Approximate Distance Sensitivity Oracles in Subquadratic Space

Davide Bilò, Shiri Chechik, Keerti Choudhary, Sarel Cohen, Tobias Friedrich, Simon Krogmann, and Martin Schirneck

55th Symposium on Theory of Computing
June 23, 2023

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

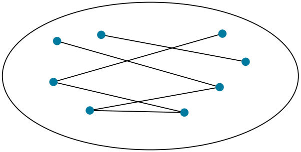

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

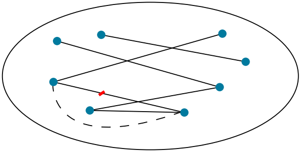

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

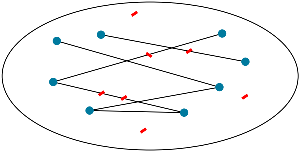

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

- Sensitivity: failures in batches, maximum number f is known.

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

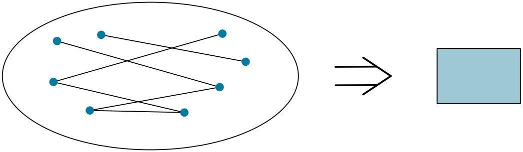

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

- Sensitivity: failures in batches, maximum number f is known.
- Data structure: preprocess once, query when needed.

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

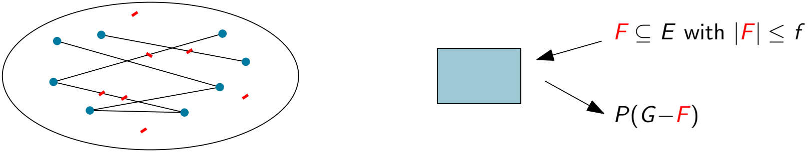

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

- Sensitivity: failures in batches, maximum number f is known.
- Data structure: preprocess once, query when needed.

Fault-Tolerant Data Structures

a.k.a. sensitivity data structures, algorithms for emergency planning, failure-prone graphs.

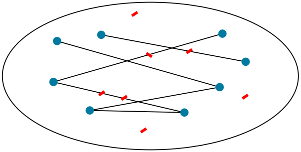

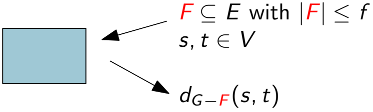

Maintain graph property $P(G)$ (distances, connectivity, ...) under edge failures.

- Sensitivity: failures in batches, maximum number f is known.
- Data structure: preprocess once, query when needed.

This talk: $P(G) = d(s, t)$ - (approximate) pairwise distances,

f-edge fault-tolerant distance sensitivity oracle (f-DSO).

Problem Parameters

Sensitivity f

- Sensitivity f: number of edge failures.

Problem Parameters

Sensitivity f

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies

  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.
- Preprocessing time.

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.
- Preprocessing time.
- Input graphs: (un-)directed, (un-)weighted, ...

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.
- Preprocessing time.
- Input graphs: (un-)directed, (un-)weighted, ...

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.
- Preprocessing time.
- Input graphs: (un-)directed, (un-)weighted, ...

Most f-DSOs in the literature take $\Omega(n^2)$ space.

[Duan & Ren STOC 2022] [Gu & Ren ICALP 2021] [Grandoni & Vassilevska Williams TALG 2020] [Chechik, Cohen, Fiat & Kaplan SODA 2017] [Weimann & Yuster JACM 2013] [Berstein & Karger STOC 2009] ...

Problem Parameters

Sensitivity f, space in $O(\log n)$-bit words.

- Sensitivity f: number of edge failures.
- Stretch σ: returned value $\widehat{d}$ satisfies
  $d_{G-F}(s,t) \leq \widehat{d}_{G-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
- Space: measured in $O(\log n)$-bit machine words.
- Query time.
- Preprocessing time.
- Input graphs: (un-)directed, (un-)weighted, ...

Most f-DSOs in the literature take $\Omega(n^2)$ space.

[Duan & Ren STOC 2022] [Gu & Ren ICALP 2021] [Grandoni & Vassilevska Williams TALG 2020] [Chechik, Cohen, Fiat & Kaplan SODA 2017] [Weimann & Yuster JACM 2013] [Berstein & Karger STOC 2009] ...

Our goal: f-DSOs in $o(n^2)$ space.

Subquadratic Space

Sensitivity f, space in $O(\log n)$-bit words,

Space lower bounds, even without failures ($f = 0$). [Thorup & Zwick JACM 2005]

- If G is undirected, any DSO must have space $\Omega(n^2)$ or stretch $\sigma \geq 3$.

Subquadratic Space

Sensitivity f, space in $O(\log n)$-bit words, undirected graphs,

Space lower bounds, even without failures ($f = 0$). [Thorup & Zwick JACM 2005]

- If G is undirected, any DSO must have space $\Omega(n^2)$ or stretch $\sigma \geq 3$.
- If G is directed, any DSO must have space $\Omega(n^2)$ for arbitrary finite stretch.

Subquadratic Space

Sensitivity f, space in $O(\log n)$-bit words, undirected graphs, $\widetilde{O}(\cdot)$ hides poly$(\log n)$ factors.

Space lower bounds, even without failures ($f = 0$). [Thorup & Zwick JACM 2005]

- If G is undirected, any DSO must have space $\Omega(n^2)$ or stretch $\sigma \geq 3$.
- If G is directed, any DSO must have space $\Omega(n^2)$ for arbitrary finite stretch.

Previous solutions (weighted graphs, any sensitivity f and integer $k \geq 1$).

- Chechik, Langberg, Peleg & Roditty. [Algorithmica 2012]
  Stretch $(8k-2)(f+1)$, space $\widetilde{O}(fk\, n^{1+1/k})$, query $\widetilde{O}(f)$.

Subquadratic Space

Sensitivity f, space in $O(\log n)$-bit words, undirected graphs, $\widetilde{O}(\cdot)$ hides poly($\log n$) factors.

Space lower bounds, even without failures ($f = 0$). [Thorup & Zwick JACM 2005]

- If G is undirected, any DSO must have space $\Omega(n^2)$ or stretch $\sigma \geq 3$.
- If G is directed, any DSO must have space $\Omega(n^2)$ for arbitrary finite stretch.

Previous solutions (weighted graphs, any sensitivity f and integer $k \geq 1$).

- Chechik, Langberg, Peleg & Roditty. [Algorithmica 2012]
  Stretch $(8k-2)(f+1)$, space $\widetilde{O}(fk\, n^{1+1/k})$, query $\widetilde{O}(f)$.
- Fault-tolerant σ-spanner: subgraph $H \subseteq G$ s.t. for all $F \subseteq E$ with $|F| \leq f$
  $d_{H-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.

Subquadratic Space

Sensitivity f, space in $O(\log n)$-bit words, undirected graphs, $\widetilde{O}(\cdot)$ hides poly($\log n$) factors.

Space lower bounds, even without failures ($f = 0$). [Thorup & Zwick JACM 2005]

- If G is undirected, any DSO must have space $\Omega(n^2)$ or stretch $\sigma \geq 3$.
- If G is directed, any DSO must have space $\Omega(n^2)$ for arbitrary finite stretch.

Previous solutions (weighted graphs, any sensitivity f and integer $k \geq 1$).

- Chechik, Langberg, Peleg & Roditty. [Algorithmica 2012]
  Stretch $(8k-2)(f+1)$, space $\widetilde{O}(fk\, n^{1+1/k})$, query $\widetilde{O}(f)$.
- Fault-tolerant σ-spanner: subgraph $H \subseteq G$ s.t. for all $F \subseteq E$ with $|F| \leq f$ $d_{H-F}(s,t) \leq \sigma \cdot d_{G-F}(s,t)$.
  Stretch $2k-1$, space $O(f^{1/2} k^{O(1)} n^{1+1/k})$, query time $\approx$ space.
  [Bodwin, Dinitz & Robelle SODA 2022]

Our Result

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Theorem

Let $0 < \alpha < 1/2$ and $\varepsilon > 0$. For any constant f, there is an f-DSO for undirected, unweighted graphs with unique shortest paths that has stretch $3 + \varepsilon$, space $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}}) \cdot O_\varepsilon(\log n)^{f+1}$, and query time $O_\varepsilon(n^\alpha)$.

Our Result

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Theorem

Let $0 < \alpha < 1/2$ and $\varepsilon > 0$. For any constant f, there is an f-DSO for undirected, unweighted graphs with unique shortest paths that has stretch $3 + \varepsilon$, space $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}}) \cdot O_\varepsilon(\log n)^{f+1}$, and query time $O_\varepsilon(n^\alpha)$.

Related result w/ different techniques: [Bilò, Choudhary, Friedrich, Krogmann & Sch. WADS 2023$^+$]

For any $f = o(\log n/\log\log n)$, there is an f-DSO for undirected, weighted graphs that has stretch $2k-1$, space $O(n^{1+\frac{1}{k}+\alpha+o(1)})$, and query time $O(n^{1+\frac{1}{k}-\frac{\alpha}{k(f+1)}})$.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides $\mathrm{poly}(\log n, 1/\varepsilon)$ factors.

Based on the f-DSO by Chechik, Cohen, Fiat & Kaplan. [SODA 2017]

- Stretch $1+\varepsilon$, space $O(n^2) \cdot O_\varepsilon(\log n)^{f+1}$, query time $\widetilde{O}(1)$.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Based on the f-DSO by Chechik, Cohen, Fiat & Kaplan. [SODA 2017]

- Stretch $1+\varepsilon$, space $O(n^2) \cdot O_\varepsilon(\log n)^{f+1}$, query time $\widetilde{O}(1)$.
- One *fault-tolerant tree* (FT-tree) for every pair of vertices.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Based on the f-DSO by Chechik, Cohen, Fiat & Kaplan. [SODA 2017]

- Stretch $1+\varepsilon$, space $O(n^2) \cdot O_\varepsilon(\log n)^{f+1}$, query time $\widetilde{O}(1)$.
- One *fault-tolerant tree* (FT-tree) for every pair of vertices.
- $FT(s,t)$ queried with set F gives exact distance $d_{G-F}(s,t)$ if the shortest s-t-path in $G-F$ satisfies some technical condition.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Based on the f-DSO by Chechik, Cohen, Fiat & Kaplan. [SODA 2017]

- Stretch $1+\varepsilon$, space $O(n^2) \cdot O_\varepsilon(\log n)^{f+1}$, query time $\widetilde{O}(1)$.
- One *fault-tolerant tree* (FT-tree) for every pair of vertices.
- $FT(s,t)$ queried with set F gives exact distance $d_{G-F}(s,t)$ if the shortest s-t-path in $G-F$ satisfies some technical condition.
- Query procedure works around the condition $\Rightarrow$ stretch $1+\varepsilon$.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Based on the f-DSO by Chechik, Cohen, Fiat & Kaplan. [SODA 2017]

- Stretch $1 + \varepsilon$, space $O(n^2) \cdot O_\varepsilon(\log n)^{f+1}$, query time $\widetilde{O}(1)$.
- One *fault-tolerant tree* (FT-tree) for every pair of vertices.
- $FT(s,t)$ queried with set F gives exact distance $d_{G-F}(s,t)$ if the shortest s-t-path in $G - F$ satisfies some technical condition.
- Query procedure works around the condition $\Rightarrow$ stretch $1 + \varepsilon$.

Cannot afford $\Omega(n^2)$ FT-trees.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides $\text{poly}(\log n, 1/\varepsilon)$ factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides $\text{poly}(\log n, 1/\varepsilon)$ factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$, the shortest s-t-path in $G - F$ is hit by some sampled vertex.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$, the shortest s-t-path in $G - F$ is hit by some sampled vertex.
- Build FT-trees only for vertex pairs in $B \times V \Rightarrow \widetilde{O}(n^2/L)$ trees.
  - Implement FT-trees in subquadratic space.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$, the shortest s-t-path in $G - F$ is hit by some sampled vertex.
- Build FT-trees only for vertex pairs in $B \times V \Rightarrow \widetilde{O}(n^2/L)$ trees.
  - Implement FT-trees in subquadratic space.
- Need a solution for small distances $d_{G-F}(s,t) \leq L$.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly$(\log n, 1/\varepsilon)$ factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$, the shortest s-t-path in $G - F$ is hit by some sampled vertex.
- Build FT-trees only for vertex pairs in $B \times V \Rightarrow \widetilde{O}(n^2/L)$ trees.
  - Implement FT-trees in subquadratic space.
- Need a solution for small distances $d_{G-F}(s,t) \leq L$.
- Naively stitching together small paths $\Rightarrow O_\varepsilon(n^{1+o(1)}/L)$ query time.
  - Improve this to $O_\varepsilon(L^{f-1+o(1)})$.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$,
  the shortest s-t-path in $G - F$ is hit by some sampled vertex.
- Build FT-trees only for vertex pairs in $B \times V \Rightarrow \widetilde{O}(n^2/L)$ trees.
  - Implement FT-trees in subquadratic space.
- Need a solution for small distances $d_{G-F}(s,t) \leq L$.
- Naively stitching together small paths $\Rightarrow O_\varepsilon(n^{1+o(1)}/L)$ query time.
  - Improve this to $O_\varepsilon(L^{f-1+o(1)})$.

Setting $L = O_\varepsilon(n^{\frac{\alpha}{f+1}})$ gives $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}})$ trees and $O_\varepsilon(n^\alpha)$ query time.

Technical Overview

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly$(\log n, 1/\varepsilon)$ factors.

Cannot afford $\Omega(n^2)$ FT-trees.

- Integer parameter $L = \omega(\log n)$: uniform set B of $\widetilde{O}(n/L)$ vertices.
- W.h.p. if $d_{G-F}(s,t) > L$ for any $F \subseteq E$ with $|F| \leq f$,
  the shortest s-t-path in $G - F$ is hit by some sampled vertex.
- Build FT-trees only for vertex pairs in $B \times V \Rightarrow \widetilde{O}(n^2/L)$ trees.
  - Implement FT-trees in subquadratic space.
- Need a solution for small distances $d_{G-F}(s,t) \leq L$.
- Naively stitching together small paths $\Rightarrow O_\varepsilon(n^{1+o(1)}/L)$ query time.
  - Improve this to $O_\varepsilon(L^{f-1+o(1)})$.

Setting $L = O_\varepsilon(n^{\frac{\alpha}{f+1}})$ gives $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}})$ trees and $O_\varepsilon(n^\alpha)$ query time.

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

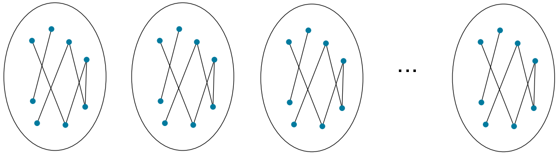

Exact distances: Weimann & Yuster [TALG 2013] graphs.

- Create $\widetilde{O}(L^f)$ copies of G,

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

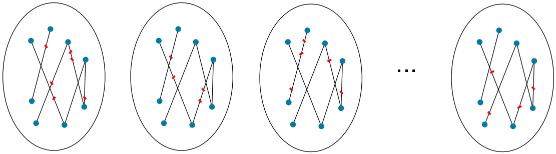

Exact distances: Weimann & Yuster [TALG 2013] graphs.

- Create $\widetilde{O}(L^f)$ copies of G, in each one fail any edge with probability $1/L$.

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

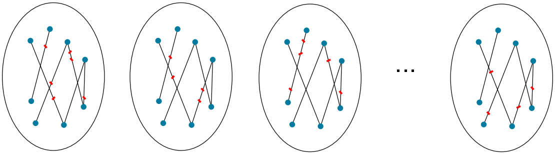

Exact distances: Weimann & Yuster [TALG 2013] graphs.

- Create $\widetilde{O}(L^f)$ copies of G, in each one fail any edge with probability $1/L$.
- W.h.p. if $d_{G-F}(s,t) \leq L$, a G_i has $E(G_i) \cap F = \emptyset$ and $d_{G_i}(s,t) = d_{G-F}(s,t)$.

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

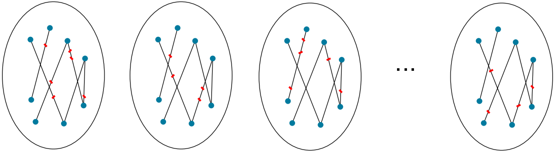

Exact distances: Weimann & Yuster [TALG 2013] graphs.

- Create $\widetilde{O}(L^f)$ copies of G, in each one fail any edge with probability $1/L$.
- W.h.p. if $d_{G-F}(s,t) \leq L$, a G_i has $E(G_i) \cap F = \emptyset$ and $d_{G_i}(s,t) = d_{G-F}(s,t)$.

Problems.

- How to find the copy G_i efficiently?

Short Distances

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

Task: f-DSO whenever real distance is $d_{G-F}(s,t) \leq L$
(arbitrary answer if $d_{G-F}(s,t) > L$).

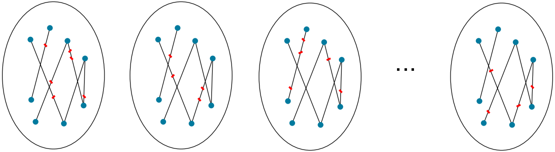

Exact distances: Weimann & Yuster [TALG 2013] graphs.

- Create $\widetilde{O}(L^f)$ copies of G, in each one fail any edge with probability $1/L$.
- W.h.p. if $d_{G-F}(s,t) \leq L$, a G_i has $E(G_i) \cap F = \emptyset$ and $d_{G_i}(s,t) = d_{G-F}(s,t)$.

Problems.

- How to find the copy G_i efficiently?
- The copies take $\Omega(L^{f-1} m)$ space.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

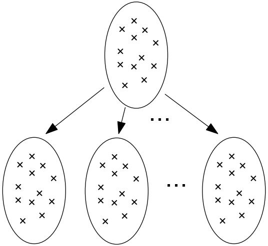

Height parameter h.

- Every inner node has $L^{f/h}$ children;

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly$(\log n, 1/\varepsilon)$ factors.

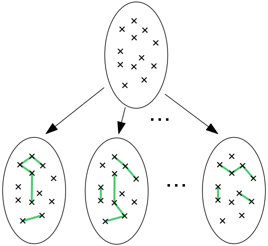

Height parameter h.

- Every inner node has $L^{f/h}$ children;
- Take edges of parent, re-insert any edge of G w/prob $1 - L^{-1/h}$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

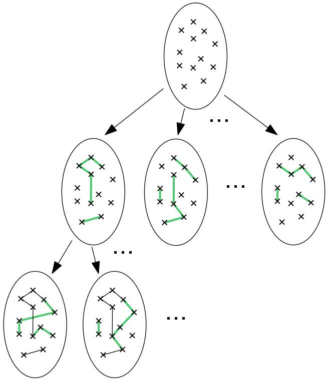

Height parameter h.

- Every inner node has $L^{f/h}$ children;
- Take edges of parent, re-insert any edge of G w/prob $1 - L^{-1/h}$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly$(\log n, 1/\varepsilon)$ factors.

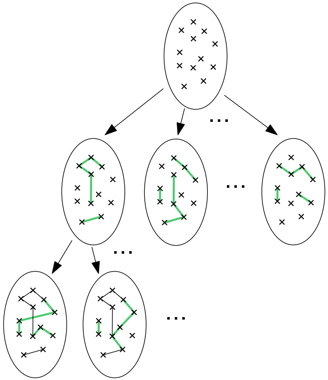

Height parameter h.

- Every inner node has $L^{f/h}$ children; $h+1$ levels $\Rightarrow O(L^{f+(f/h)})$ nodes.
- Take edges of parent, re-insert any edge of G w/prob $1 - L^{-1/h}$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

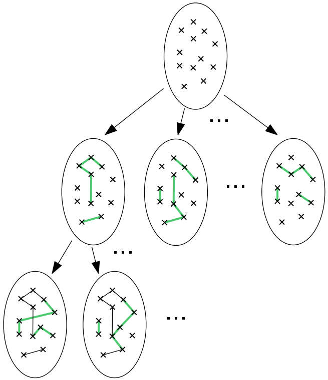

Height parameter h.

- Every inner node has $L^{f/h}$ children; $h+1$ levels $\Rightarrow O(L^{f+(f/h)})$ nodes.
- Take edges of parent, re-insert any edge of G w/prob $1 - L^{-1/h}$.
- Leaves: same probabilities as Weimann & Yuster, no independence.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides $\mathrm{poly}(\log n, 1/\varepsilon)$ factors.

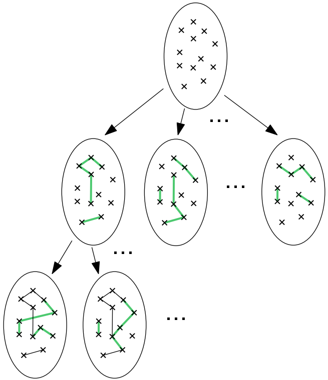

Height parameter h.

- Queried with F: follow any child that has no edge of F $\Rightarrow$ time $O(h L^{f/h})$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides $\mathrm{poly}(\log n, 1/\varepsilon)$ factors.

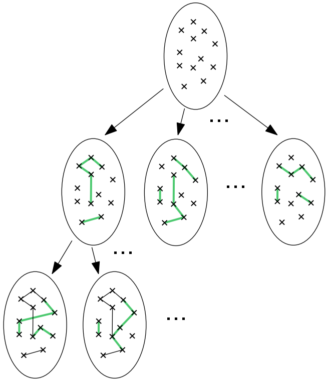

Height parameter h.

- Queried with F: follow any child that has no edge of F $\Rightarrow$ time $O(h L^{f/h})$.
- Reaching a leaf: correct copy with probability $1/c^h$ for $c > 0$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

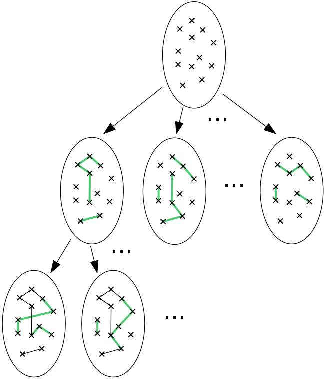

Height parameter h.

- Queried with F: follow any child that has no edge of F $\Rightarrow$ time $O(h L^{f/h})$.
- Reaching a leaf: correct copy with probability $1/c^h$ for $c > 0$.
- Repeat in $\widetilde{O}(c^h)$ independent trees for high probability.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides $\mathrm{poly}(\log n, 1/\varepsilon)$ factors.

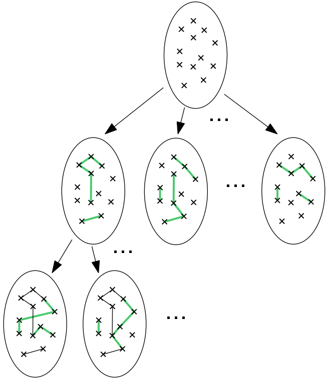

Set $h = \sqrt{f \ln L}$.

- $\widetilde{O}(L^{o(1)})$ trees, each with $L^{f+o(1)}$ nodes.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

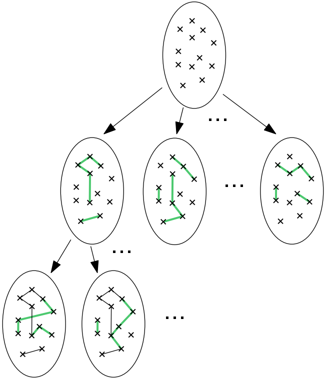

Set $h = \sqrt{f \ln L}$.

- $\widetilde{O}(L^{o(1)})$ trees, each with $L^{f+o(1)}$ nodes.
- Query time is $\widetilde{O}(L^{o(1)})$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

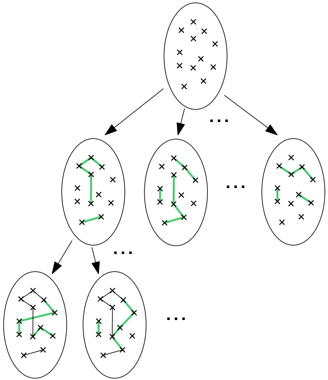

Set $h = \sqrt{f \ln L}$.

- $\widetilde{O}(L^{o(1)})$ trees, each with $L^{f+o(1)}$ nodes.
- Query time is $\widetilde{O}(L^{o(1)})$.
- Still one leaf alone has $\Omega(m/L)$ edges.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_{\varepsilon}(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

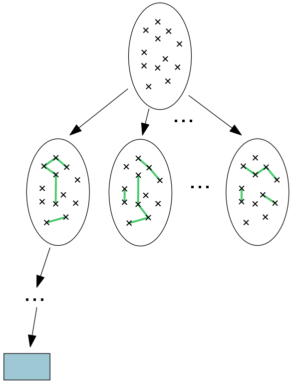

- Replace the leaves with Thorup & Zwick [JACM 2005] distance oracles
  $\Rightarrow$ stretch $2k-1$, size $O(k\, n^{1+1/k})$, query $O(k)$.

Hierarchical Tree Sampling

Constant sensitivity f, space in $O(\log n)$-bit words, undirected unweighted graphs, unique shortest paths, $\widetilde{O}_\varepsilon(\cdot)$ hides poly($\log n$, $1/\varepsilon$) factors.

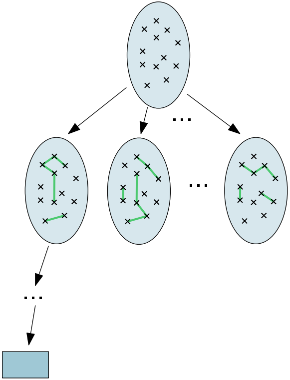

- Replace the leaves with Thorup & Zwick [JACM 2005] distance oracles
  $\Rightarrow$ stretch $2k-1$, size $O(k\,n^{1+1/k})$, query $O(k)$.
- Replace inner nodes with spanners
  $\Rightarrow$ total space of all sampling trees $\widetilde{O}(L^{f+o(1)} n^{1+1/k})$.

Summary

- Distance sensitivity oracle for constant f with stretch $3 + \varepsilon$, space $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}}) \cdot O_\varepsilon(\log n)^{f+1}$, and query time $O_\varepsilon(n^\alpha)$.
- Distance sensitivity oracle for distances $\leq L$ with stretch $2k - 1$, space $\widetilde{O}(L^{f+o(1)} n^{1+1/k})$, and query time $\widetilde{O}(L^{o(1)})$.

Summary

- Distance sensitivity oracle for constant f with stretch $3+\varepsilon$, space $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}}) \cdot O_\varepsilon(\log n)^{f+1}$, and query time $O_\varepsilon(n^\alpha)$.

- Distance sensitivity oracle for distances $\leq L$ with stretch $2k-1$, space $\widetilde{O}(L^{f+o(1)} n^{1+1/k})$, and query time $\widetilde{O}(L^{o(1)})$.

Open Problems

- Reduce the query time to $\text{poly}(\log n, 1/\varepsilon)$.

- Extend solution to weighted graphs.

Summary

- Distance sensitivity oracle for constant f with stretch $3 + \varepsilon$, space $\widetilde{O}_\varepsilon(n^{2-\frac{\alpha}{f+1}}) \cdot O_\varepsilon(\log n)^{f+1}$, and query time $O_\varepsilon(n^\alpha)$.
- Distance sensitivity oracle for distances $\leq L$ with stretch $2k - 1$, space $\widetilde{O}(L^{f+o(1)} n^{1+1/k})$, and query time $\widetilde{O}(L^{o(1)})$.

Open Problems

- Reduce the query time to $\text{poly}(\log n, 1/\varepsilon)$.
- Extend solution to weighted graphs.

Thanks!