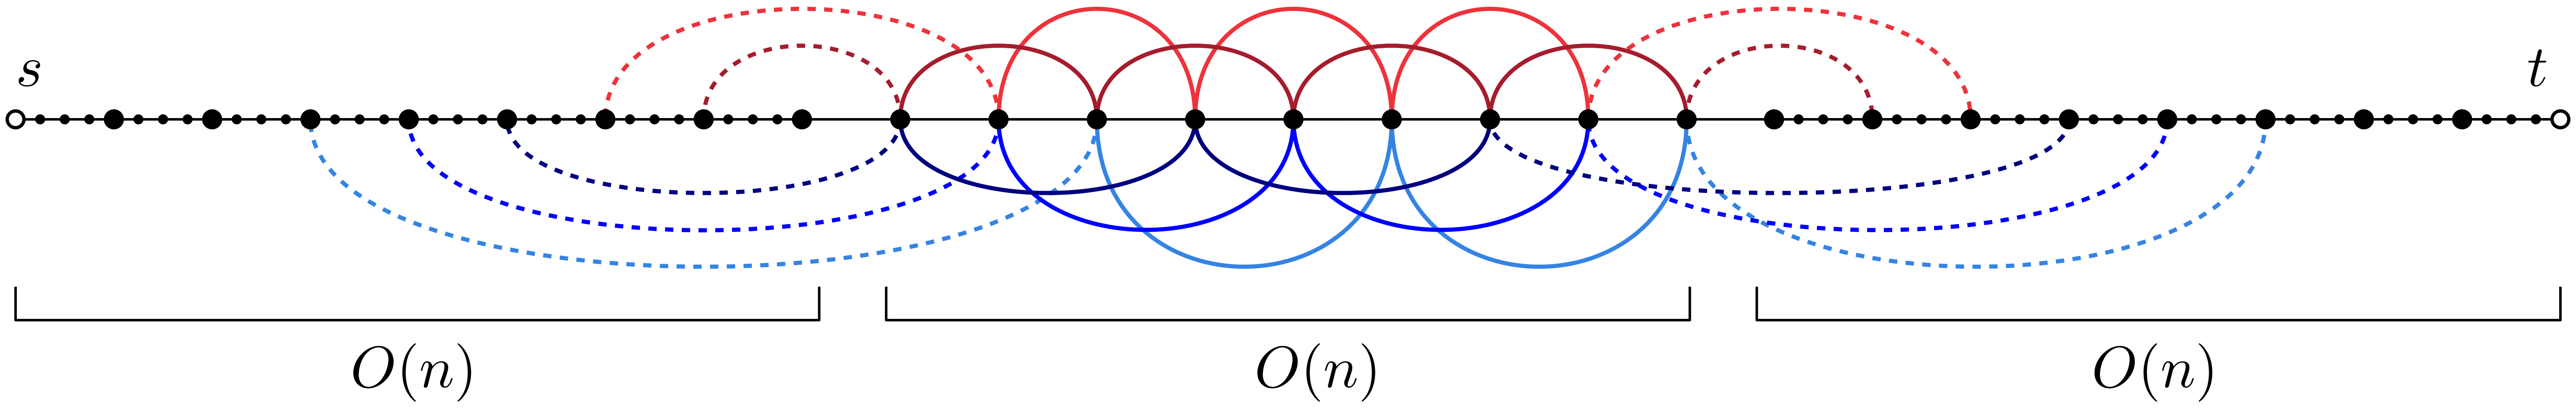

$\Theta(n)$ long, edge-disjoint, unavoidable paths $\implies$ $\Omega(n\sqrt{n})$ edges for st-preserver

Lower bound with additive stretch

Single-pair β-addivite $O(n\sqrt{\frac{n}{\beta+1}})$ edges

Lower bound with additive stretch

concatenate $2\beta+1$ gadgets $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \ldots, \mathcal{X}^{(2\beta+1)}$

Single-pair β-addivite $O(n\sqrt{\frac{n}{\beta+1}})$ edges

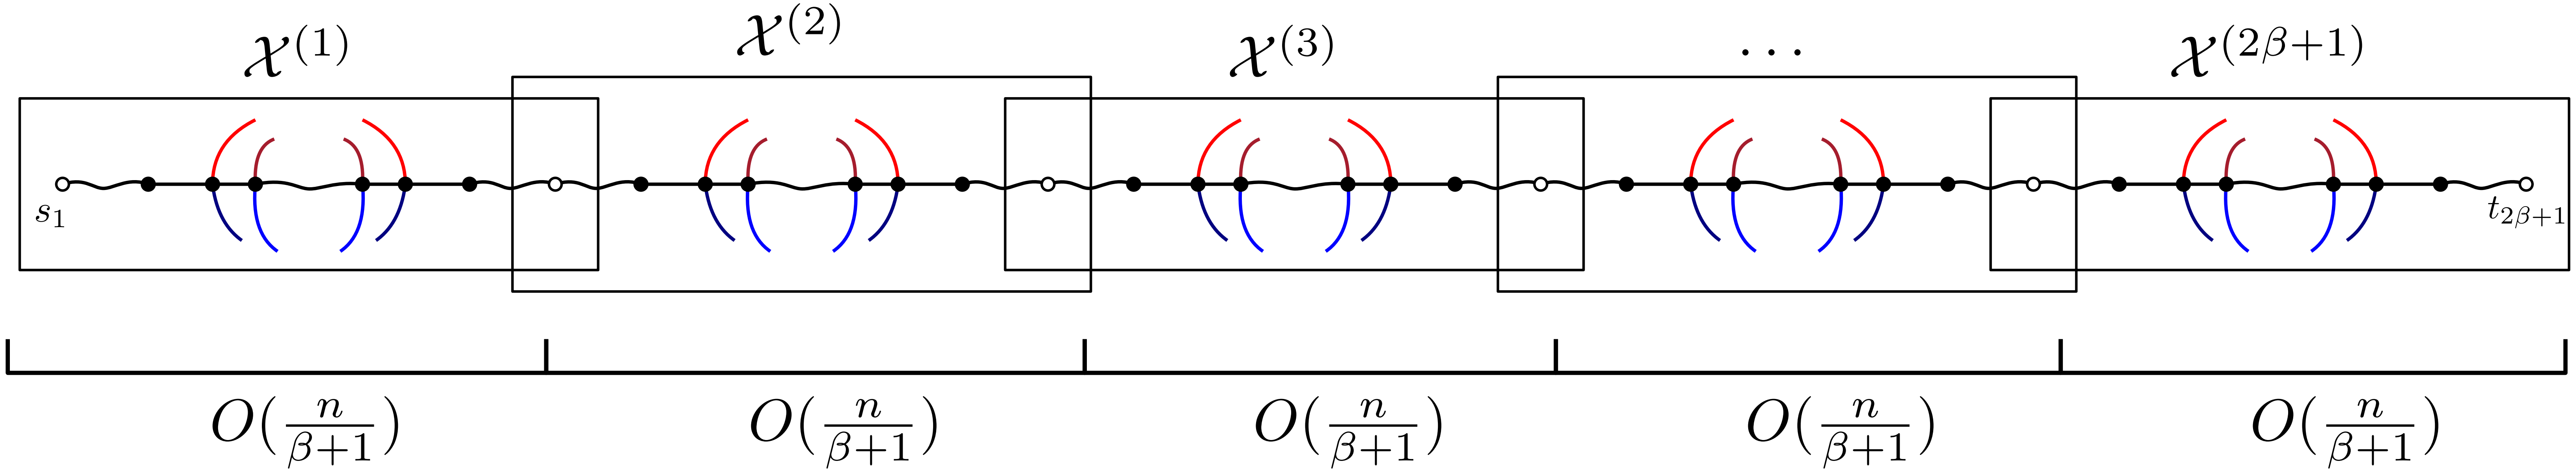

Lower bound with additive stretch

concatenate $2\beta+1$ gadgets $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \ldots, \mathcal{X}^{(2\beta+1)}$

Single-pair β-addivite $O(n\sqrt{\frac{n}{\beta+1}})$ edges

missing π_ℓ in one copy $\implies$ $+1$ error

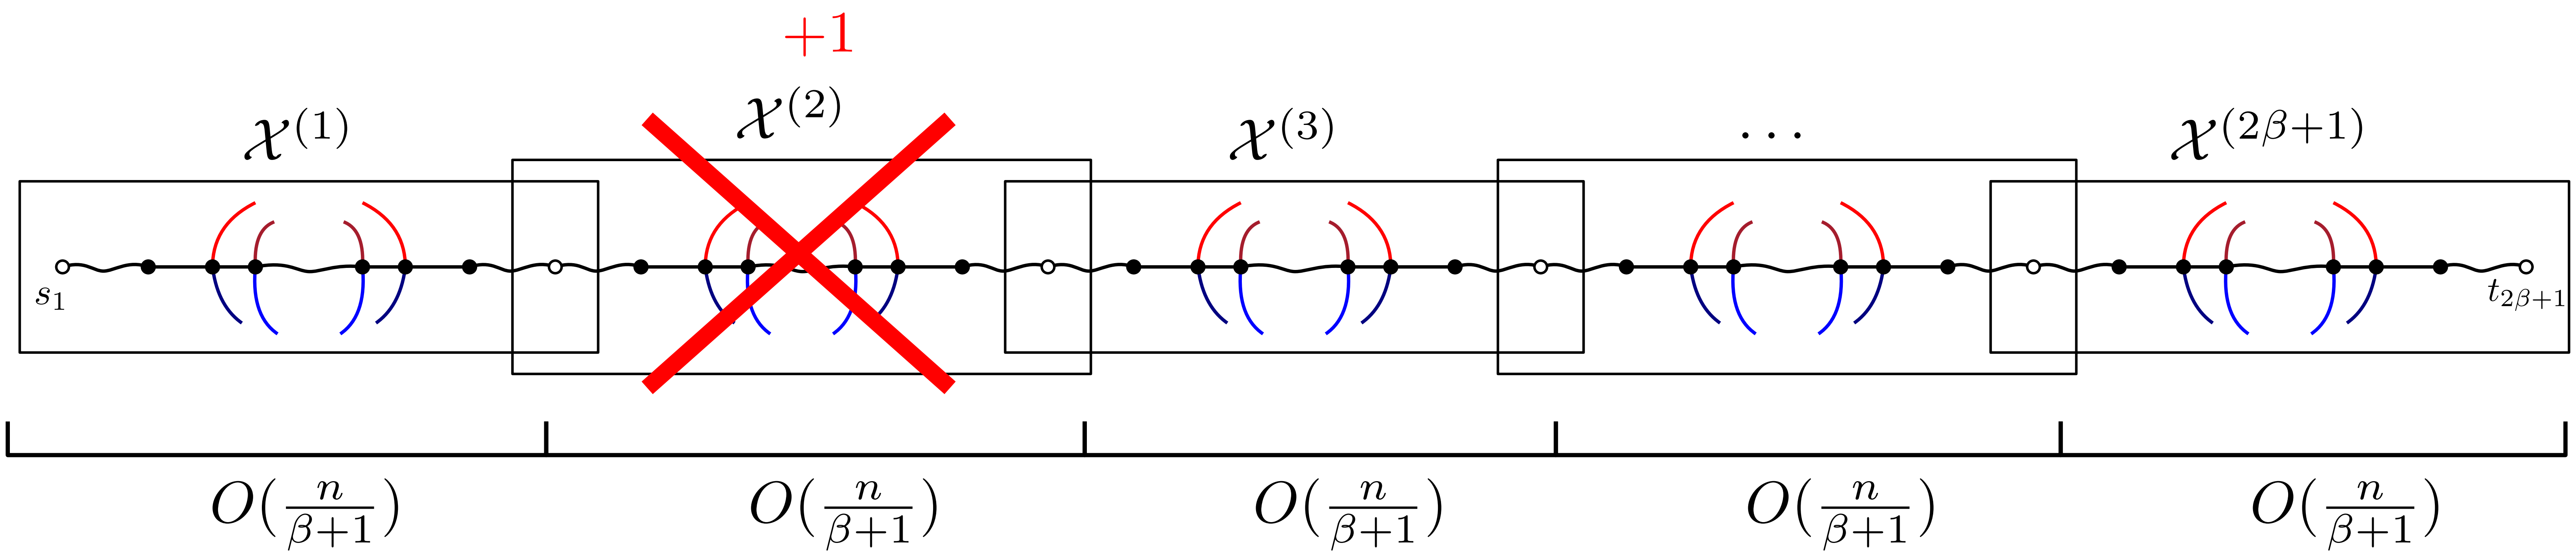

Lower bound with additive stretch

concatenate $2\beta+1$ gadgets $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \ldots, \mathcal{X}^{(2\beta+1)}$

Single-pair β-addivite $O(n\sqrt{\frac{n}{\beta+1}})$ edges

missing π_ℓ in one copy $\implies$ $+1$ error

β-additive $\implies$ π_ℓ must survive in at least $\beta+1$ copies

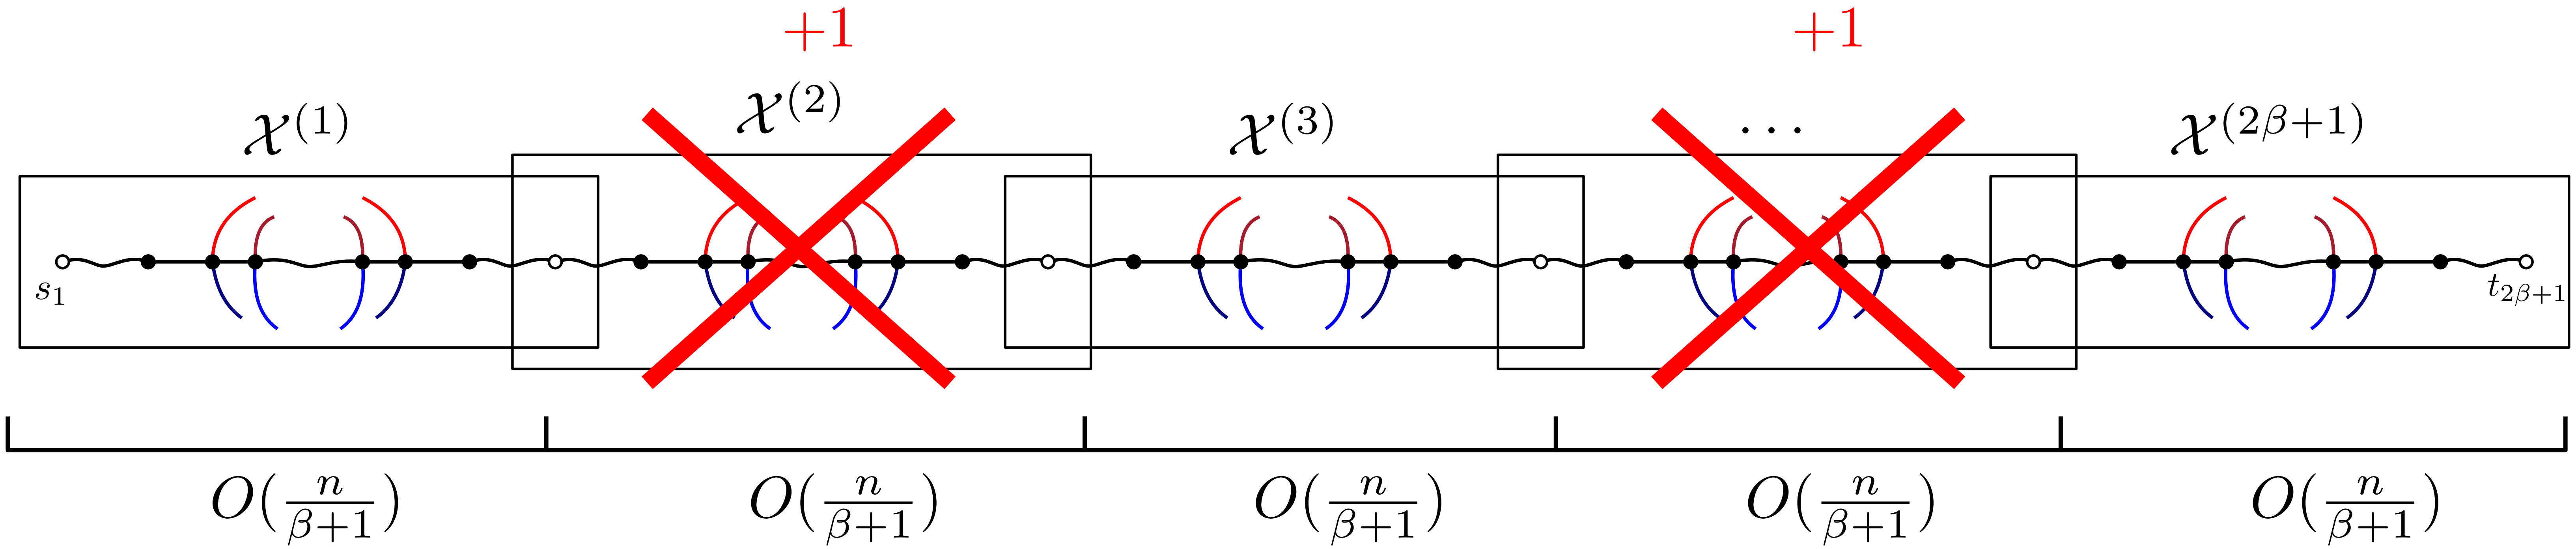

Lower bound with additive stretch

concatenate $2\beta+1$ gadgets $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \ldots, \mathcal{X}^{(2\beta+1)}$

Single-pair β-addivite $O(n\sqrt{\frac{n}{\beta+1}})$ edges

missing π_ℓ in one copy $\implies$ $+1$ error

β-additive $\implies$ π_ℓ must survive in at least $\beta+1$ copies

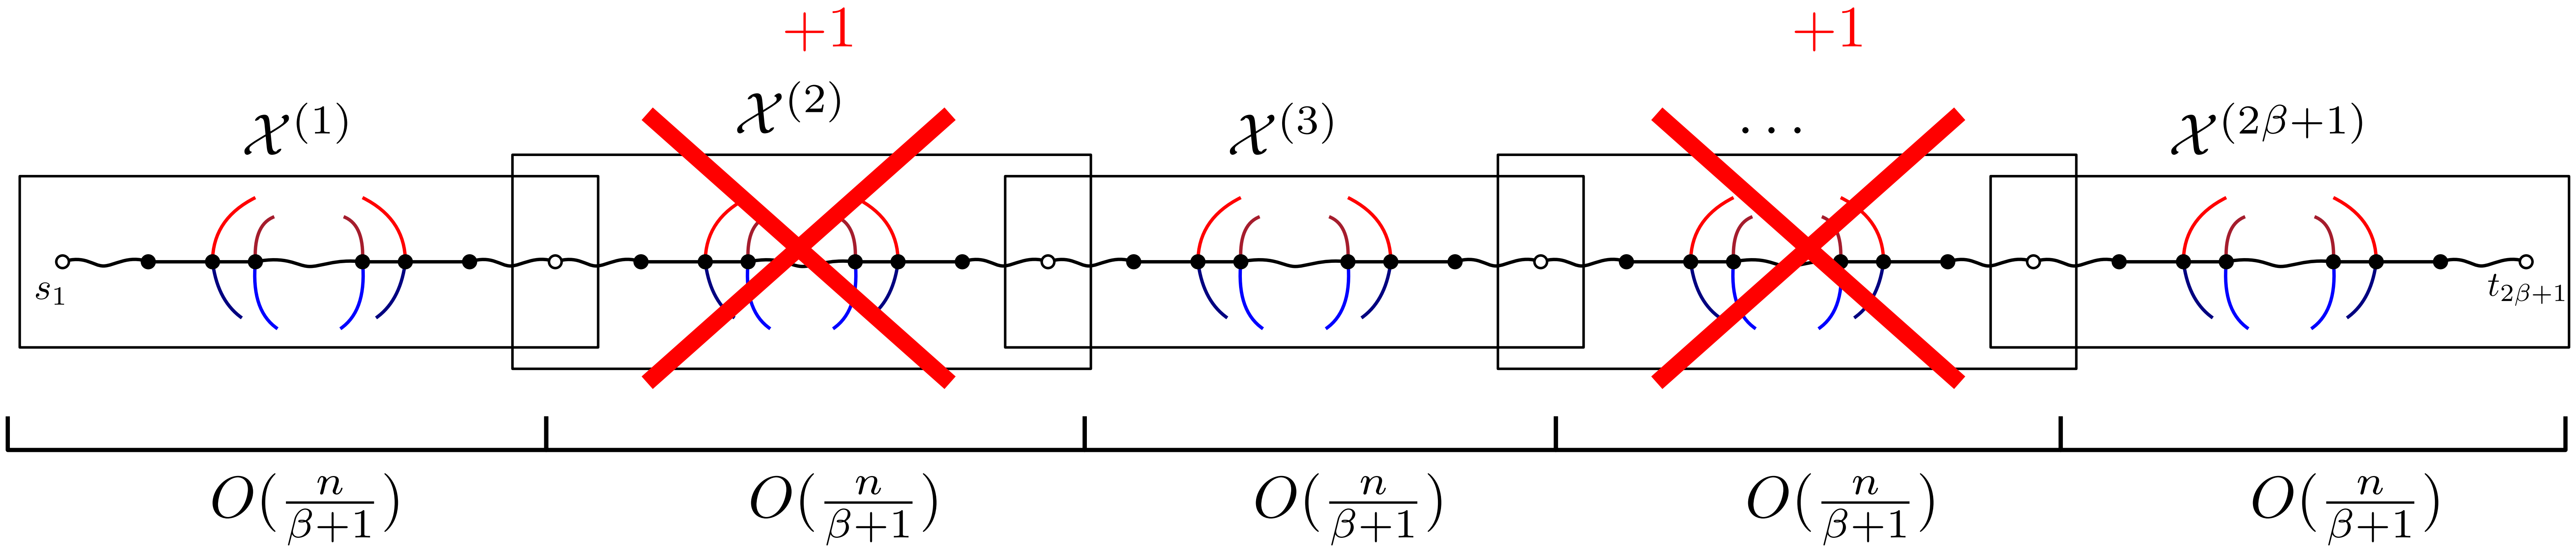

$\Omega(n\sqrt{\frac{n}{\beta+1}})$ edges

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

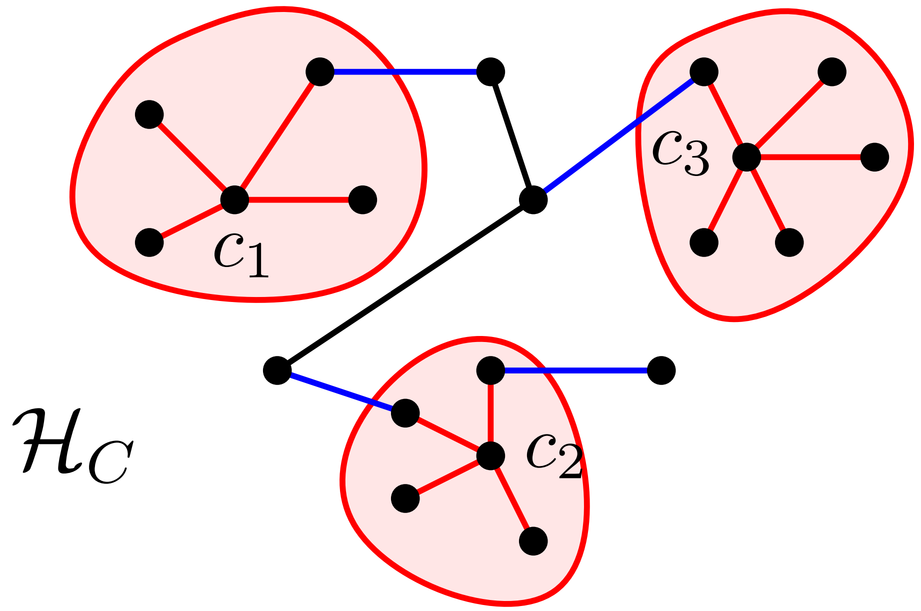

Hierarchical clustering:

- $C = \{c_1, c_2, \ldots\}$, cluster centers
- $|C_i| = \Delta$
- $|C| \leq \frac{n}{\Delta}$
- $|\mathcal{H}_C| = O(n\Delta)$

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

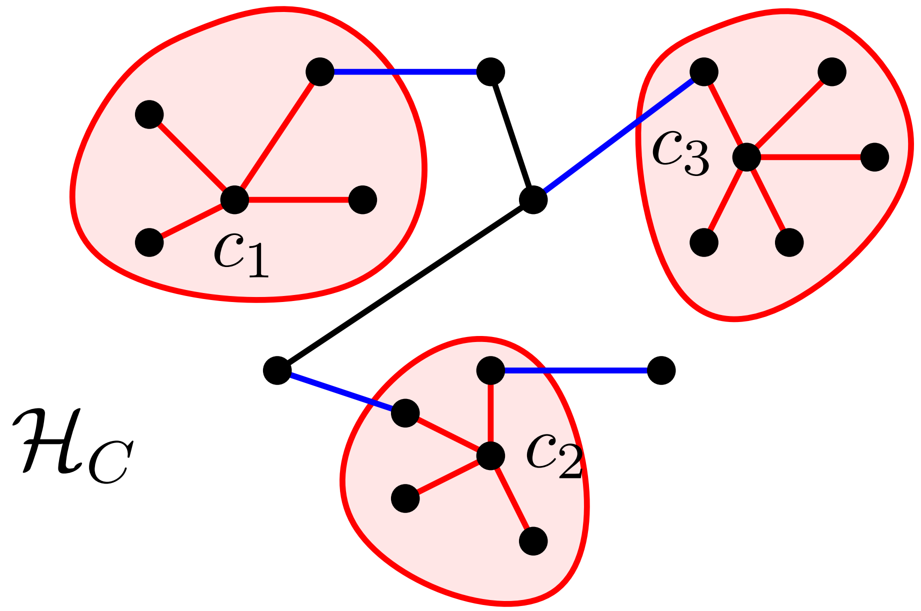

Hierarchical clustering:

- $C = \{c_1, c_2, \dots\}$, cluster centers
- $|C_i| = \Delta$
- $|C| \le \frac{n}{\Delta}$
- $|\mathcal{H}_C| = O(n\Delta)$

Build a hierarchical spanner only involving the centers

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

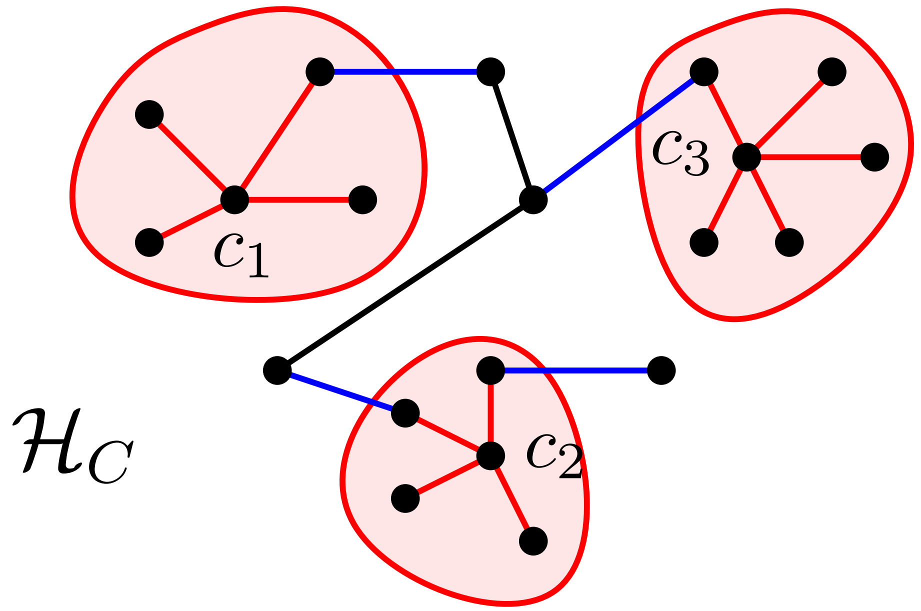

Hierarchical clustering:

- $C = \{c_1, c_2, \dots\}$, cluster centers
- $|C_i| = \Delta$
- $|C| \leq \frac{n}{\Delta}$
- $|\mathcal{H}_C| = O(n\Delta)$

Build a hierarchical spanner only involving the centers

Thm 9. $\mathcal{H}_C$ plus **subsetwise** $C \times C$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha,\ 2\alpha + \beta + 2)$.

Thm 11. $\mathcal{H}_C$ plus **sourcewise** $C \times V$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha, \alpha + \beta + 1)$.

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

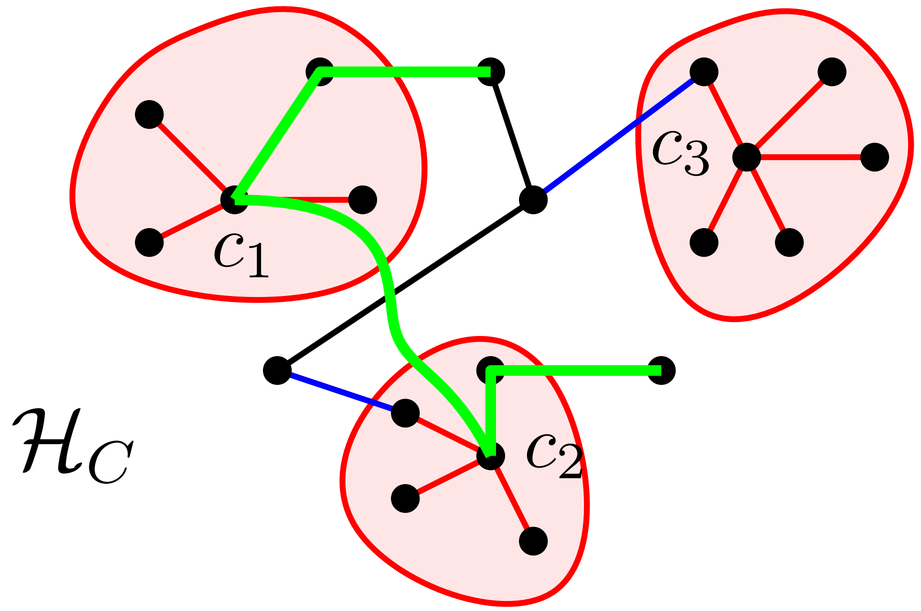

Hierarchical clustering:

- $C = \{c_1, c_2, \dots\}$, cluster centers
- $|C_i| = \Delta$
- $|C| \leq \frac{n}{\Delta}$
- $|\mathcal{H}_C| = O(n\Delta)$

Proof: detour analysis through the appropriate cluster centers in the appropriate layer

Build a hierarchical spanner only involving the centers

Thm 9. $\mathcal{H}_C$ plus **subsetwise** $C \times C$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha,\ 2\alpha + \beta + 2)$.

Thm 11. $\mathcal{H}_C$ plus **sourcewise** $C \times V$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha, \alpha + \beta + 1)$.

From one-pair to all-pairs

Idea: If we can preserve distances among a few cluster centers, we can approximately preserve all pairs.

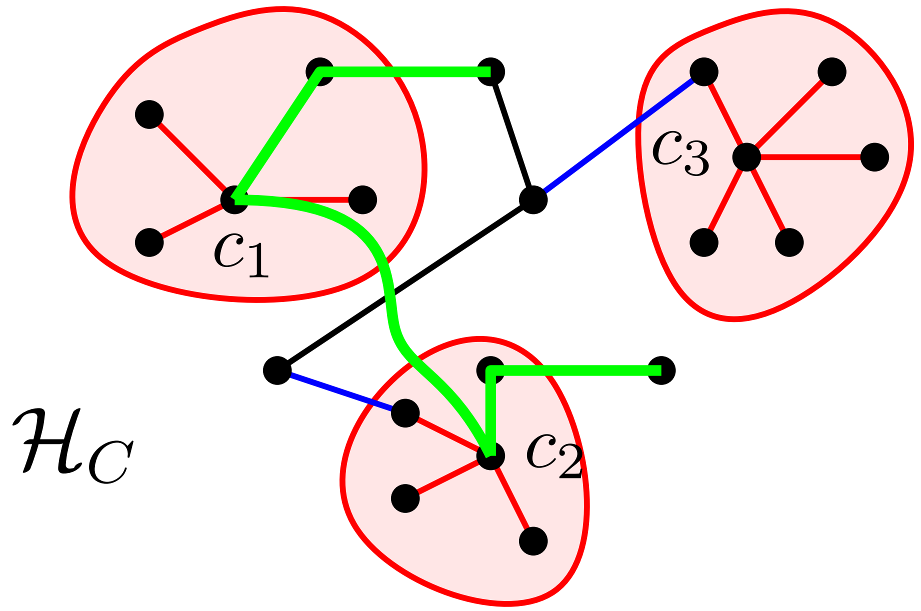

Hierarchical clustering:

- $C = \{c_1, c_2, \dots\}$, cluster centers
- $|C_i| = \Delta$
- $|C| \leq \frac{n}{\Delta}$
- $|\mathcal{H}_C| = O(n\Delta)$

Proof: detour analysis through the appropriate cluster centers in the appropriate layer

Build a hierarchical spanner only involving the centers

Thm 9. $\mathcal{H}_C$ plus **subsetwise** $C \times C$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha,\ 2\alpha + \beta + 2)$.

$$S \times S \text{ preserver } \implies 4\text{-additive}, \widetilde{O}(n^{5/3})$$

Thm 11. $\mathcal{H}_C$ plus **sourcewise** $C \times V$ hierarchical (α, β)-spanner $\implies$ all-pairs $(\alpha, \alpha + \beta + 1)$.

$$S \times V\ (1+\varepsilon)\text{-spanner } \implies (1+\varepsilon, 2)\text{-stretch}, \widetilde{O}(n\sqrt{n/\varepsilon})$$

Takeaways

Hierarchical spanners behave very differently from classical spanners

Takeaways

Hierarchical spanners behave very differently from classical spanners

- **Multiplicative stretch is robust to hierarchy.**

Classical $(2k-1)$-spanners with $O(n^{1+1/k})$ edges extend essentially unchanged to the hierarchical setting.

Takeaways

Hierarchical spanners behave very differently from classical spanners

- **Multiplicative stretch is robust to hierarchy.**

Classical $(2k-1)$-spanners with $O(n^{1+1/k})$ edges extend essentially unchanged to the hierarchical setting.

- **Exact and additive preservation can become harder.**

Even for single pair $\Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$ tight bound, in particular, preservers require $\Theta(n^{3/2})$.

Takeaways

Hierarchical spanners behave very differently from classical spanners

- **Multiplicative stretch is robust to hierarchy.**

Classical $(2k-1)$-spanners with $O(n^{1+1/k})$ edges extend essentially unchanged to the hierarchical setting.

- **Exact and additive preservation can become harder.**

Even for single pair $\Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$ tight bound, in particular, preservers require $\Theta(n^{3/2})$.

- **A small multiplicative relaxation restores sparsity.**

single-source preserver $\Omega(n^2)$ $\longrightarrow$ single-source $(1+\varepsilon)$-spanner essentialy linear
all-pair 2-additive unlikely size $O(n\sqrt{n})$ $\longrightarrow$ $(1+\varepsilon, 2)$-spanner essentially $O(n\sqrt{n})$

Takeaways

Hierarchical spanners behave very differently from classical spanners

- **Multiplicative stretch is robust to hierarchy.**

Classical $(2k-1)$-spanners with $O(n^{1+1/k})$ edges extend essentially unchanged to the hierarchical setting.

- **Exact and additive preservation can become harder.**

Even for single pair $\Theta\left(n\sqrt{\frac{n}{\beta+1}}\right)$ tight bound, in particular, preservers require $\Theta(n^{3/2})$.

- **A small multiplicative relaxation restores sparsity.**

single-source preserver $\Omega(n^2)$ $\longrightarrow$ single-source $(1+\varepsilon)$-spanner essentialy linear
all-pair 2-additive unlikely size $O(n\sqrt{n})$ $\longrightarrow$ $(1+\varepsilon, 2)$-spanner essentially $O(n\sqrt{n})$

Hierarchy hurts exacts / additive preservation more then multiplicative approximation

Open problems

1. How hard are all-pairs additive hierarchical spanners?

Open problems

1. How hard are all-pairs additive hierarchical spanners?

single-pair $O(1)$-additive $\Omega(n^{3/2})$ **[gap]** all-pair 4-additive $O(n^{5/3})$

Open problems

1. How hard are all-pairs additive hierarchical spanners?

single-pair $O(1)$-additive $\Omega(n^{3/2})$ **[gap]** all-pair 4-additive $O(n^{5/3})$

Can we prove that all-pairs additive spanners are harder than the single-pair barrier?

For example, $\Omega(n^{3/2+\delta})$ for 2-additive all-pairs?

Open problems

1. How hard are all-pairs additive hierarchical spanners?

single-pair $O(1)$-additive $\Omega(n^{3/2})$ **[gap]** all-pair 4-additive $O(n^{5/3})$

Can we prove that all-pairs additive spanners are harder than the single-pair barrier?

For example, $\Omega(n^{3/2+\delta})$ for 2-additive all-pairs?

2. What is the tight bound for $s \times T$ preservers?

Open problems

1. How hard are all-pairs additive hierarchical spanners?

single-pair $O(1)$-additive $\Omega(n^{3/2})$ **[gap]** all-pair 4-additive $O(n^{5/3})$

Can we prove that all-pairs additive spanners are harder than the single-pair barrier?

For example, $\Omega(n^{3/2+\delta})$ for 2-additive all-pairs?

2. What is the tight bound for $s \times T$ preservers?

$\Omega(n|T|)$ **[gap]** $O(n\sqrt{n|T|})$

Open problems

1. How hard are all-pairs additive hierarchical spanners?

single-pair $O(1)$-additive $\Omega(n^{3/2})$ **[gap]** all-pair 4-additive $O(n^{5/3})$

Can we prove that all-pairs additive spanners are harder than the single-pair barrier?

For example, $\Omega(n^{3/2+\delta})$ for 2-additive all-pairs?

2. What is the tight bound for $s \times T$ preservers?

$\Omega(n|T|)$ **[gap]** $O(n\sqrt{n|T|})$

What's here?

Thank you for your attention!