Compact representation of distances in a graph : a tour around 2-hop labelings

Laurent Viennot (Univ. Paris – Inria – Irif)

Joint work with Siddharth Gupta (Univ. of Warwick),
Adrian Kosowski (NavAlgo)
and Przemysław Uznański (NavAlgo)

The beginning of the story

SODA 2017 :

- Chepoi, Dragan, Vaxès : Core congestion is inherent in hyperbolic networks.
- Kosowski, V. : Beyond highway dimension : small distance labels using tree skeletons.

The beginning of the story

SODA 2017 :

- Chepoi, Dragan, Vaxès : Core congestion is inherent in hyperbolic networks.
- Kosowski, V. : Beyond highway dimension : small distance labels using tree skeletons.

The beginning of the story

SODA 2017 :

- Chepoi, Dragan, Vaxès : Core congestion is inherent in hyperbolic networks.
- Kosowski, V. : Beyond highway dimension : small distance labels using tree skeletons.

Encoding distances in a graph

We are given a (weighted) (di-) graph $G = (V, E)$ with n nodes and m edges.

Make any useful pre-computation to answer efficiently online distance queries : what is distance $d(u_1, v_1)$?, $d(u_2, v_2)$?,$d(u_3, v_3)$?,...

Encoding distances in a graph

We are given a (weighted) (di-) graph $G = (V, E)$ with n nodes and m edges.

Make any useful pre-computation to answer efficiently online distance queries : what is distance $d(u_1, v_1)$?, $d(u_2, v_2)$?,$d(u_3, v_3)$?,...

Encoding a graph metric : distance oracles

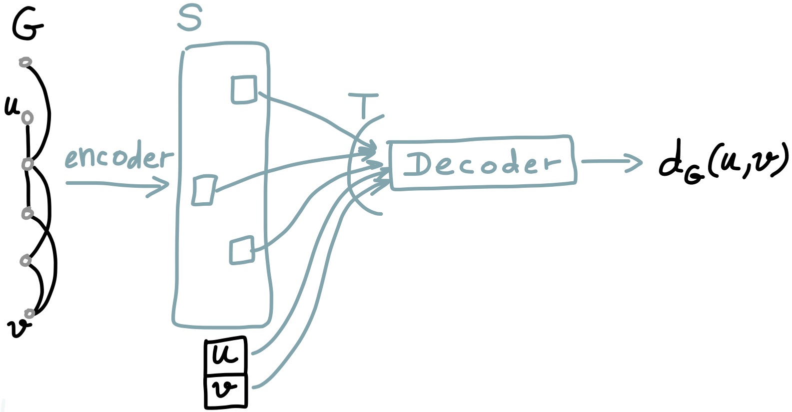

Encoding a graph metric : distance labelings

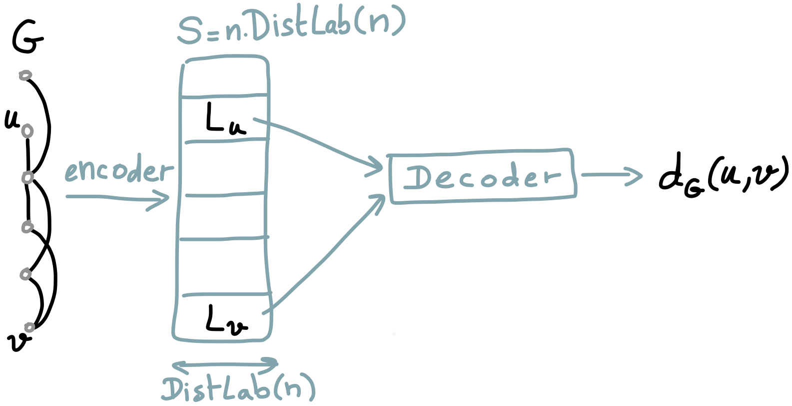

Encoding a graph metric : 2-hop labelings

A 2-hop labelings is a very simple kind of distance labeling.

The main idea is to associate a set $H_u \subseteq V$ of "hubs" to each node u and to store the distances $d(u, v)$ for all $v \in H_u$.

Also known as hub labeling, or landmark labeling.

Encoding a graph metric : 2-hop labelings

A 2-hop labelings is a very simple kind of distance labeling.

The main idea is to associate a set $H_u \subseteq V$ of "hubs" to each node u and to store the distances $d(u, v)$ for all $v \in H_u$.

Also known as hub labeling, or landmark labeling.

Encoding a graph metric : 2-hop labelings

A 2-hop labelings is a very simple kind of distance labeling.

The main idea is to associate a set $H_u \subseteq V$ of "hubs" to each node u and to store the distances $d(u, v)$ for all $v \in H_u$.

Also known as hub labeling, or landmark labeling.

2-hop labeling

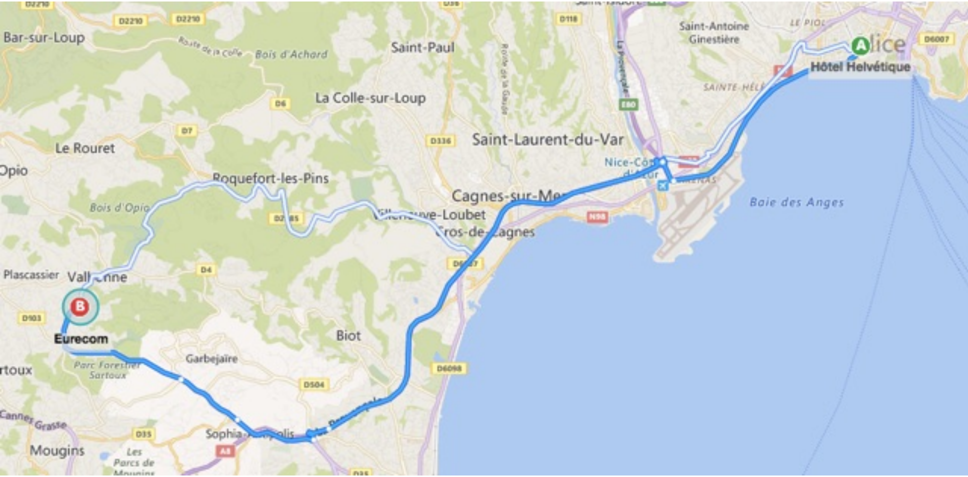

2-hop labeling

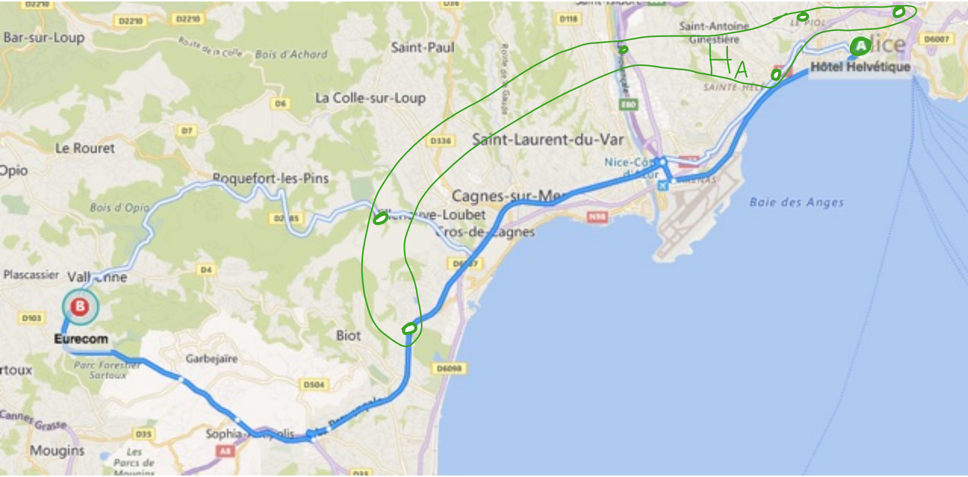

2-hop labeling

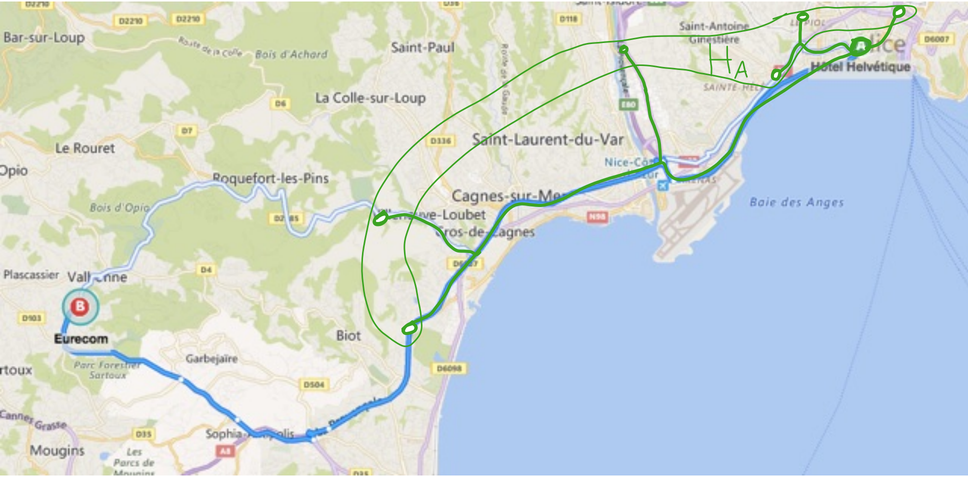

2-hop labeling

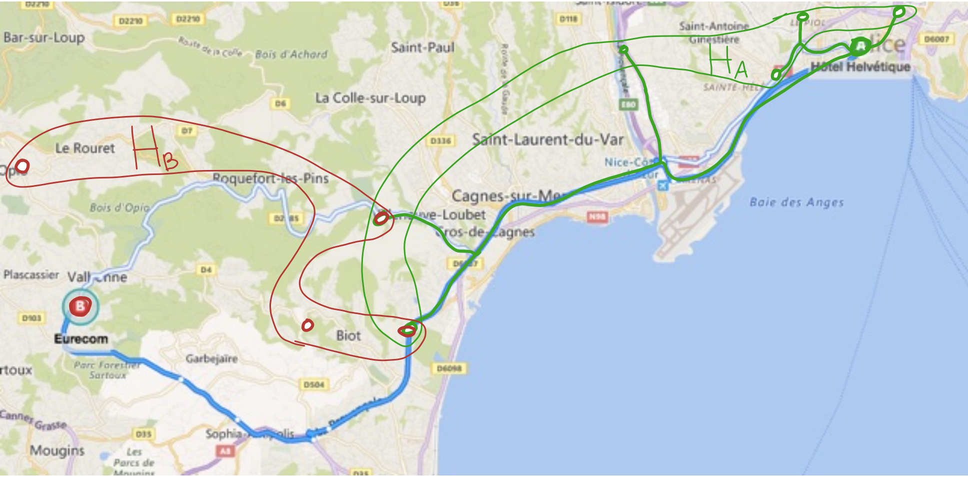

2-hop labeling

2-hop labeling

Hub sets

Covering property :
A collection of hub sets $H_u \subseteq V$ for all $u \in V$ is said to cover graph G if for all u, v, there exists $w \in H_u \cap H_v$ with $d(u,v) = d(u,w) + d(w,v)$.

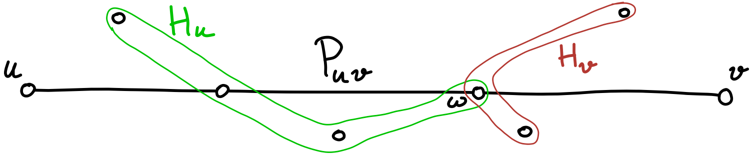

Distance labels : $L_u = \{(w, d(u,w)) : w \in H_u\}$

Distance query : Dist $(L_u, L_v) = \min_{w \in H_u \cap H_v} d(u,w) + d(w,v)$

Introduced by [Gavoille et al. '04; Cohen et al. 2003], applied to road networks [Abraham et al. 2010-2013], and other practical networks [Akiba et al. 2013]. Approximability results : [Babenko et al. 2013, Angelidakis et al. 2017].

Hub sets

Covering property :
A collection of hub sets $H_u \subseteq V$ for all $u \in V$ is said to cover graph G if for all u, v, there exists $w \in H_u \cap H_v$ with $d(u, v) = d(u, w) + d(w, v)$.

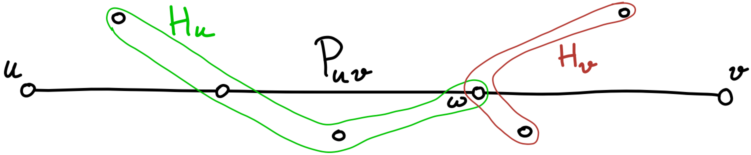

Distance labels : $L_u = \{(w, d(u, w)) : w \in H_u\}$

Distance query : Dist $(L_u, L_v) = \min_{w \in H_u \cap H_v} d(u, w) + d(w, v)$

Introduced by [Gavoille et al. '04; Cohen et al. 2003], applied to road networks [Abraham et al. 2010-2013], and other practical networks [Akiba et al. 2013]. Approximability results : [Babenko et al. 2013, Angelidakis et al. 2017].

Hub sets

Covering property :
A collection of hub sets $H_u \subseteq V$ for all $u \in V$ is said to cover graph G if for all u, v, there exists $w \in H_u \cap H_v$ with $d(u,v) = d(u,w) + d(w,v)$ (equivalently, $H_u \cap H_v \cap I(u,v) \neq \emptyset$).

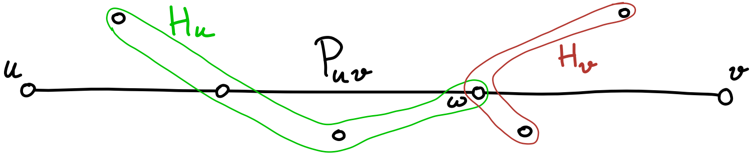

Equivalently, for all u, v, $H_u \cap H_v \cap I(u,v) \neq \emptyset$ where the interval $I(u,v)$ is the union of shortest paths from u to v.

Note that using Sergio Cabello framework :
$V \times V = \bigcup_{w \in V} H_w^{-1} \times H_w^{-1}$ where $H_w = \{u : w \in H_u\}$.

Hub sets

Covering property :
A collection of hub sets $H_u \subseteq V$ for all $u \in V$ is said to cover graph G if for all u, v, there exists $w \in H_u \cap H_v$ with $d(u,v) = d(u,w) + d(w,v)$ (equivalently, $H_u \cap H_v \cap I(u,v) \neq \emptyset$).

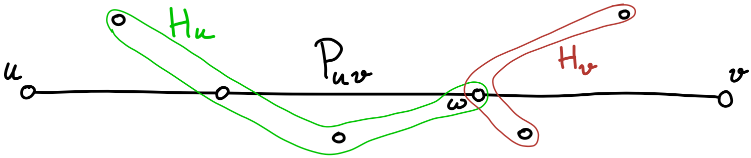

Equivalently, for all u, v, $H_u \cap H_v \cap I(u,v) \neq \emptyset$ where the interval $I(u,v)$ is the union of shortest paths from u to v.

Note that using Sergio Cabello framework :
$V \times V = \bigcup_{w \in V} H_w^{-1} \times H_w^{-1}$ where $H_w = \{u : w \in H_u\}$.

Hub sets covering a path

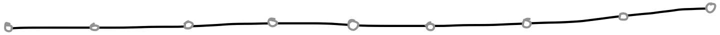

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

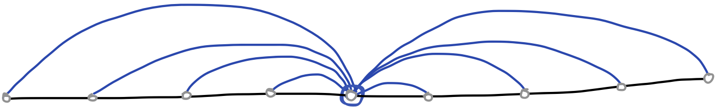

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

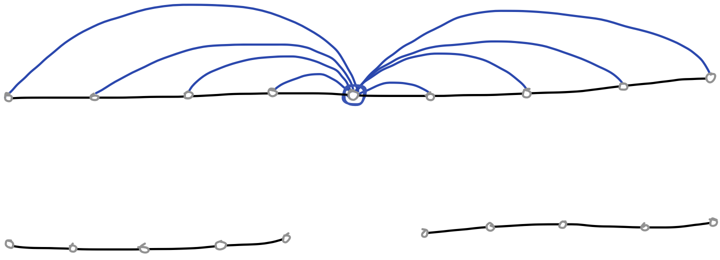

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

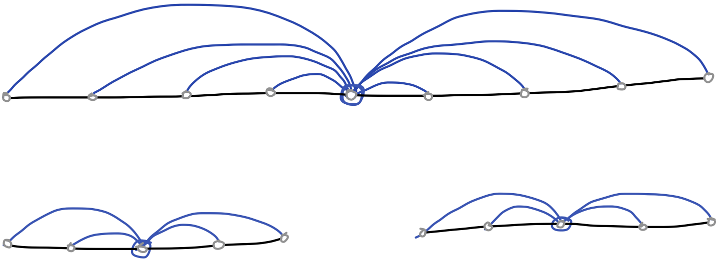

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

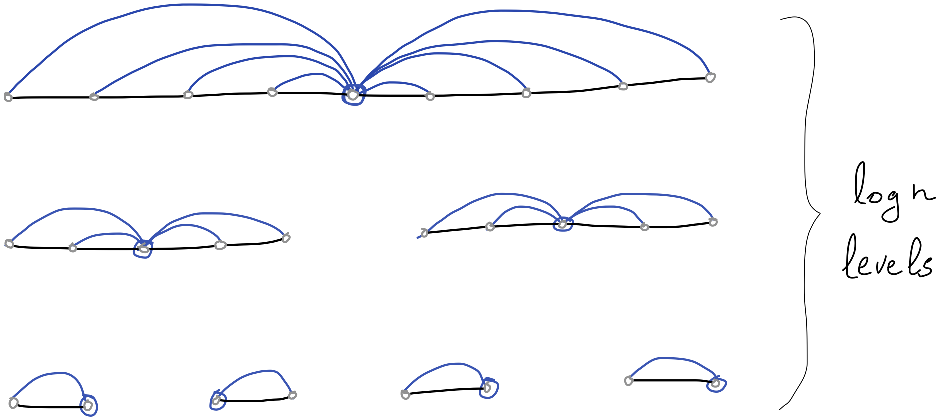

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

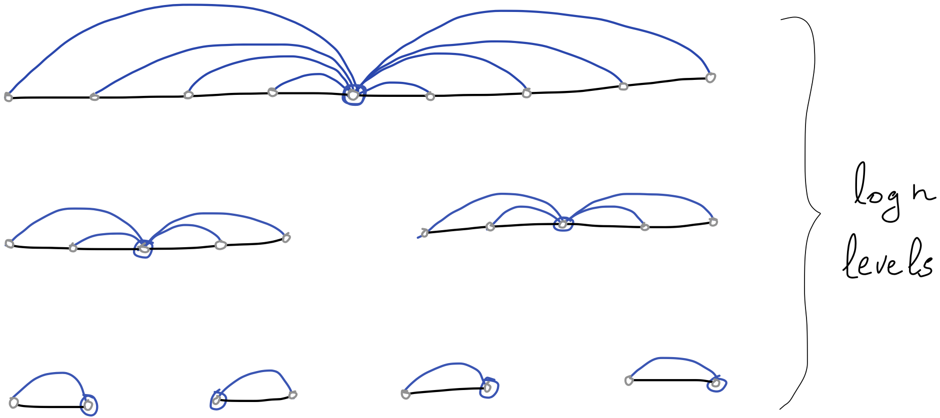

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

Hub sets covering a path

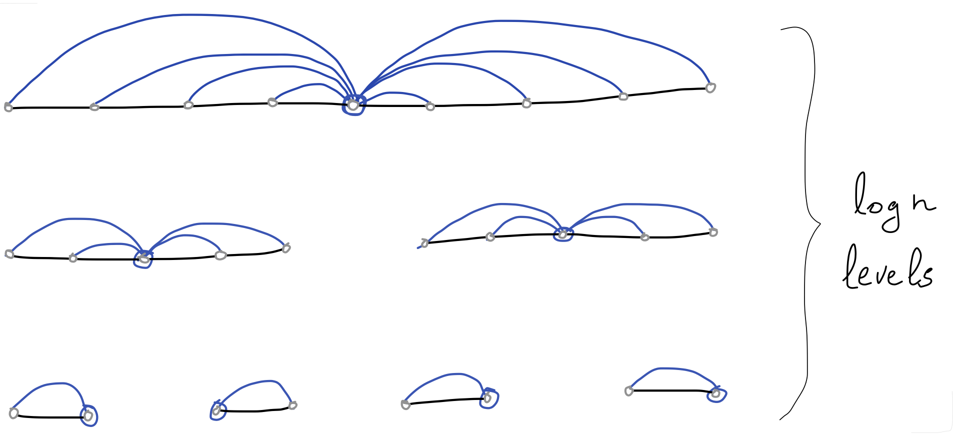

This results in covering hub sets of size $O(\log n)$.

A similar construction works for trees and bounded-treewidth graphs.

This talk is about

What graphs do have small hubsets?

No hope for dense graphs :

- average hub-set size is at least $\frac{m}{2n}$ as :
- for each edge $uv \in E$, we must have $u \in H_v$ or $v \in H_u$.

Planar graphs have covering hub sets of size $O(\sqrt{n})$, with a best known lower bound of $\Omega(n^{1/3})$ (unweighted). [Gavoille, Peleg, Pérennes, Raz '04].

This talk is about

What graphs do have small hubsets?

No hope for dense graphs :

- average hub-set size is at least $\frac{m}{2n}$ as :
- for each edge $uv \in E$, we must have $u \in H_v$ or $v \in H_u$.

Planar graphs have covering hub sets of size $O(\sqrt{n})$, with a best known lower bound of $\Omega(n^{1/3})$ (unweighted). [Gavoille, Peleg, Pérennes, Raz '04].

This talk is about

What graphs do have small hubsets?

No hope for dense graphs :

- average hub-set size is at least $\frac{m}{2n}$ as :
- for each edge $uv \in E$, we must have $u \in H_v$ or $v \in H_u$.

Planar graphs have covering hub sets of size $O(\sqrt{n})$, with a best known lower bound of $\Omega(n^{1/3})$ (unweighted). [Gavoille, Peleg, Pérennes, Raz '04].

Part I : Do sparse graphs have covering hub sets with $o(n)$ size?

Can we have sublinear size for sparse graphs ($m = O(n)$)?

Or even constant degree graphs?

Best known upper bound is $O(\frac{n}{\log n})$. [Alstrup, Dahlgaard, Beck, Knudsen, Porah '16] [Gawrychowski, Kosowski, Uznanski '16]

Best known lower bound was $\Omega(\sqrt{n})$ [Gavoille, Peleg, Pérennes, Raz '04].

Can we have sublinear size for sparse graphs ($m = O(n)$)?

Or even constant degree graphs?

Best known upper bound is $O(\frac{n}{\log n})$. [Alstrup, Dahlgaard, Beck, Knudsen, Porah '16] [Gawrychowski, Kosowski, Uznanski '16]

Best known lower bound was $\Omega(\sqrt{n})$ [Gavoille, Peleg, Pérennes, Raz '04].

Can we have sublinear size for sparse graphs ($m = O(n)$)?

Or even constant degree graphs?

Best known upper bound is $O(\frac{n}{\log n})$. [Alstrup, Dahlgaard, Beck, Knudsen, Porah '16] [Gawrychowski, Kosowski, Uznanski '16]

Best known lower bound was $\Omega(\sqrt{n})$ [Gavoille, Peleg, Pérennes, Raz '04].

Can we have sublinear size for sparse graphs ($m = O(n)$)?

Or even constant degree graphs?

Best known upper bound is $O(\frac{n}{\log n})$. [Alstrup, Dahlgaard, Beck, Knudsen, Porah '16] [Gawrychowski, Kosowski, Uznanski '16]

Best known lower bound was $\Omega(\sqrt{n})$ [Gavoille, Peleg, Pérennes, Raz '04].

Theorem (Kosowski, Uznański, V. '19)

(1) There exists graphs of degree at most 3 where any collection of covering hub sets has average size $\frac{n}{2^{O(\sqrt{\log n})}}$.

(2) Any graph has a collection of hub sets of $O(\frac{n}{RQ(n)^{1/7}})$ size where $2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$ is a number related to Ruzsa-Szemerédi graphs.

Proof : cov. hub sets of this graph have size $\frac{n}{2^{O(\sqrt{\log n})}}$

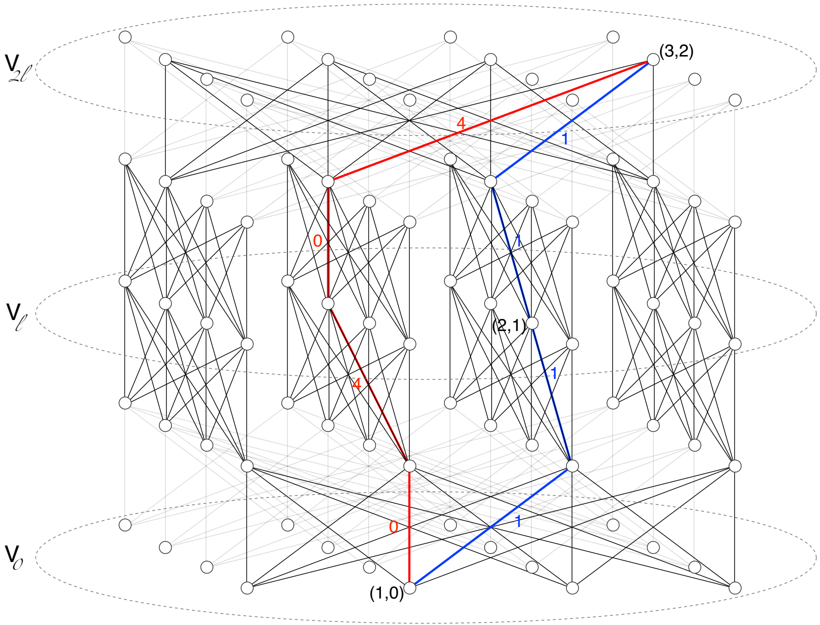

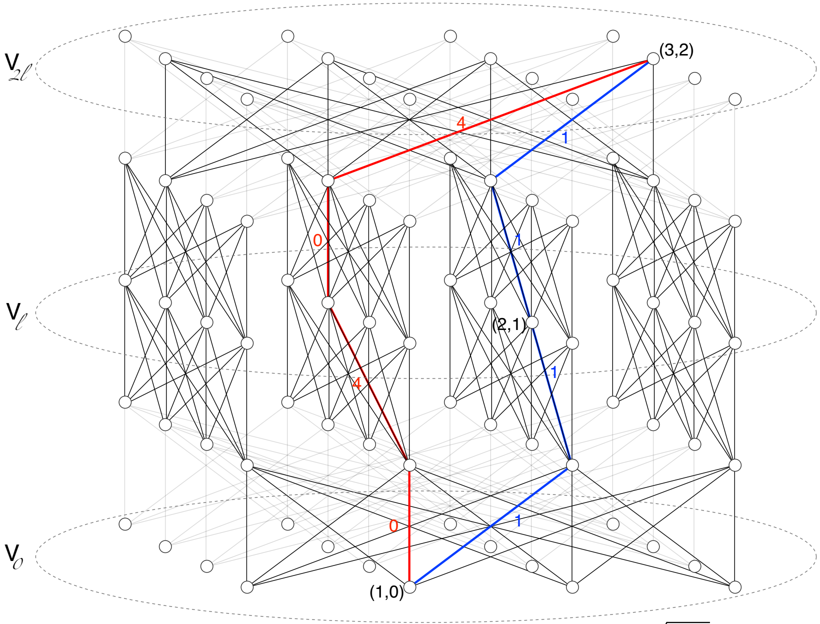

Each V_i is a regular $2\ell \times \cdots \times 2\ell$ lattice of dim. $\ell \approx 2^{\sqrt{\log n}}$ (here $\ell = 2$).
Edges from V_{i-1} to V_i connect nodes differing on ith coordinate.

Ruzsa-Szemerédi

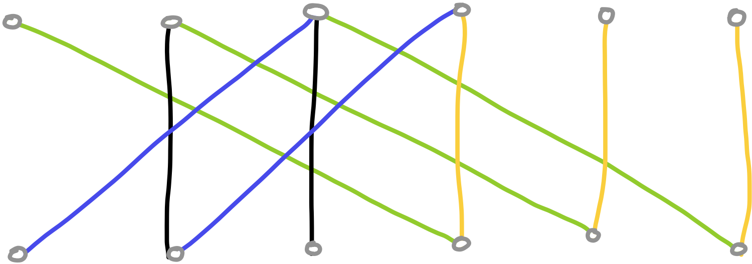

A graph is an RS-graph if it can be decomposed into n induced matchings.

Ruzsa-Szemerédi

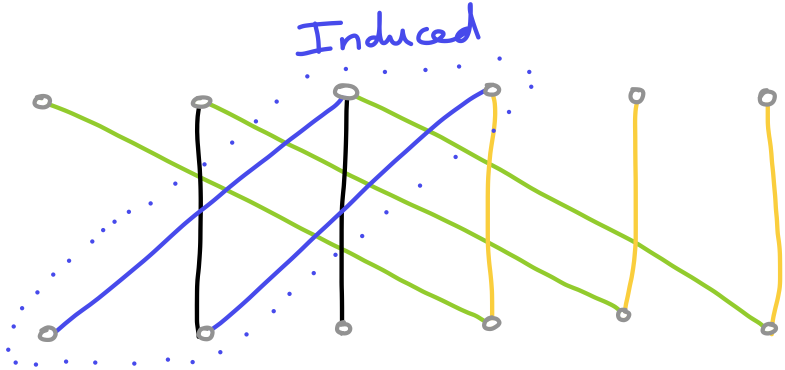

A graph is an RS-graph if it can be decomposed into n induced matchings.

Ruzsa-Szemerédi

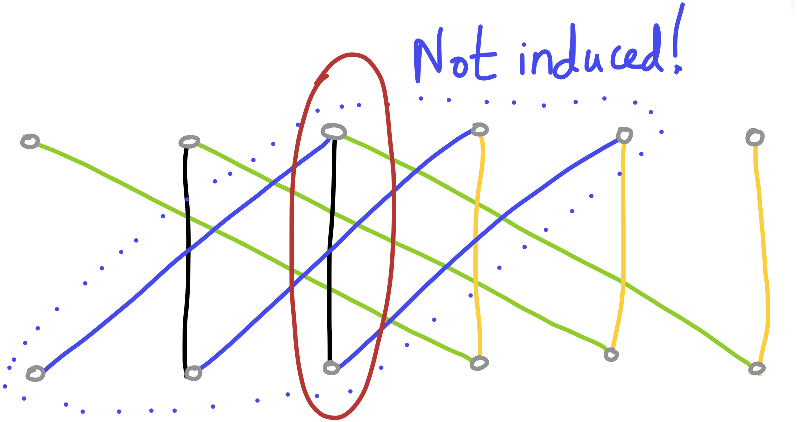

A graph is an RS-graph if it can be decomposed into n induced matchings.

Ruzsa-Szemerédi

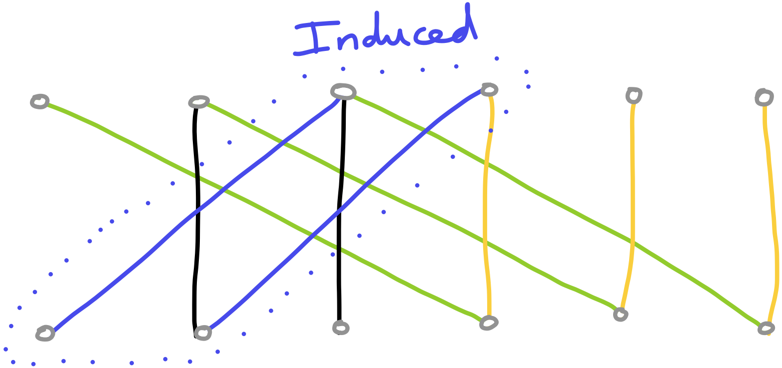

A graph is an RS-graph if it can be decomposed into n induced matchings.

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.
(Using dense subsets of $\{1, \ldots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turán '36].)

Define $RS(n)$ is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.

(Using dense subsets of $\{1, \ldots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turan '36].)

Define $RS(n)$ is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.
(Using dense subsets of $\{1, \ldots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turan '36].)

Define RS(n) is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.
(Using dense subsets of $\{1, \dots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turan '36].)

Define RS(n) is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \le RS(n) \le 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.
(Using dense subsets of $\{1, \ldots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turan '36].)

Define RS(n) is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

A graph is an RS-graph if it can be decomposed into n induced matchings.

What are the densest RS-graphs?

Theorem ([Ruzsa, Szemerédi '78])
Any RS-graph has at most $\frac{n^2}{2^{O(\log^* n)}}$ edges.
(Using dense subsets of $\{1, \ldots, n\}$ with no arithmetic triples [Behrand '46] after [Erdős and Turan '36].)

Define RS(n) is the largest integer such that there exists an RS-graph with n nodes and $\frac{n^2}{RS(n)}$ edges.

$$2^{\Omega(\log^* n)} \leq RS(n) \leq 2^{O(\sqrt{\log n})}$$

[Ruzsa, Szemerédi '78] [Elkin '10] [Fox '11]

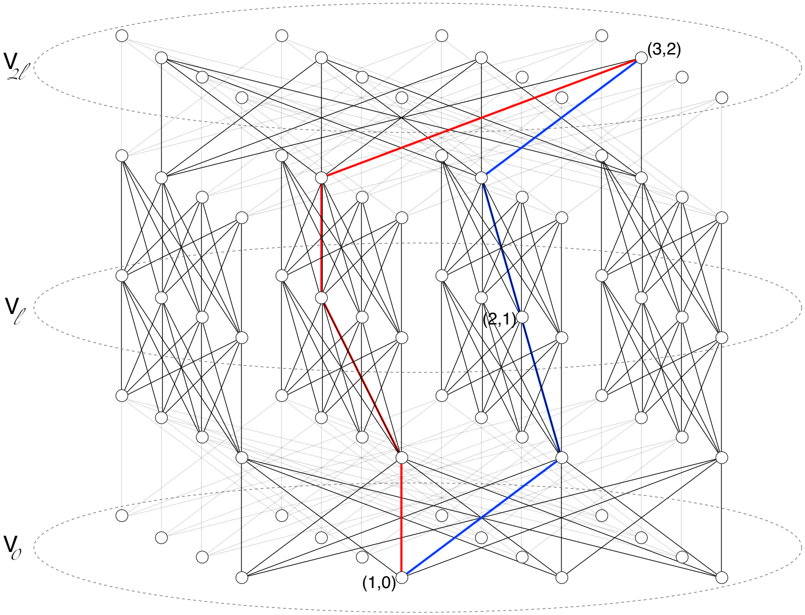

$$G_y^D = \{x_0 z_{2\ell} \mid y = \tfrac{x+z}{2} \text{ and } d_G(x,z) = D\} \quad \exists D \text{ s.t. } |\cup_y G_y^D| \geq \frac{n^2}{2^{O(\sqrt{\log n})}}$$

Converse

Any cst. deg. graph G has hub sets of av. size $O(\frac{n}{\mathsf{RS}(n)^{1/7}})$.

Idea : use a vertex cover of each G_y^D (VC ≤ 2MM).

Connection with SumIndex problem (comm. complexity)

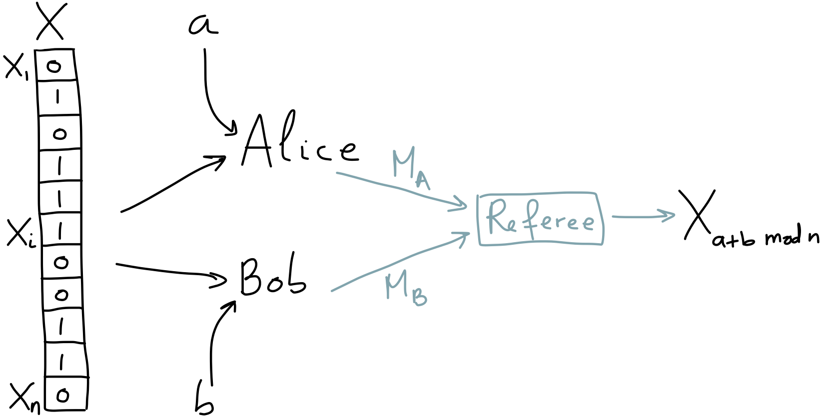

$\text{SUMINDEX}(n) = \min_{\text{Encoder}} \max_X |M_A| + |M_B|$

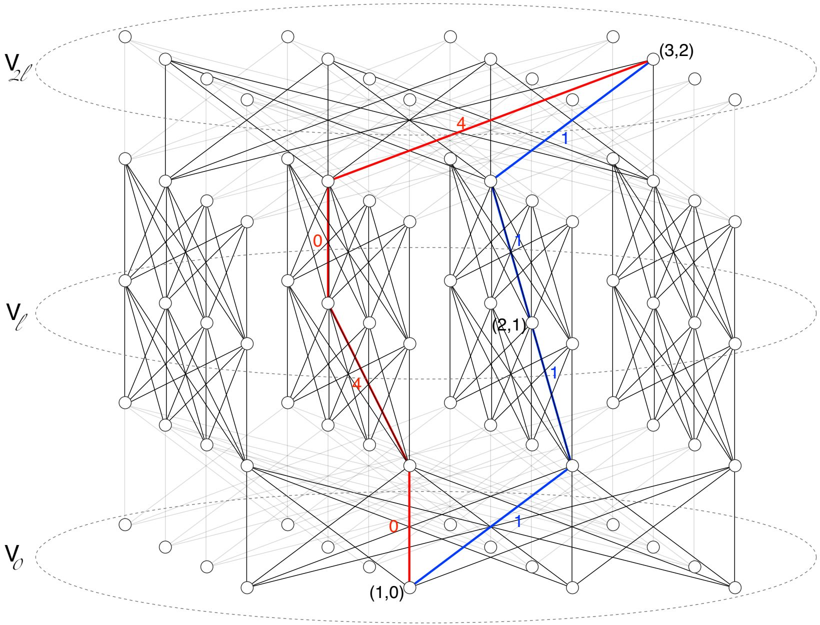

$G_X = G \setminus \{y_\ell \mid X_y = 0\}$, send $x = 2a, L_{x_0}$, $z = 2b, L_{z_{2\ell}}$, check $d(x_0, z_{2\ell})$.

Part II : what about practical graphs ?

Yes! practical graphs tend to have small covering hub sets. [Akiba et al. '13] [Delling et al. 14]

What kind of property they have enables that?

Small highway dimension. [Abraham, Fiat, Goldberg, Werneck '10-13]

More generally, small skeleton dimension. [Kosowski, V. '17]

Yes! practical graphs tend to have small covering hub sets.
[Akiba et al. '13] [Delling et al. 14]

What kind of property they have enables that?

Small highway dimension. [Abraham, Fiat, Goldberg, Werneck '10-13]

More generally, small skeleton dimension. [Kosowski, V. '17]

Yes! practical graphs tend to have small covering hub sets. [Akiba et al. '13] [Delling et al. 14]

What kind of property they have enables that?

Small highway dimension. [Abraham, Fiat, Goldberg, Werneck '10-13]

More generally, small skeleton dimension. [Kosowski, V. '17]

Yes! practical graphs tend to have small covering hub sets. [Akiba et al. '13] [Delling et al. 14]

What kind of property they have enables that?

Small highway dimension. [Abraham, Fiat, Goldberg, Werneck '10-13]

More generally, small skeleton dimension. [Kosowski, V. '17]

Skeleton dimension

The skeleton dimension k of G is the maximum “width” of a “pruned” shortest path tree.

Barcelona shortest path tree

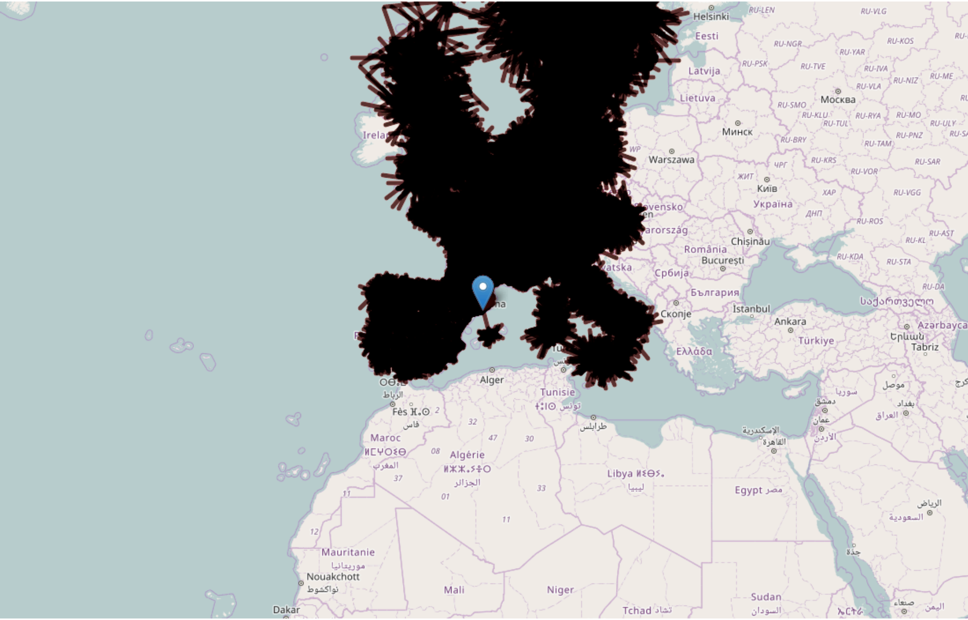

Barcelona tree skeleton : prune last third

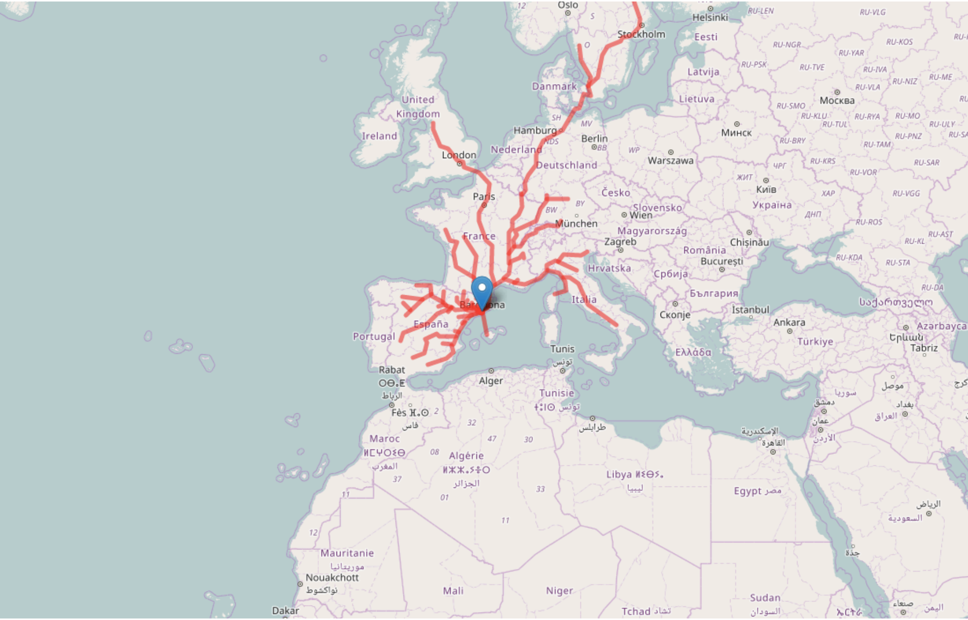

Barcelona tree skeleton : prune last third

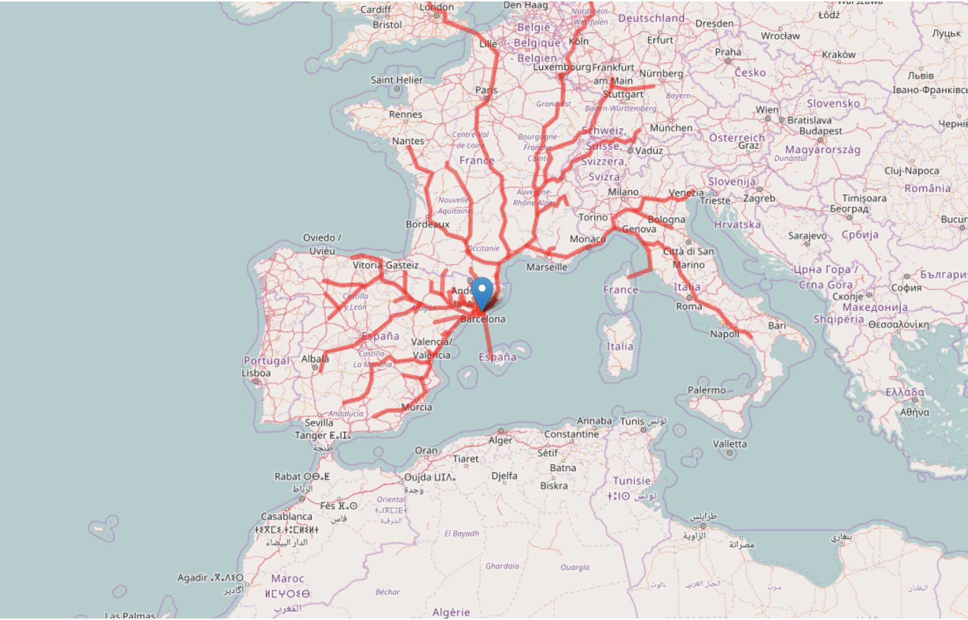

Barcelona tree skeleton : prune last third

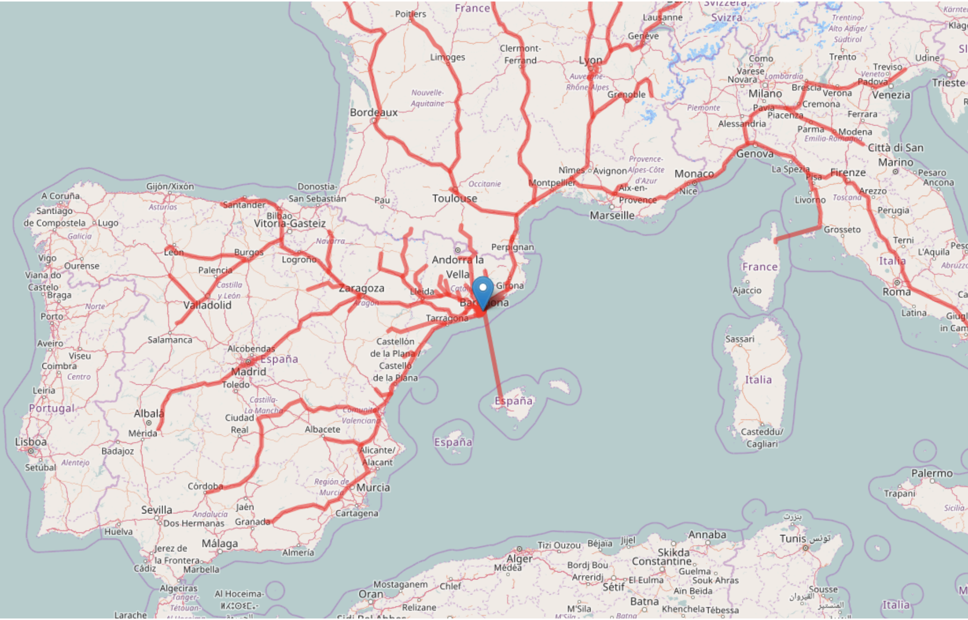

Barcelona tree skeleton : prune last third

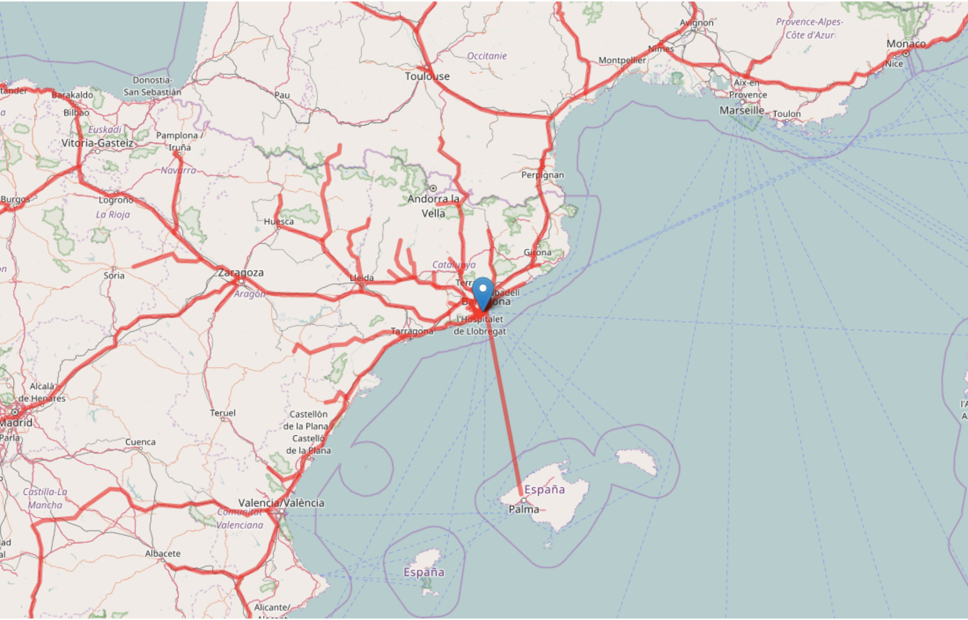

Barcelona tree skeleton : prune last third

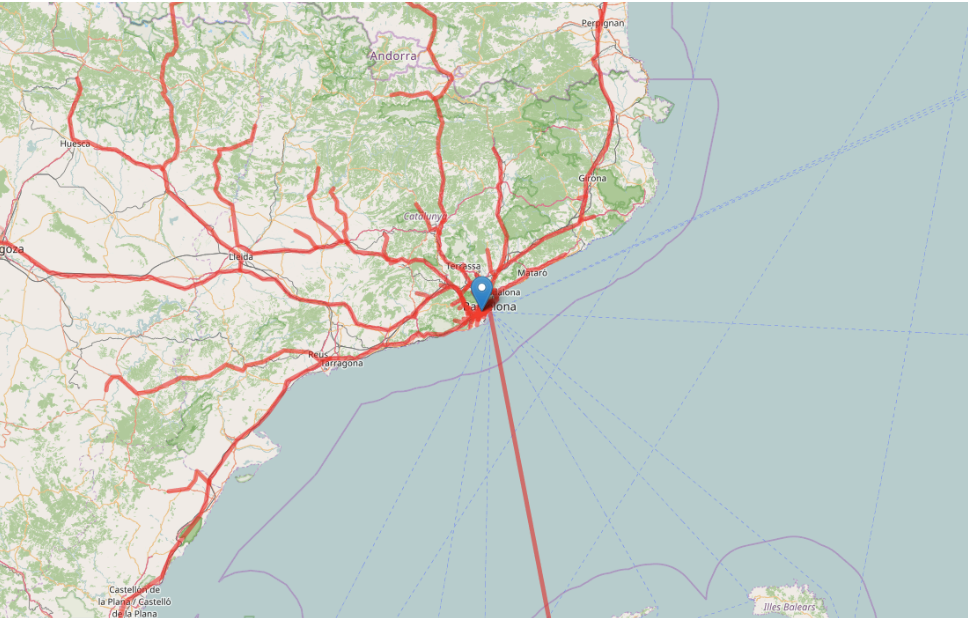

Barcelona tree skeleton : prune last third

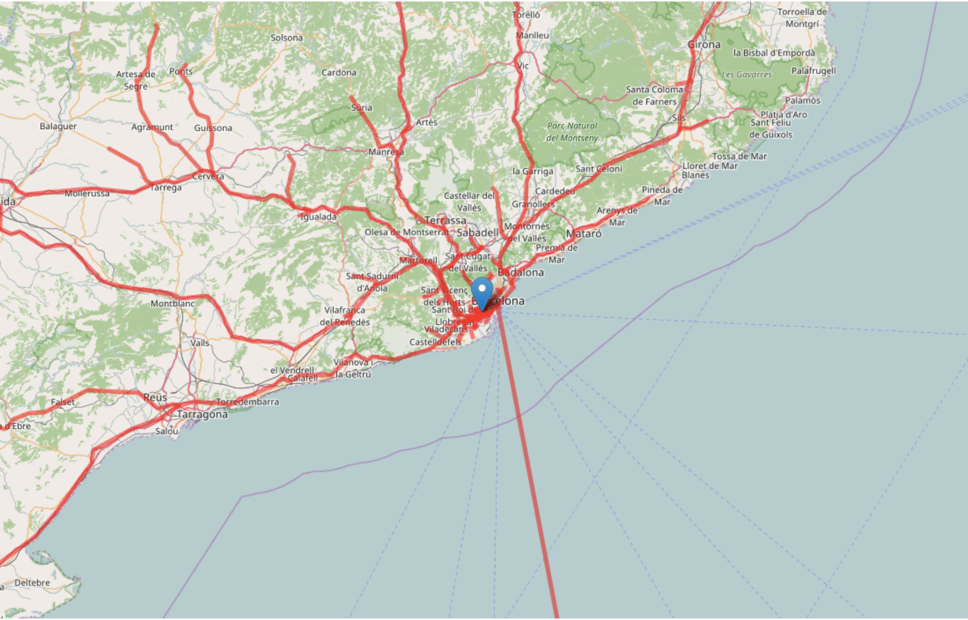

Barcelona tree skeleton : prune last third

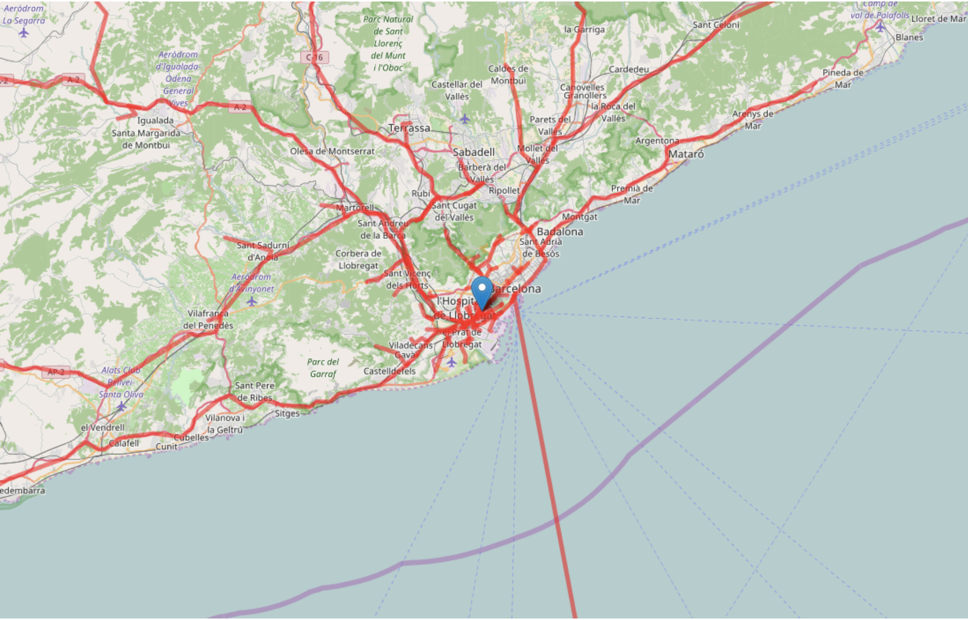

Barcelona tree skeleton : prune last third

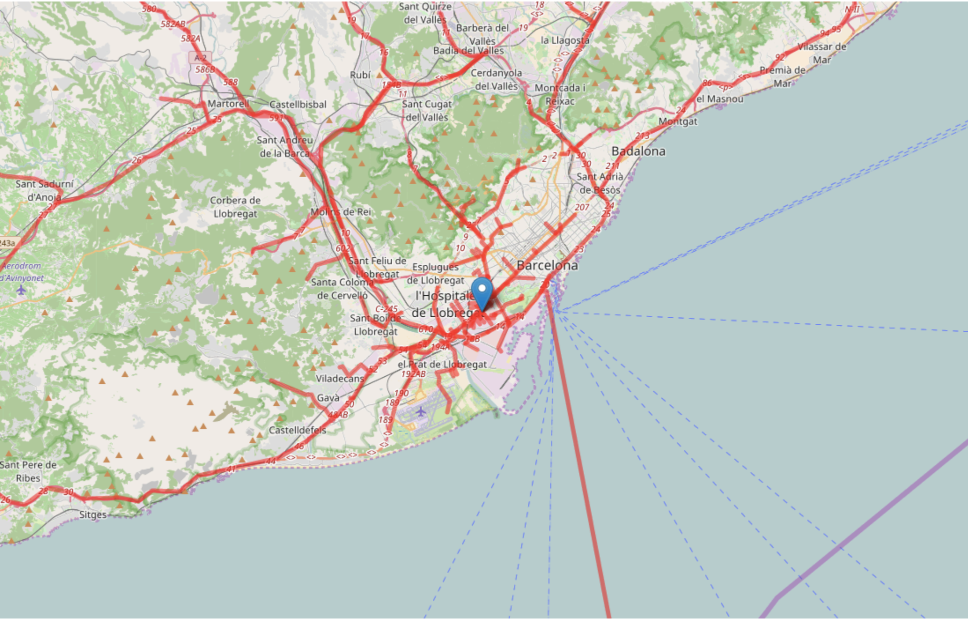

Barcelona tree skeleton : prune last third

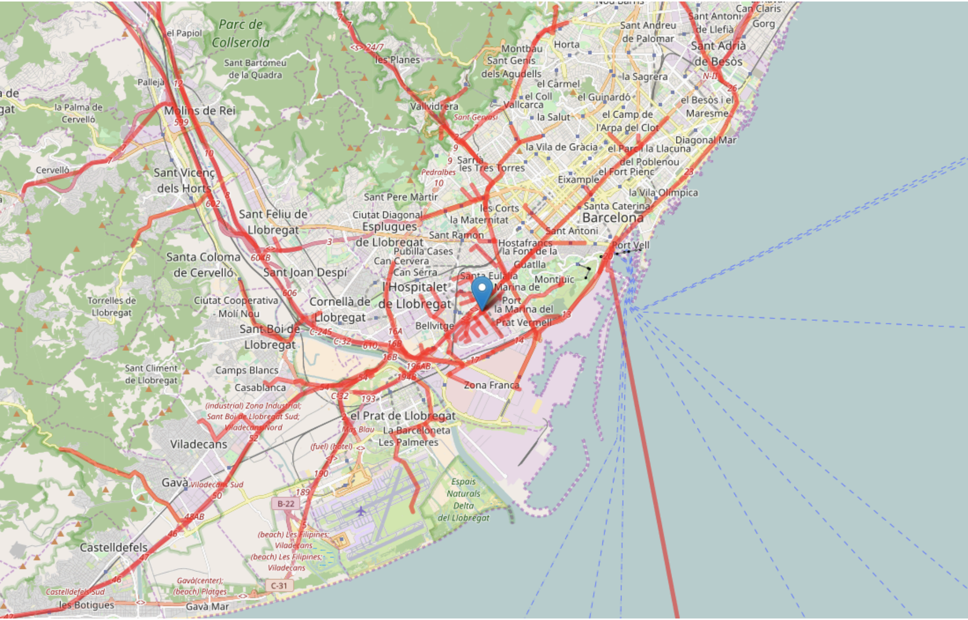

Barcelona tree skeleton : prune last third

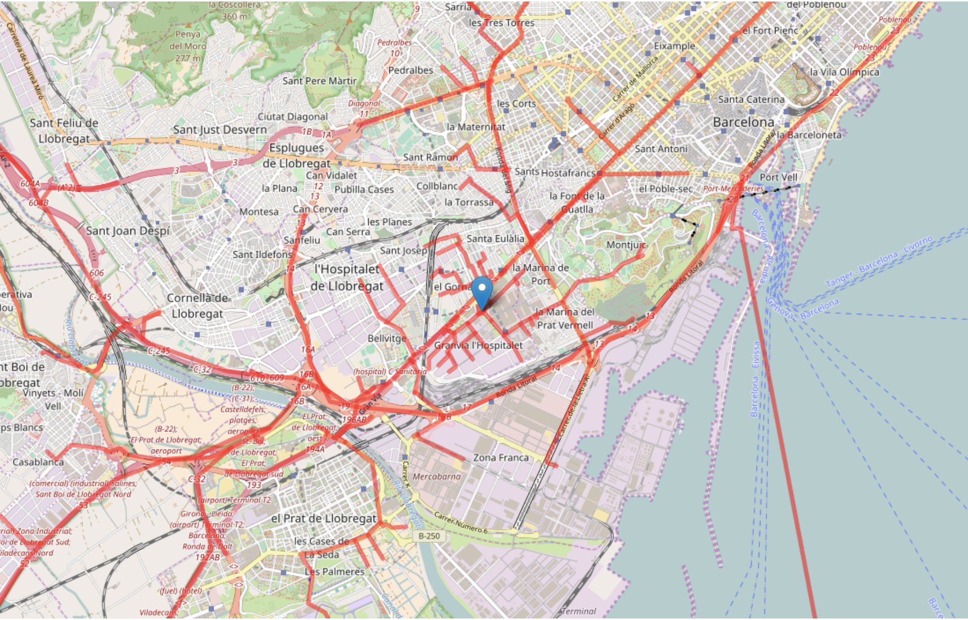

Barcelona tree skeleton : prune last third

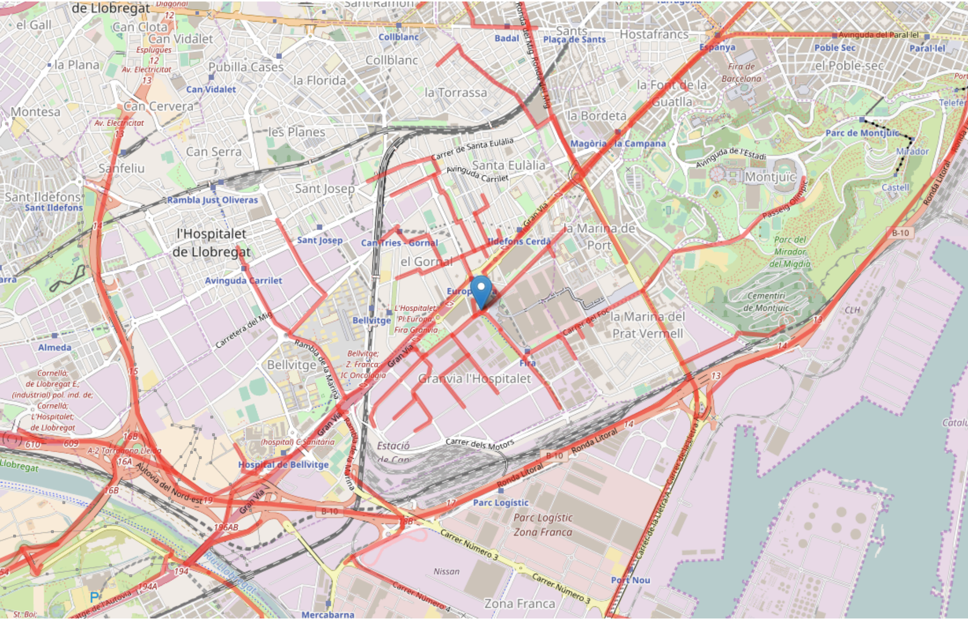

Tree skeleton

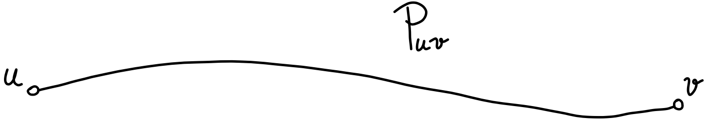

Tree skeleton

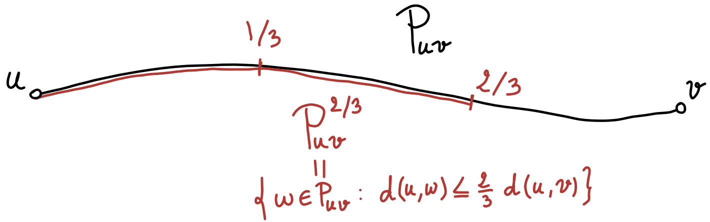

Tree skeleton

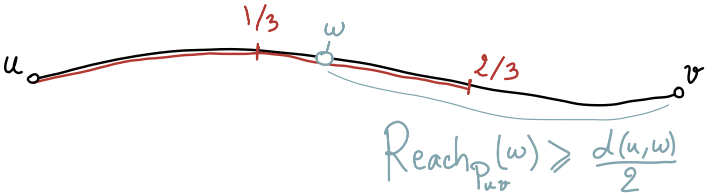

Tree skeleton

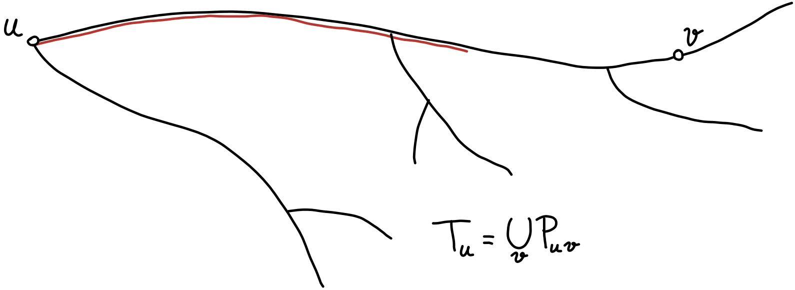

$T_u = \bigcup_v P_{uv}$

Tree skeleton

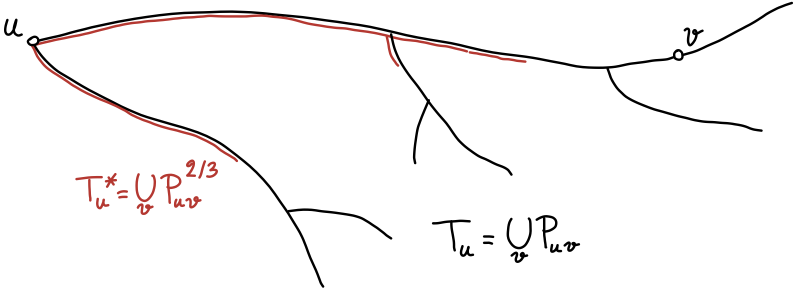

Tree skeleton

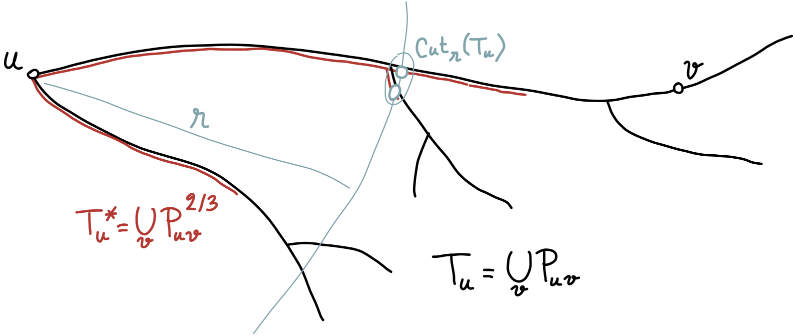

Tree skeleton

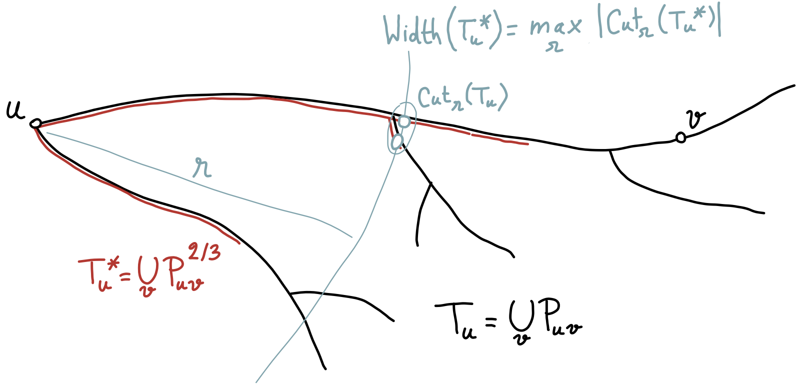

Tree skeleton

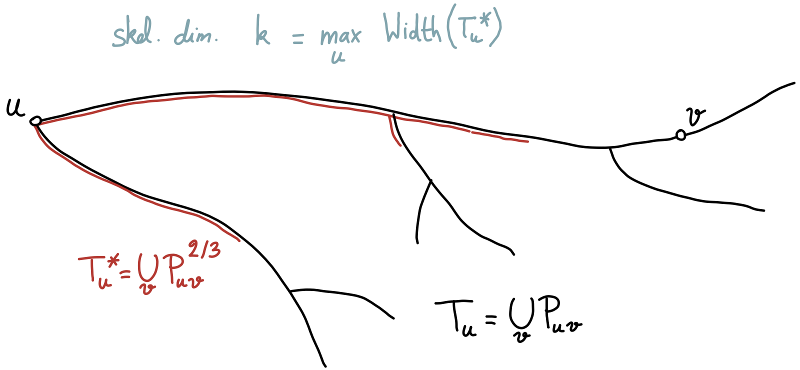

Theorem (Kosowski, V. 2017)

Given a graph G with skeleton dimension k and diameter D, a simple random sampling technique allows to find in polynomial time hub sets with size $O(k \log D)$ on average and maximum size $O(k \log \log k \log D)$ with high probability.

Hub set selection : random sampling

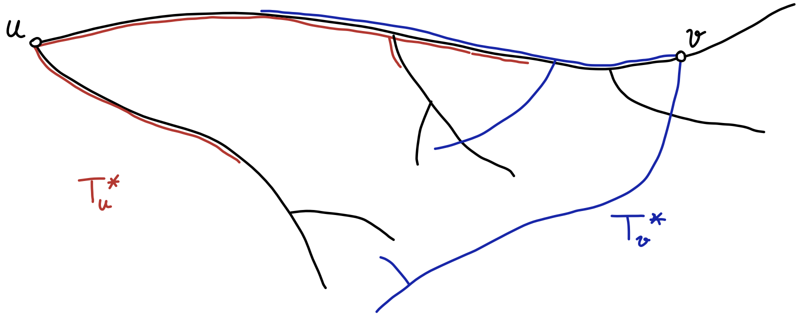

The probability to select a node x is $\propto \frac{1}{d(u,x)}$.

Hub set selection : random sampling

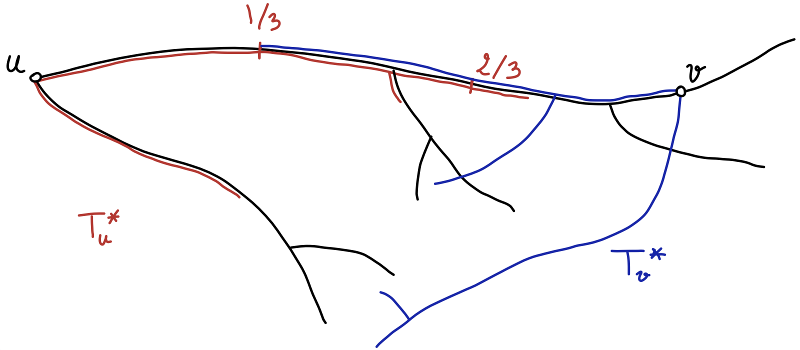

The probability to select a node x is $\propto \frac{1}{d(u,x)}$.

Hub set selection : random sampling

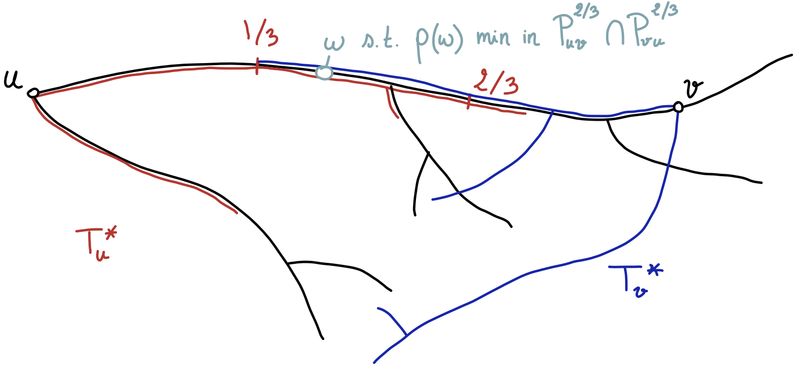

The probability to select a node x is $\propto \frac{1}{d(u,x)}$.

Hub set selection : random sampling

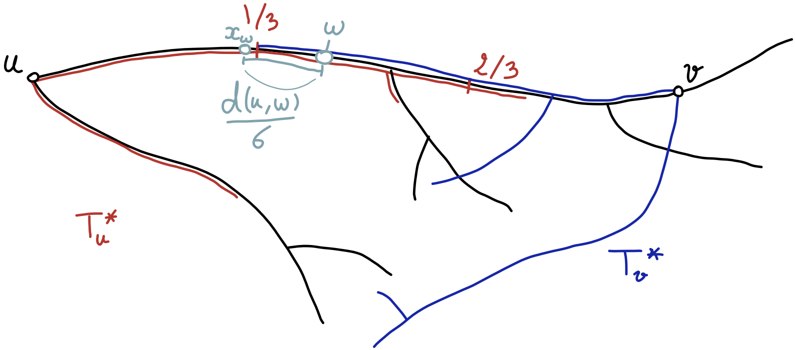

The probability to select a node x is $\propto \frac{1}{d(u,x)}$.

Road networks : two tree skeletons

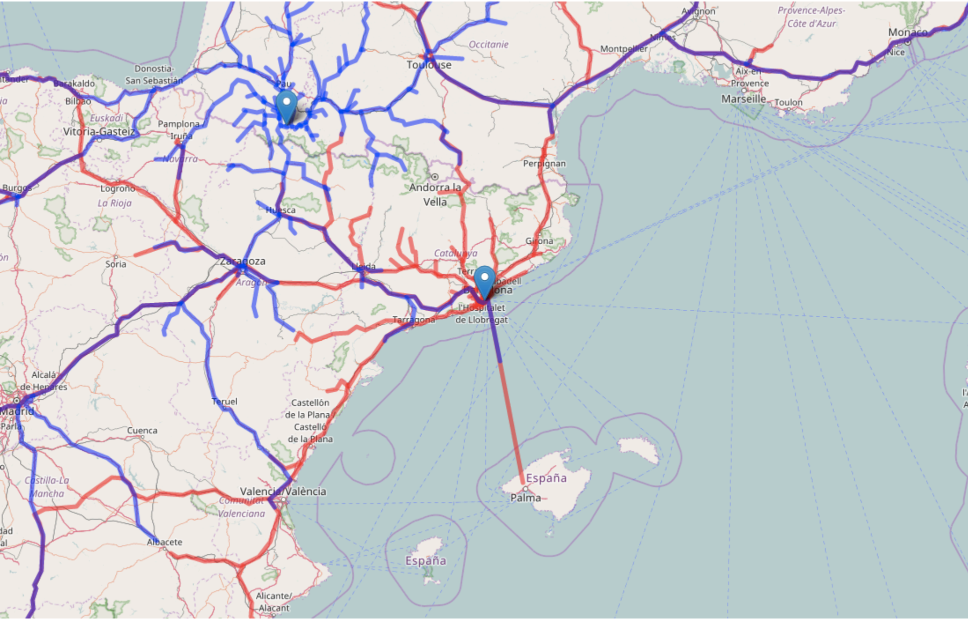

What ...maps do?

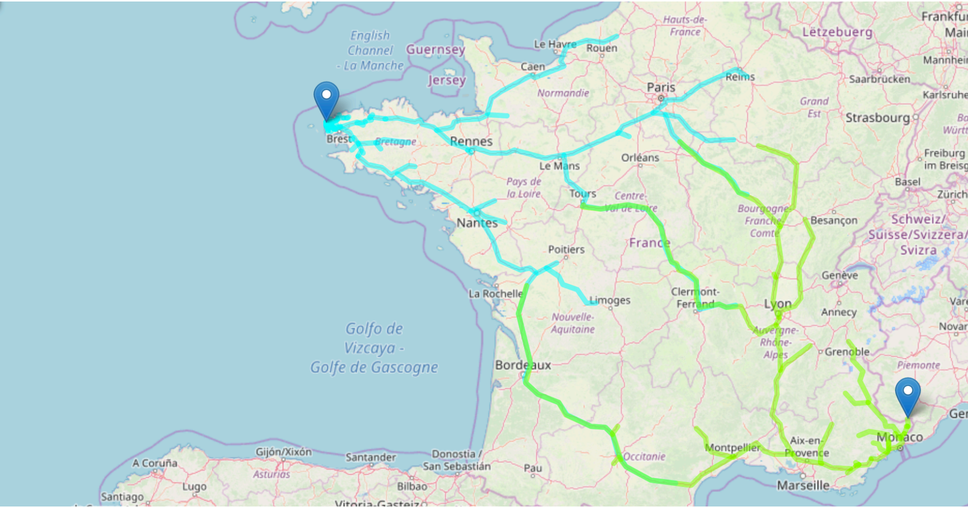

English Channel - La Manche
Guernsey
Jersey
Le Havre
Rouen
Caen
Hauts-de-France
Lëtzebuerg
Frankfurt
Mannheim
Saarbrücken
Karlsruhe
Reims
Paris
Normandie
Grand Est
Strasbourg
Brest
Bretagne
Rennes
Orléans
Le Mans
Freiburg im Breisgau
Basel
Zürich
Pays de la Loire
Tours
Centre-Val de Loire
Bourgogne-Franche-Comté
Besançon
Schweiz/Suisse/Svizzera/Svizra
Nantes
France
Poitiers
Genève
La Rochelle
Limoges
Clermont-Ferrand
Lyon
Annecy
Nouvelle-Aquitaine
Auvergne-Rhône-Alpes
Grenoble
Golfo de Vizcaya - Golfe de Gascogne
Bordeaux
Piemonte
Monaco
Gijón/Xixón
Santander
Donostia / San Sebastián
Pau
Occitanie
Montpellier
Aix-en-Provence
Marseille
A Coruña
Asturias
Lugo
Bilbao
Santiago

What ...maps do ?

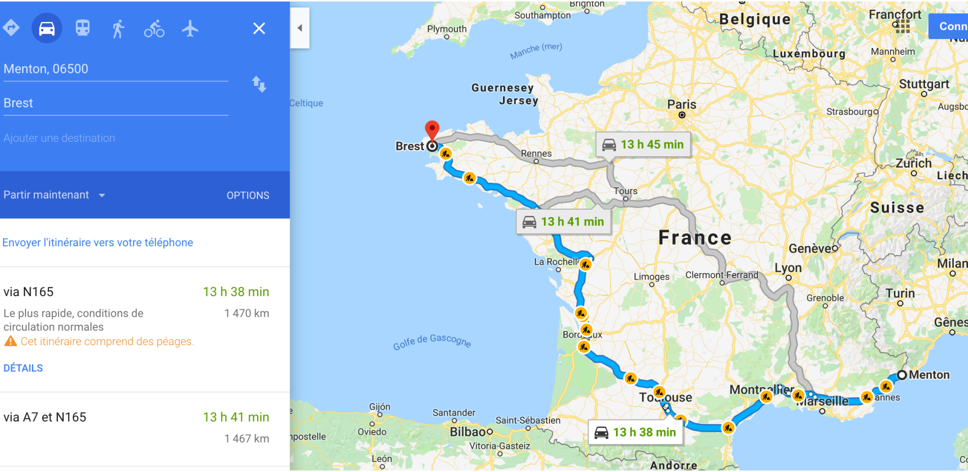

Menton, 06500
Brest
Ajouter une destination
Partir maintenant
OPTIONS
Envoyer l'itinéraire vers votre téléphone
via N165
13 h 38 min
Le plus rapide, conditions de circulation normales
1 470 km
Cet itinéraire comprend des péages.
DÉTAILS
via A7 et N165
13 h 41 min
1 467 km
13 h 45 min
13 h 41 min
13 h 38 min
Belgique
Luxembourg
Francfort
Mannheim
Stuttgart
Strasbourg
Zurich
Suisse
Genève
Milan
Turin
Gênes
Grenoble
Lyon
Clermont-Ferrand
Limoges
Paris
Tours
Rennes
Brest
La Rochelle
France
Montpellier
Marseille
Menton
Andorre
Saint-Sébastien
Vitoria-Gasteiz
Bilbao
Santander
Gijón
Oviedo
León
Golfe de Gascogne
Manche (mer)
Guernesey
Jersey
Plymouth
Southampton
Brighton
Celtique

What ...maps do ?

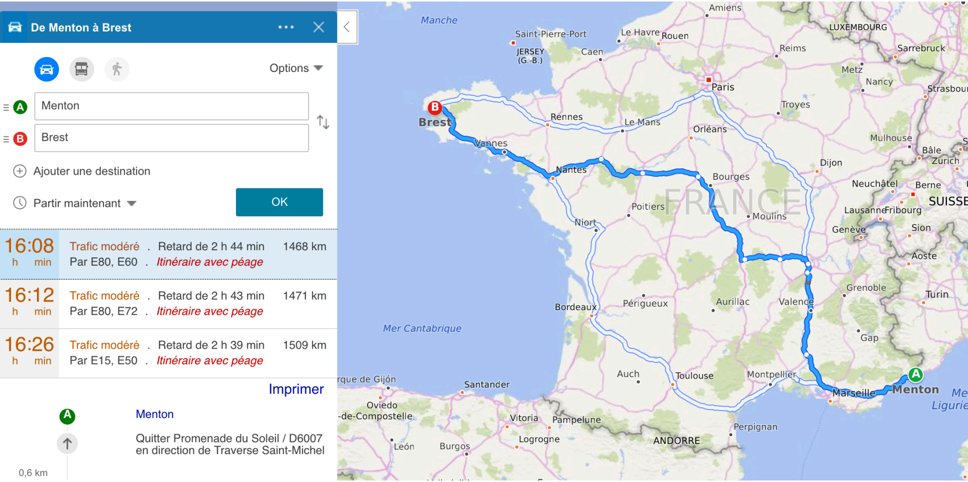
De Menton à Brest
Options
Menton
Brest
Ajouter une destination
Partir maintenant
OK
16:08 h min
Trafic modéré . Retard de 2 h 44 min
1468 km
Par E80, E60 . Itinéraire avec péage
16:12 h min
Trafic modéré . Retard de 2 h 43 min
1471 km
Par E80, E72 . Itinéraire avec péage
16:26 h min
Trafic modéré . Retard de 2 h 39 min
1509 km
Par E15, E50 . Itinéraire avec péage
Imprimer
Menton
Quitter Promenade du Soleil / D6007 en direction de Traverse Saint-Michel
0,6 km
Manche
Saint-Pierre-Port
JERSEY (G.-B.)
Le Havre
Rouen
Amiens
Caen
Reims
LUXEMBOURG
Sarrebruck
Metz
Nancy
Strasbou
Paris
Troyes
Brest
Rennes
Le Mans
Orléans
Vannes
Nantes
Mulhouse
Bâle
Zurich
Dijon
Bourges
Neuchâtel
Berne
SUISSE
Poitiers
FRANCE
Moulins
Lausanne
Fribourg
Niort
Genève
Sion
Aoste
Grenoble
Turin
Bordeaux
Périgueux
Aurillac
Valence
Mer Cantabrique
Gap
Auch
Toulouse
Montpellier
Marseille
Menton
Santander
Oviedo
-de-Compostelle
Vitoria
Pampelune
Logrogne
ANDORRE
Perpignan
León
Burgos

Highway vs skeleton in Brooklyn

Packing of 172 paths

Skeleton width 48

Skeleton dimension of grids

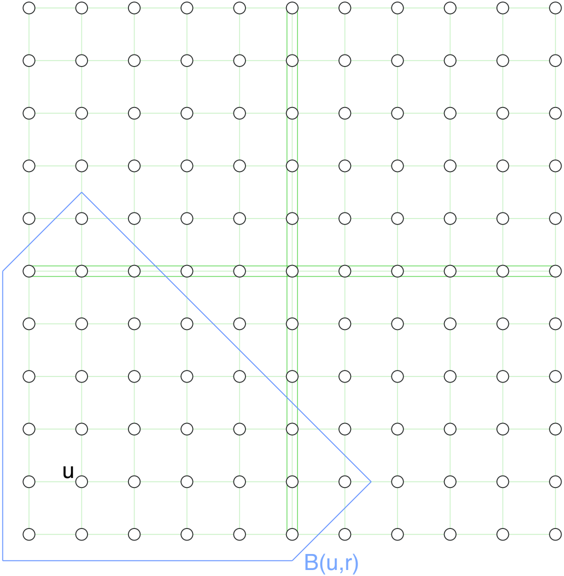

$k = \Theta(\log n)$

Skeleton dimension of grids

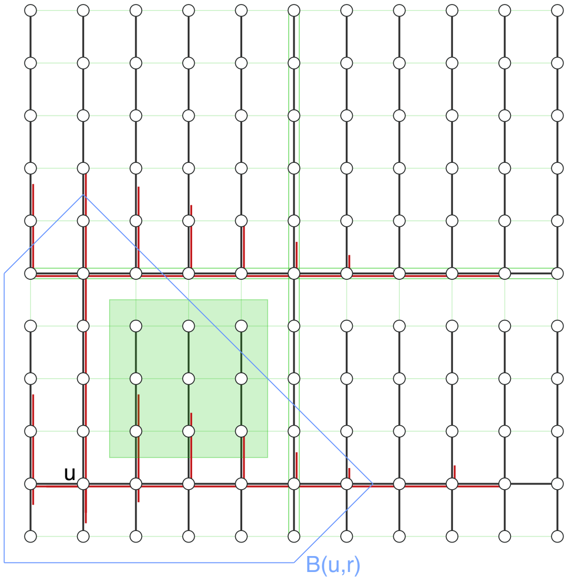

$k = \Theta(\log n)$

Skeleton dimension of grids

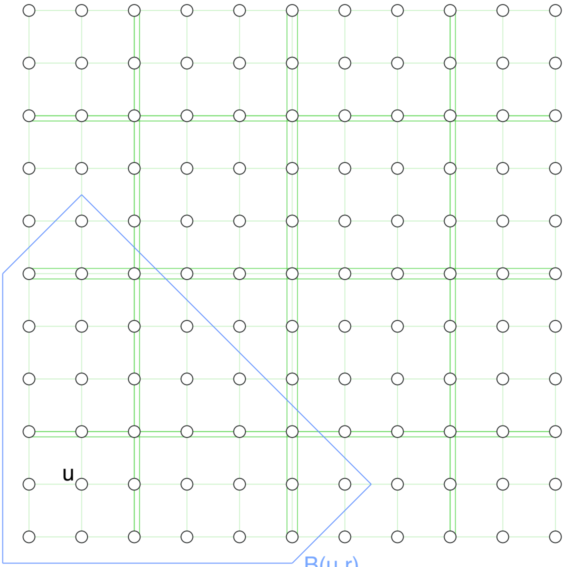

$$k = \Theta(\log n)$$

Open : random grid

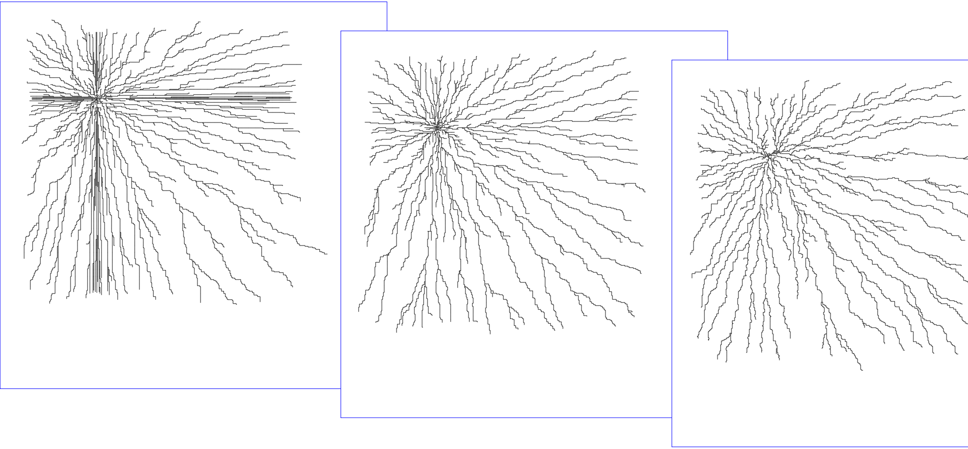

k $= 70$ k $= 49$ (fpp $[1, 4)$) k $= 49$ (prob 2/3)

Related to first-passage percolation [Licea, Newman, Piza '96] [Aldous '14].

Part III : what about 3 hops?

3-hopset of a path

3-hopset of a path

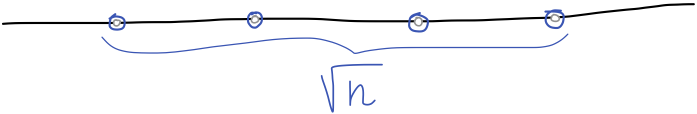

3-hopset of a path

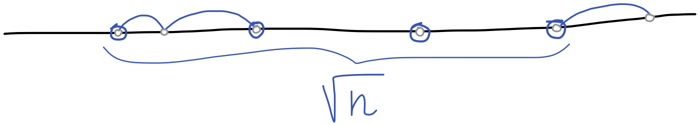

3-hopset of a path

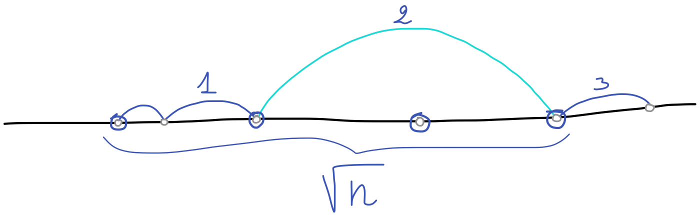

3-hopset of a path

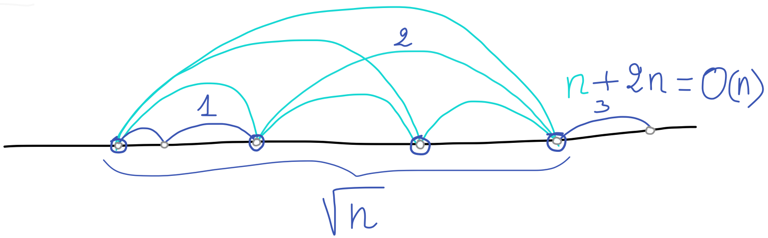

3-hopset of a path

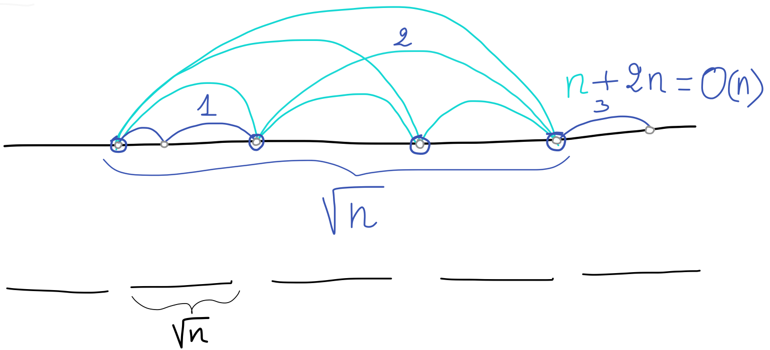

3-hopset of a path

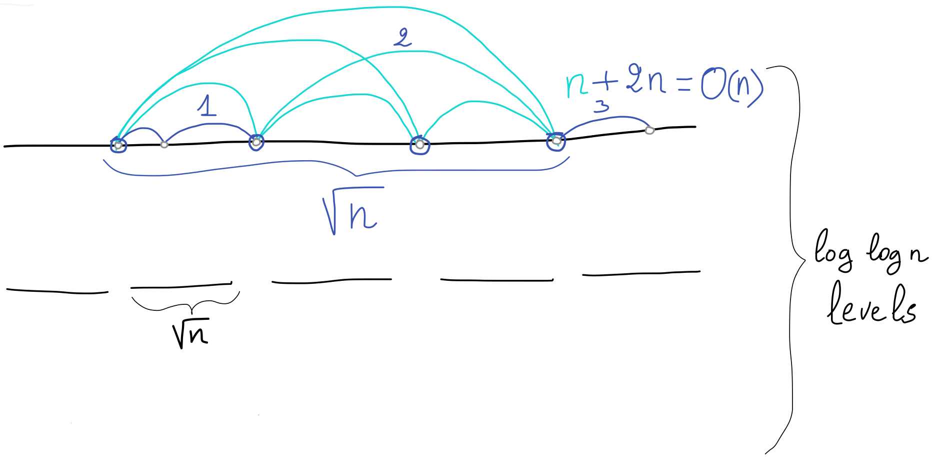

3-hopset distance oracle

Store $x, d_G(u,x)$ for $x \in N_{13}(u)$ ($2 \log \log n$ per node).

Store midle links in a hashtable H_2 ($O(n \log \log n)$ size).

Query for $d_G(u,v)$: best 3-hop path length is

$$\min_{x \in N_{13}(u), y \in N_{13}(v), xy \in H_2} d_G(u,x) + d_G(x,y) + d_G(y,v)$$

($0((\log \log n)^2)$ time).

3-hopset distance oracle

Store $x, d_G(u, x)$ for $x \in N_{13}(u)$ ($2 \log \log n$ per node).

Store midle links in a hashtable H_2 ($O(n \log \log n)$ size).

Query for $d_G(u, v)$: best 3-hop path length is

$$\min_{x \in N_{13}(u), y \in N_{13}(v), xy \in H_2} d_G(u, x) + d_G(x, y) + d_G(y, v)$$

($0((\log \log n)^2)$ time).

3-hopset distance oracle

Store $x, d_G(u,x)$ for $x \in N_{13}(u)$ ($2 \log\log n$ per node).

Store midle links in a hashtable H_2 ($O(n \log\log n)$ size).

Query for $d_G(u,v)$: best 3-hop path length is

$$\min_{x \in N_{13}(u), y \in N_{13}(v), xy \in H_2} d_G(u,x) + d_G(x,y) + d_G(y,v)$$

($0((\log\log n)^2)$ time).

Theorem (Kosowski, Gupta, V. '19)

For a unique-shortest-path graph with skeleton dimension k and average link length $L \geq 1$, there exists a randomized construction of a 3-hopset distance oracle of size $|H| = O(nk \log k(\log \log n + \log L))$, which performs distance queries in expected time $O(k^2 \log^2 k(\log^2 \log n + \log^2 L))$.

End of the story?

What is the skeleton dimension of a random grid?

Improve lower-bounds on sparse graphs for general distance oracles.

What graphs have covering hub sets of size $O(1)$?

End of the story?

What is the skeleton dimension of a random grid?

Improve lower-bounds on sparse graphs for general distance oracles.

What graphs have covering hub sets of size $O(1)$?

End of the story?

What is the skeleton dimension of a random grid?

Improve lower-bounds on sparse graphs for general distance oracles.

What graphs have covering hub sets of size $O(1)$?

Thanks.