Covering Metric Spaces by Few Trees

Yair Bartal Nova Fandina **Ofer neiman**

Tree Covers

- Let (X,d_X) be a metric space.
- A *dominating* tree T on a vertex set containing X satisfies for all $u, v \in X$,
$$d_T(u, v) \geq d_X(u, v)$$

Tree Covers

- Let (X,d_X) be a metric space.
- A *dominating* tree T on a vertex set containing X satisfies for all $u, v \in X$,
$$d_T(u,v) \geq d_X(u,v)$$
- The tree has distortion D for the pair u,v if
$$d_T(u,v) \leq D \cdot d_X(u,v)$$

Tree Covers

- Let (X,d_X) be a metric space.
- A *dominating* tree T on a vertex set containing X satisfies for all $u, v \in X$,
$$d_T(u,v) \geq d_X(u,v)$$
- The tree has *distortion* D for the pair u,v if
$$d_T(u,v) \leq D \cdot d_X(u,v)$$
- A *(D,k)-tree cover* for (X,d) is a collection of k trees $T_1,\ldots,T_k$, such that any pair u,v has a tree with distortion at most D
  - The union of the k tree is a *D-spanner* .

Ramsey Tree Covers

- Recall: A *(D,k)-tree cover* for (X,d) is a collection of k trees $T_1, \ldots, T_k$, such that any pair u,v has a tree with distortion at most D

Ramsey Tree Covers

- Recall: A *(D,k)-tree cover* for (X,d) is a collection of k trees $T_1,\dots,T_k$, such that any pair u,v has a tree with distortion at most D
- In a *Ramsey tree cover*, we want that each point has a "home" tree with distortion at most D to all other points.

Ramsey Tree Covers

- Recall: A *(D,k)-tree cover* for (X,d) is a collection of k trees $T_1,\dots,T_k$, such that any pair u,v has a tree with distortion at most D
- In a *Ramsey tree cover*, we want that each point has a "home" tree with distortion at most D to all other points.
- This is very useful for routing:
  - Since routing in a tree is easy, we can route towards u in its home tree.

Other notions of approximation via trees

- It is known that a single tree must incur linear distortion (e.g. for the cycle graph).

Other notions of approximation via trees

- It is known that a single tree must incur linear distortion (e.g. for the cycle graph).
- Most previous research focused on random embedding, and bound the expected distortion.
  - Useful in various settings, such as approximation and online algorithms.
  - A tight $\Theta(\log n)$ bound is known, and $O(\log n \log\log n)$ for spanning trees.

Other notions of approximation via trees

- It is known that a single tree must incur linear distortion (e.g. for the cycle graph).
- Most previous research focused on random embedding, and bound the expected distortion.
  - Useful in various settings, such as approximation and online algorithms.
  - A tight $\Theta(\log n)$ bound is known, and $O(\log n \log\log n)$ for spanning trees.
- There are approximation algorithm giving a tree which has distortion at most 6 times larger that the best possible tree, for unweighted graphs.

Tree Covers for General Metrics

- Small distortion regime:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees

Tree Covers for General Metrics

- Small distortion regime:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees
    - This follows from the girth lower bound: there are graphs with girth >D and $\Omega(n^{1+1/D})$ edges
    - (Gives a lower bound for Ramsey tree cover as well)

Tree Covers for General Metrics

- Small distortion regime:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees
    - This follows from the girth lower bound: there are graphs with girth >D and $\Omega(n^{1+1/D})$ edges
    - (Gives a lower bound for Ramsey tree cover as well)
  - A Ramsey tree cover with distortion D and $O(D \cdot n^{1/D})$ trees was given in [MN07].
    - Recently extended to spanning trees, with distortion O(D loglog n)

Tree Covers for General Metrics

- <u>Small distortion regime</u>:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees
    - This follows from the girth lower bound: there are graphs with girth >D and $\Omega(n^{1+1/D})$ edges
    - (Gives a lower bound for Ramsey tree cover as well)
  - A Ramsey tree cover with distortion D and $O(D \cdot n^{1/D})$ trees was given in [MN07].
    - Recently extended to spanning trees, with distortion O(D loglog n)
- <u>Small number of trees regime</u>:
  - The girth lower bound: for k trees, distortion $\Omega(\log_k n)$ is needed

Tree Covers for General Metrics

- Small distortion regime:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees
    - This follows from the girth lower bound: there are graphs with girth >D and $\Omega(n^{1+1/D})$ edges
    - (Gives a lower bound for Ramsey tree cover as well)
  - A Ramsey tree cover with distortion D and $O(D \cdot n^{1/D})$ trees was given in [MN07].
    - Recently extended to spanning trees, with distortion O(D loglog n)
- Small number of trees regime:
  - The girth lower bound: for k trees, distortion $\Omega(\log_k n)$ is needed
  - For Ramsey tree cover, we show that the technique of [MN07] with only k trees can provide distortion $O(n^{1/k} \cdot \log n)$

Tree Covers for General Metrics

- Small distortion regime:
  - Any tree cover with distortion D must contain at least $\Omega(n^{1/D})$ trees
    - This follows from the girth lower bound: there are graphs with girth >D and $\Omega(n^{1+1/D})$ edges
    - (Gives a lower bound for Ramsey tree cover as well)
  - A Ramsey tree cover with distortion D and $\tilde{O}(n^{1/D})$ trees was given in [MN07].
    - Recently extended to spanning trees, with distortion O(D loglog n)
- Small number of trees regime:
  - The girth lower bound: for k trees, distortion $\Omega(\log_k n)$ is needed
  - For Ramsey tree cover, we show that the technique of [MN07] with only k trees can provide distortion $O(n^{1/k} \cdot \log n)$
    - A large gap!

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is log λ.

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is log λ.
  - This is a standard notion of dimension for arbitrary metrics.
    - The d-dimensional space has doubling dimension Θ(d).

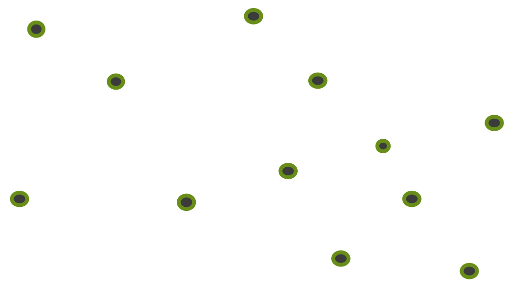

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is log λ.
  - This is a standard notion of dimension for arbitrary metrics.
    - The d-dimensional space has doubling dimension Θ(d).

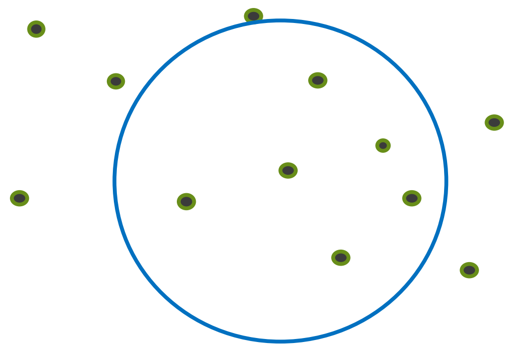

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is log λ.
  - This is a standard notion of dimension for arbitrary metrics.
    - The d-dimensional space has doubling dimension Θ(d).

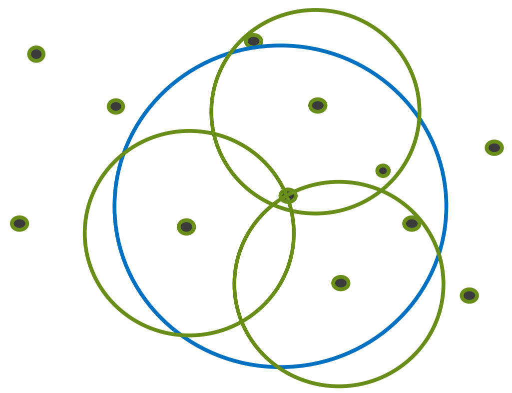

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is log λ.
  - This is a standard notion of dimension for arbitrary metrics.
    - The d-dimensional space has doubling dimension Θ(d).
- Thm: For every ε>0, every metric with doubling dimension log λ has a tree cover with $\lambda^{O(\log 1/\varepsilon)}$ trees and distortion 1+ε.
  - The number of trees is optimal (up to the constant in the O-notation)
  - Generalizes a result of [ADMSS'95] for Euclidean metrics.

Tree Covers for Doubling Metrics

- Doubling metric: every radius 2r ball can be covered by λ balls of radius r.
  - The *doubling dimension* is $\log \lambda$.
  - This is a standard notion of dimension for arbitrary metrics.
    - The d-dimensional space has doubling dimension $\Theta(d)$.
- Thm: For every $\varepsilon>0$, every metric with doubling dimension $\log \lambda$ has a tree cover with $\lambda^{O(\log 1/\varepsilon)}$ trees and distortion $1+\varepsilon$.
  - The number of trees is optimal (up to the constant in the O-notation)
  - Generalizes a result of [ADMSS'95] for Euclidean metrics.
- With distortion D we can achieve only $\tilde{O}(\lambda^{1/D})$ trees
  - Generalizes and improves the previous result of [CGMZ'05]

Lower bound for Ramsey Tree Cover

- Thm: For all n,k, there exists a *doubling* metric on n points, such that any Ramsey tree cover with k trees incurs distortion at least $\Omega(n^{1/k})$.

Lower bound for Ramsey Tree Cover

- Thm: For all n,k, there exists a *doubling* metric on n points, such that any Ramsey tree cover with k trees incurs distortion at least $\Omega(n^{1/k})$.
  - That metric is also planar (series-parallel).
  - A significant improvement over the $\Omega(\log_k n)$ girth lower bound.

Lower bound for Ramsey Tree Cover

- Thm: For all n,k, there exists a *doubling* metric on n points, such that any Ramsey tree cover with k trees incurs distortion at least $\Omega(n^{1/k})$.
  - That metric is also planar (series-parallel).
  - A significant improvement over the $\Omega(\log_k n)$ girth lower bound.
- Conclusions:
  1. Ramsey spanning trees are essentially well-understood in general/doubling/planar metrics.

Lower bound for Ramsey Tree Cover

- Thm: For all n,k, there exists a *doubling* metric on n points, such that any Ramsey tree cover with k trees incurs distortion at least $\Omega(n^{1/k})$.
  - That metric is also planar (series-parallel).
  - A significant improvement over the $\Omega(\log_k n)$ girth lower bound.
- Conclusions:
  1. Ramsey spanning trees are essentially well-understood in general/doubling/planar metrics.
  2. A large difference between tree cover and Ramsey tree cover in doubling metrics:
     - With O(1) trees we can achieve constant distortion for the former, while the latter requires polynomial distortion.

Planar and Minor-free Graphs

- Thm: For every ε>0, every planar graph has a tree cover with $O\left(\frac{\log n}{\varepsilon}\right)^2$ trees and distortion 1+ε.
  - Also holds for graphs excluding a fixed minor.

Planar and Minor-free Graphs

- Thm: For every ε>0, every planar graph has a tree cover with $O\left(\frac{\log n}{\varepsilon}\right)^2$ trees and distortion 1+ε.
  - Also holds for graphs excluding a fixed minor.
- Previous results of [GKR01] (obtained spanning trees):
  - Exact tree cover for such graphs with $\Theta(\sqrt{n})$ trees
  - Tree cover with distortion 3 and $O(\log n)$ trees.

Proof for Doubling Metrics

- Nets: An r-net is a set $N \subseteq X$ such that
  1. For all $x \in X$ there is $u \in N$ such that $d(x, u) \leq r$.
  2. For all $u, v \in N$, $d(u, v) > r$.

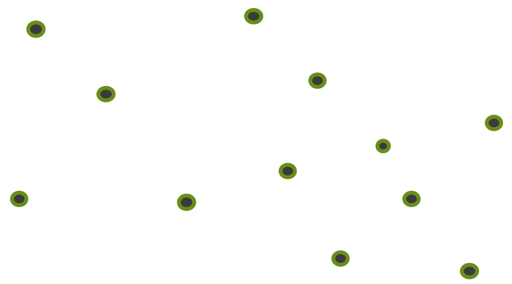

Proof for Doubling Metrics

- Nets: An r-net is a set $N \subseteq X$ such that
  1. For all $x \in X$ there is $u \in N$ such that $d(x, u) \leq r$.
  2. For all $u, v \in N$, $d(u, v) > r$.

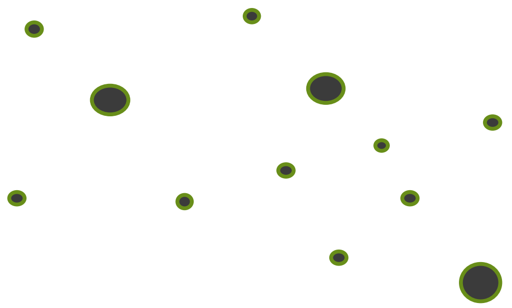

Proof for Doubling Metrics

- Nets: An r-net is a set $N \subseteq X$ such that
  1. For all $x \in X$ there is $u \in N$ such that $d(x, u) \leq r$.
  2. For all $u, v \in N$, $d(u, v) > r$.

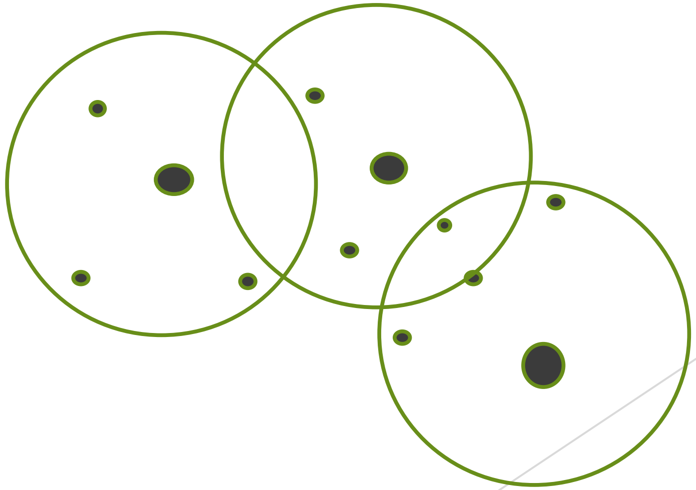

Proof for Doubling Metrics

- Nets: An r-net is a set $N \subseteq X$ such that
  1. For all $x \in X$ there is $u \in N$ such that $d(x,u) \leq r$.
  2. For all $u, v \in N$, $d(u,v) > r$.

- Nets in doubling metrics are *locally sparse*:
  - Every ball of radius R contains at most $\lambda^{2\log(R/r)}$ net points.

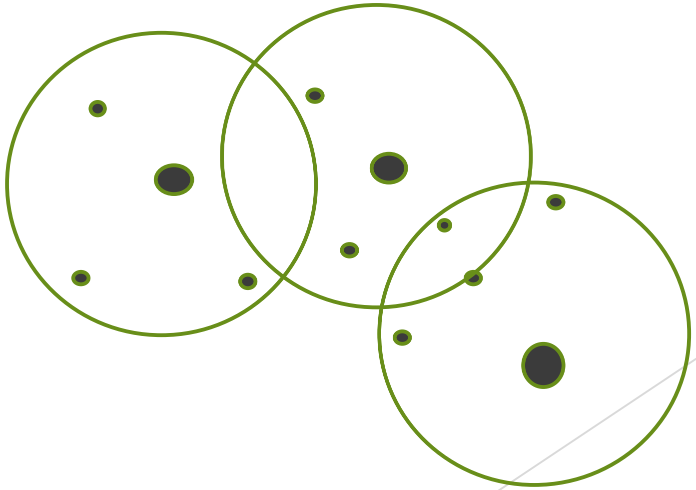

Proof for Doubling Metrics

- <u>Nets</u>: An r-net is a set $N \subseteq X$ such that
  1. For all $x \in X$ there is $u \in N$ such that $d(x, u) \leq r$.
  2. For all $u, v \in N$, $d(u, v) > r$.

- Nets in doubling metrics are *locally sparse*:
  - Every ball of radius R contains at most $\lambda^{2 \log(R/r)}$ net points.

- A simple greedy algorithm can give 2^i-nets N_i for all I that are hierarchical; $N_i \subseteq N_{i-1}$

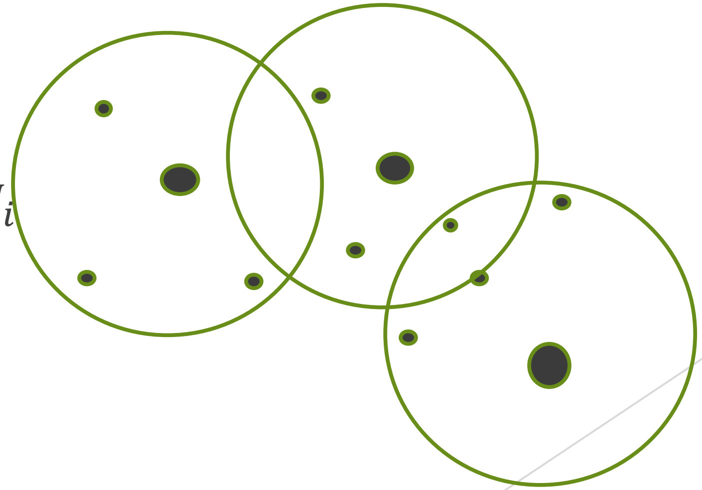

Well Separated Sub-Nets

- Fix $t = \lambda^{O(\log 1/\varepsilon)}$, a 2^i–nets N_i can be partitioned to t sets $N_{i1}, N_{i2}, \ldots, N_{it}$ that are $10 \cdot 2^i/\varepsilon$ - separated

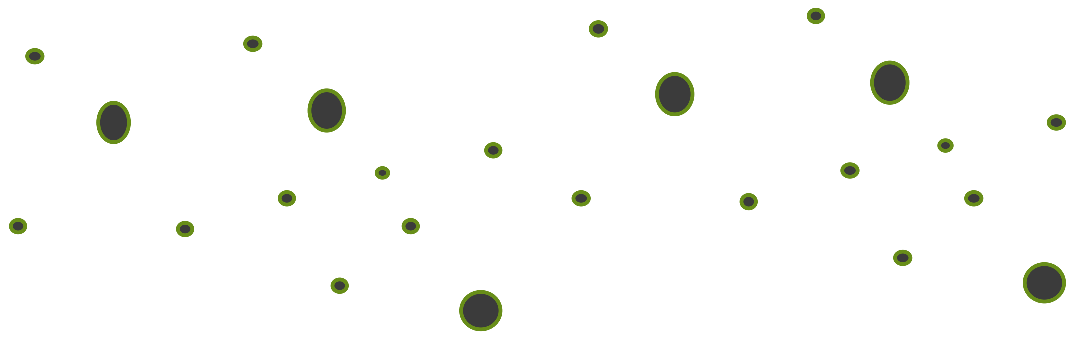

Well Separated Sub-Nets

- Fix $t = \lambda^{O(\log 1/\varepsilon)}$, a 2^i–nets N_i can be partitioned to t sets $N_{i1}, N_{i2}, \ldots, N_{it}$ that are $10 \cdot 2^i/\varepsilon$ - separated

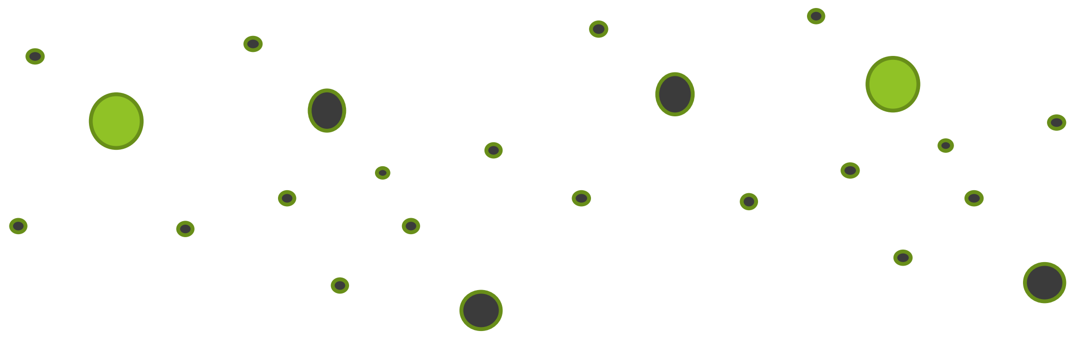

Well Separated Sub-Nets

- Fix $t = \lambda^{O(\log 1/\varepsilon)}$, a 2^i–nets N_i can be partitioned to t sets $N_{i1}, N_{i2}, \dots, N_{it}$ that are $10 \cdot 2^i/\varepsilon$ - separated

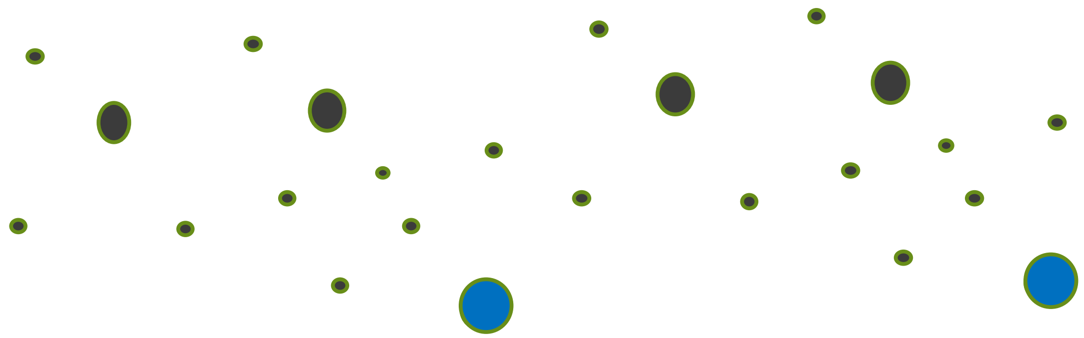

Well Separated Sub-Nets

- Fix $t = \lambda^{O(\log 1/\varepsilon)}$, a 2^i–nets N_i can be partitioned to t sets $N_{i1}, N_{i2}, \dots, N_{it}$ that are $10 \cdot 2^i/\varepsilon$ - separated

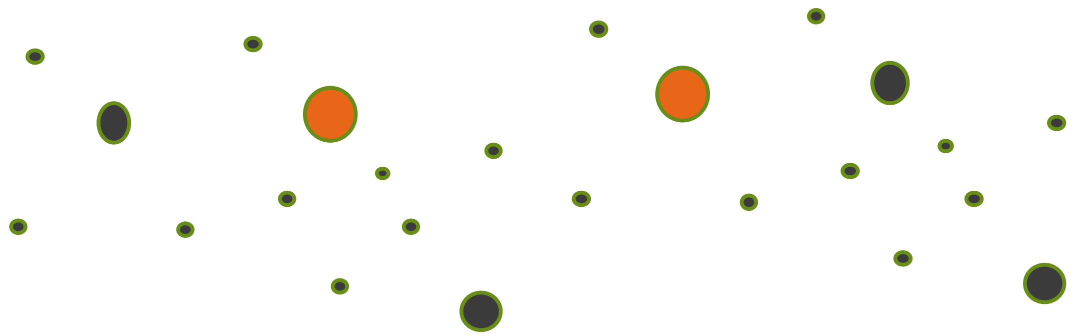

Clustering the Subnets $\{N_{ij}\}_i$

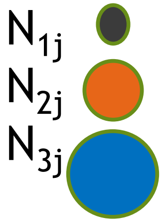

- Initially all points are not clustered.
- Go over all indices i in increasing order
  - Every $x \in N_{ij}$ creates a cluster of all unclustered points within $3 \cdot 2^i/\varepsilon$
  - All these point (except x) are now clustered

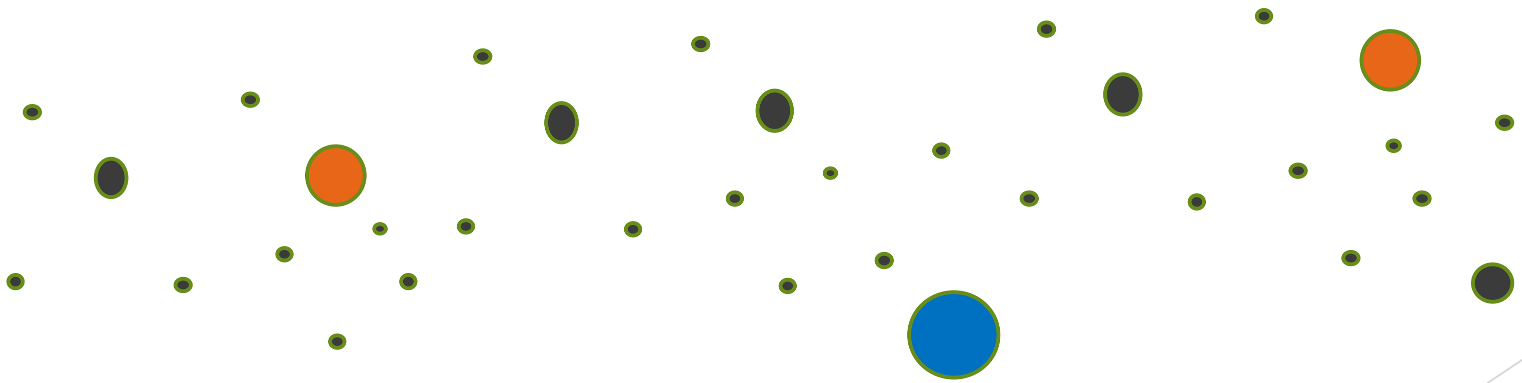

Clustering the Subnets $\{N_{ij}\}_i$

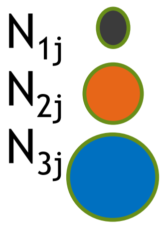

- Initially all points are not clustered.
- Go over all indices i in increasing order
  - Every $x \in N_{ij}$ creates a cluster of all unclustered points within $3 \cdot 2^i/\varepsilon$
  - All these point (except x) are now clustered

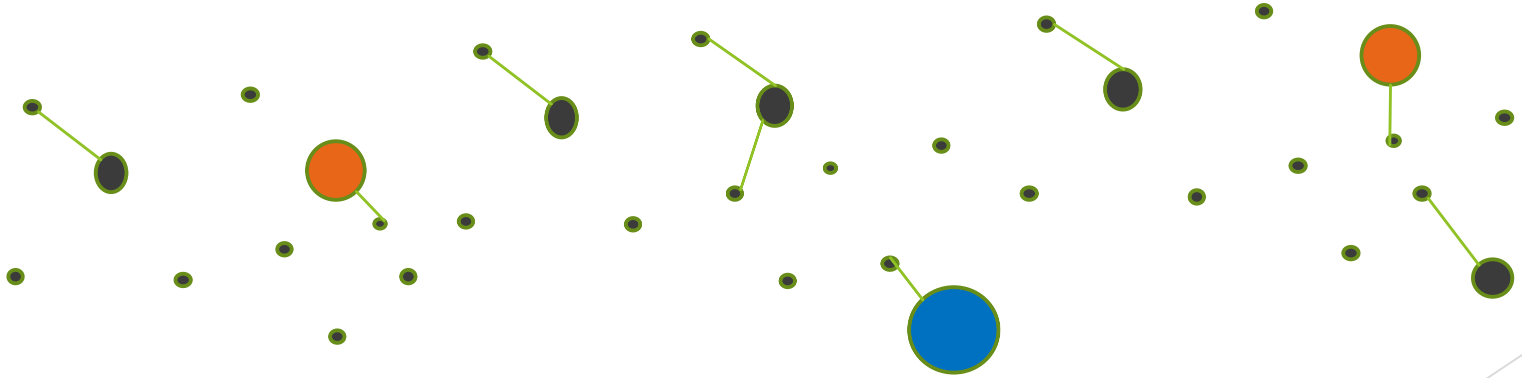

Clustering the Subnets $\{N_{ij}\}_i$

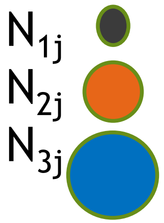

- Initially all points are not clustered.
- Go over all indices i in increasing order
  - Every $x \in N_{ij}$ creates a cluster of all unclustered points within $3 \cdot 2^i/\varepsilon$
  - All these point (except x) are now clustered

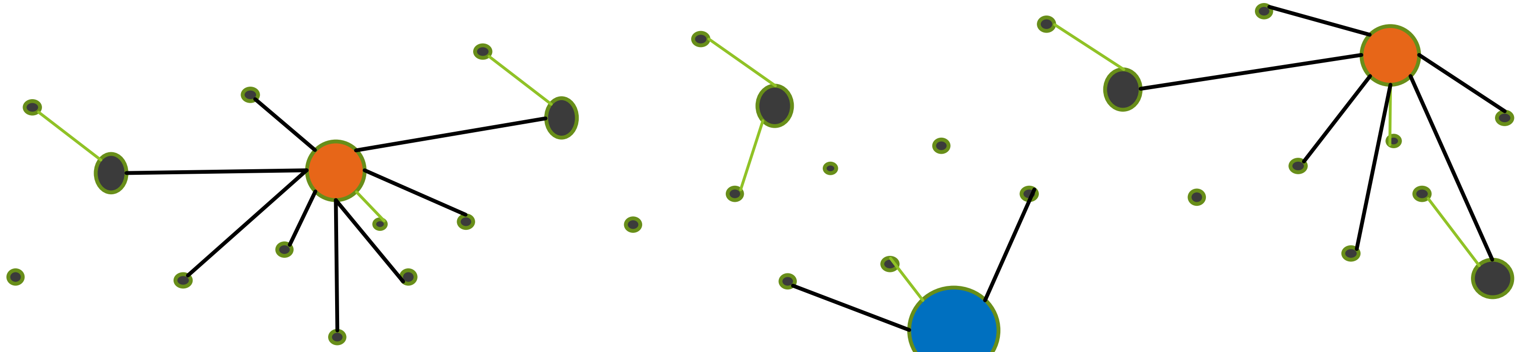

Clustering the Subnets $\{N_{ij}\}_i$

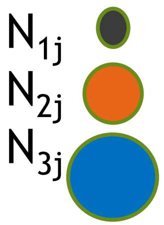

- Initially all points are not clustered.
- Go over all indices i in increasing order
  - Every $x \in N_{ij}$ creates a cluster of all unclustered points within $3 \cdot 2^i/\varepsilon$
  - All these point (except x) are now clustered

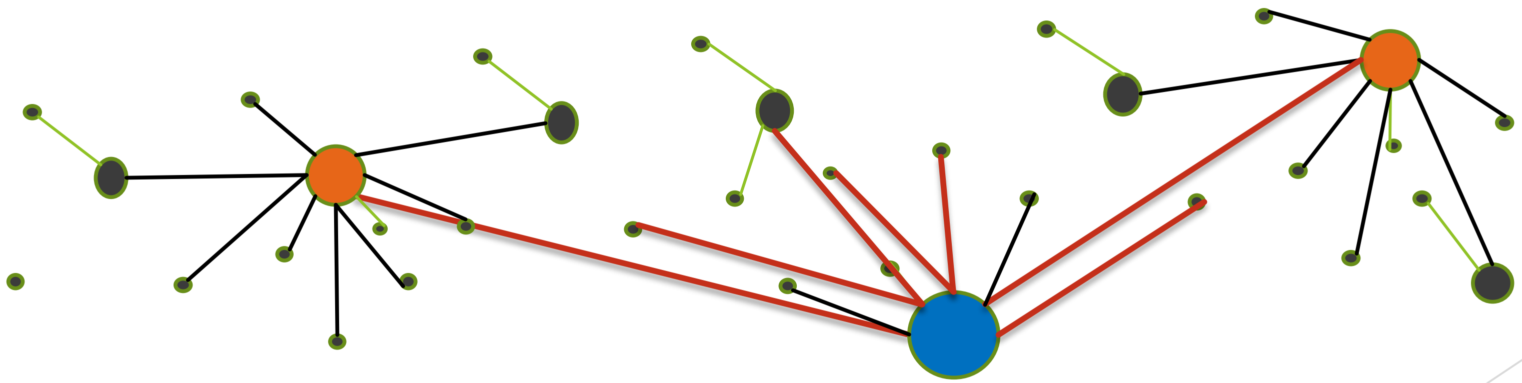

Trees Construction

- The clustering induces a forest for every j=1,2,…,t

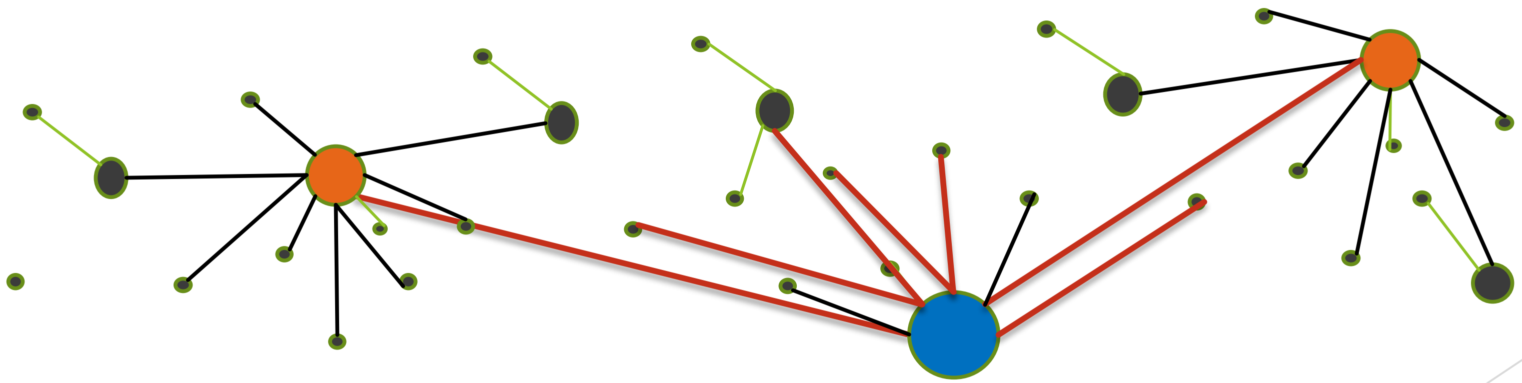

Trees Construction

- The clustering induces a forest for every j=1,2,…,t
- In fact, we cluster with ε-gaps, i.e. if p=log 1/ε
  - N_{1j}, $N_{(p+1)j}$, $N_{(2p+1)j}$,…
  - N_{2j}, $N_{(p+2)j}$, $N_{(2p+2)j}$,…
  - N_{3j}, $N_{(p+3)j}$, $N_{(2p+3)j}$,…

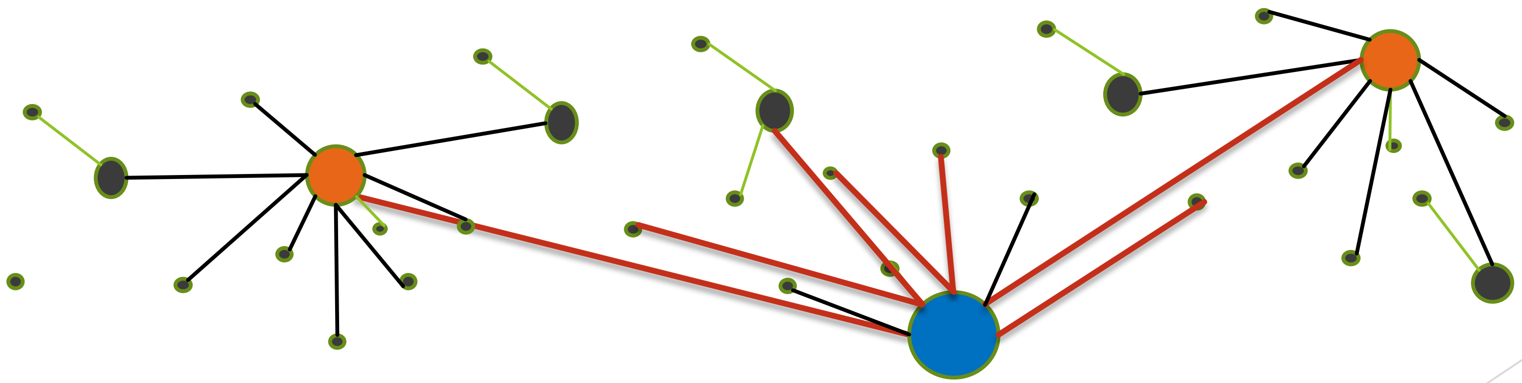

Trees Construction

- The clustering induces a forest for every j=1,2,…,t
- In fact, we cluster with ε-gaps, i.e. if p=log 1/ε
  - $N_{1j}, N_{(p+1)j}, N_{(2p+1)j}, \ldots$
  - $N_{2j}, N_{(p+2)j}, N_{(2p+2)j}, \ldots$
  - $N_{3j}, N_{(p+3)j}, N_{(2p+3)j}, \ldots$
- So we have total of $t \cdot p = \lambda^{O(\log 1/\varepsilon)}$ forests.

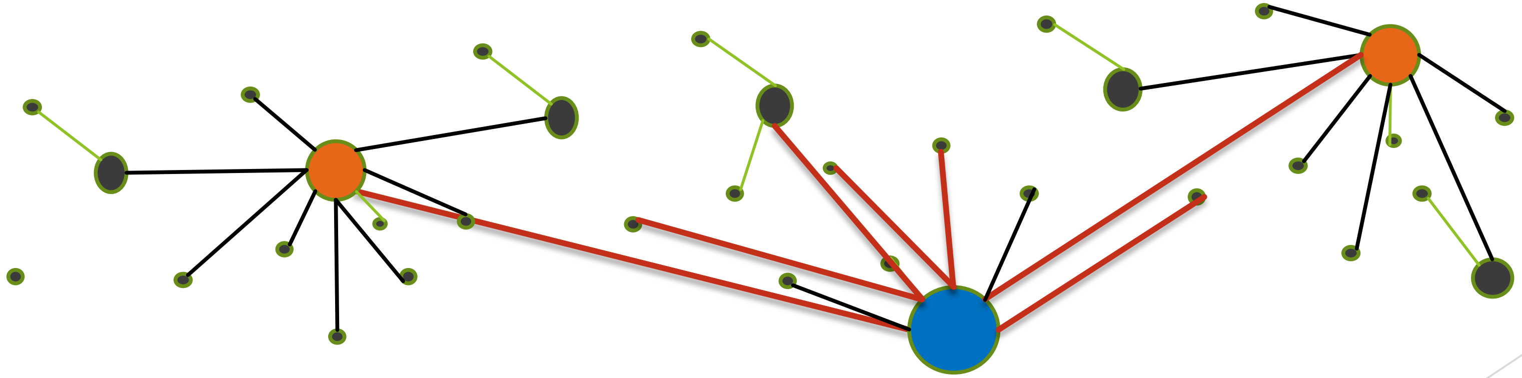

Observations on Clustering

1. Since N_{ij} is $10 \cdot 2^i/\varepsilon$ -separated , and the clustering is done to distance $3 \cdot 2^i/\varepsilon$, no point is clustered more than once.

Observations on Clustering

1. Since N_{ij} is $10 \cdot 2^i/\varepsilon$ -separated , and the clustering is done to distance $3 \cdot 2^i/\varepsilon$, no point is clustered more than once.
2. The diameter of a level i cluster is at most $8 \cdot 2^i/\varepsilon$.

   Can be proved by a simple induction

Observations on Clustering

1. Since N_{ij} is $10 \cdot 2^i/\varepsilon$ -separated , and the clustering is done to distance $3 \cdot 2^i/\varepsilon$, no point is clustered more than once.
2. The diameter of a level i cluster is at most $8 \cdot 2^i/\varepsilon$.

   Can be proved by a simple induction
3. For every y in the level i cluster centered at x, there is a path of length at most $d(x, y) + O(2^i)$ between them

Observations on Clustering

1. Since N_{ij} is $10 \cdot 2^i/\varepsilon$ -separated , and the clustering is done to distance $3 \cdot 2^i/\varepsilon$, no point is clustered more than once.
2. The diameter of a level i cluster is at most $8 \cdot 2^i/\varepsilon$.

   Can be proved by a simple induction
3. For every y in the level i cluster centered at x, there is a path of length at most $d(x,y) + O(2^i)$ between them
   - Pf: suppose y belongs to z's cluster just before level i
   - z's cluster is level at most i-p, so has diam

     $8 \cdot 2^{i-p}/\varepsilon = 8 \cdot 2^i$
   - Thus $d(x,z) \leq d(x,y) + O(2^i)$

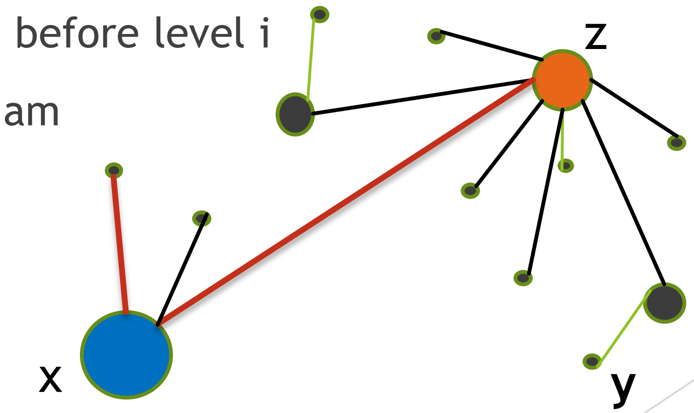

Observations on Clustering

1. For every y in the level i cluster centered at x, there is a path of length at most $d(x,y)+O\left(2^i\right)$ between them
   - Pf: suppose y belongs to z's cluster just before level i
   - z's cluster is level at most i-p, so has diam

     $8 \cdot 2^{i-p}/\varepsilon = 8 \cdot 2^i$
   - Thus $d(x,z) \le d(x,y)+O\left(2^i\right)$
   - Finally, $d_T(x,y)=d(x,z)+d_T(z,y) \le d(x,y)+O\left(2^i\right)$

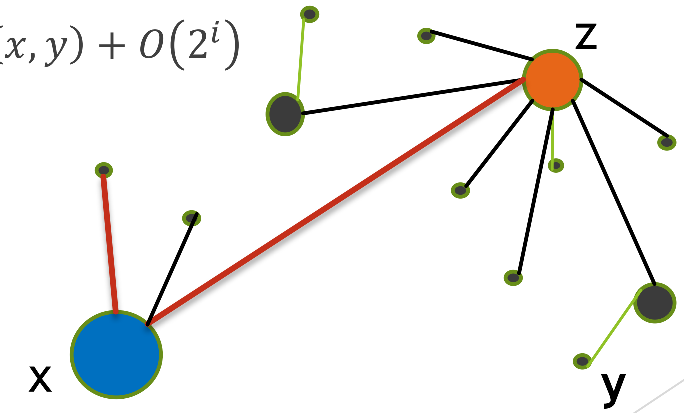

Bounding the Stretch

- Suppose u,v are such that $d(u,v) \approx 2^i/\varepsilon$.
- There is a net point $x \in N_{ij}$ such that $d(u,x) \leq 2^i$ (for some $1 \leq j \leq t$).

Bounding the Stretch

- Suppose u,v are such that $d(u,v) \approx 2^i/\varepsilon$.
- There is a net point $x \in N_{ij}$ such that $d(u,x) \le 2^i$ (for some $1 \le j \le t$).
- Suppose u in y's cluster, and v in z's cluster just before level i

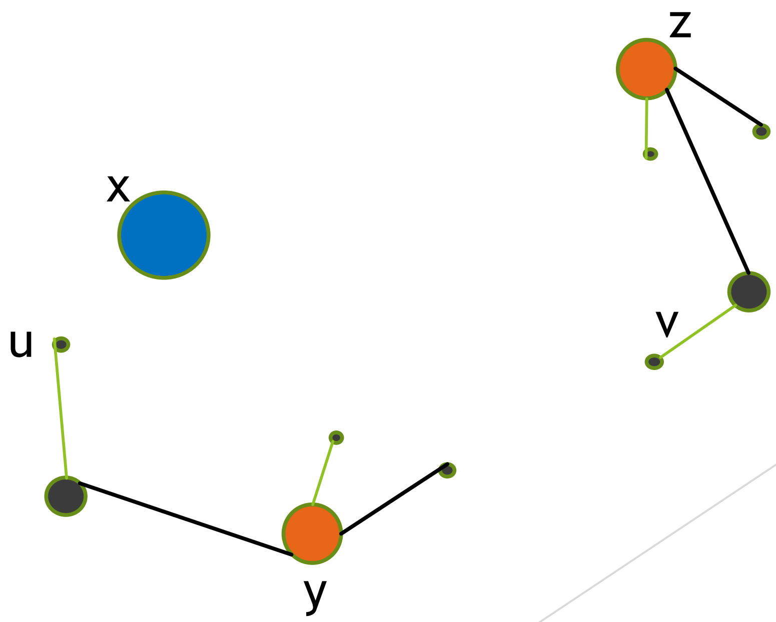

Bounding the Stretch

- Suppose u,v are such that $d(u,v) \approx 2^i/\varepsilon$.
- There is a net point $x \in N_{ij}$ such that $d(u,x) \leq 2^i$ (for some $1 \leq j \leq t$).
- Suppose u in y's cluster, and v in z's cluster just before level i.
- Since both are clusters of level at most i-p, their diameters are $O(2^i)$

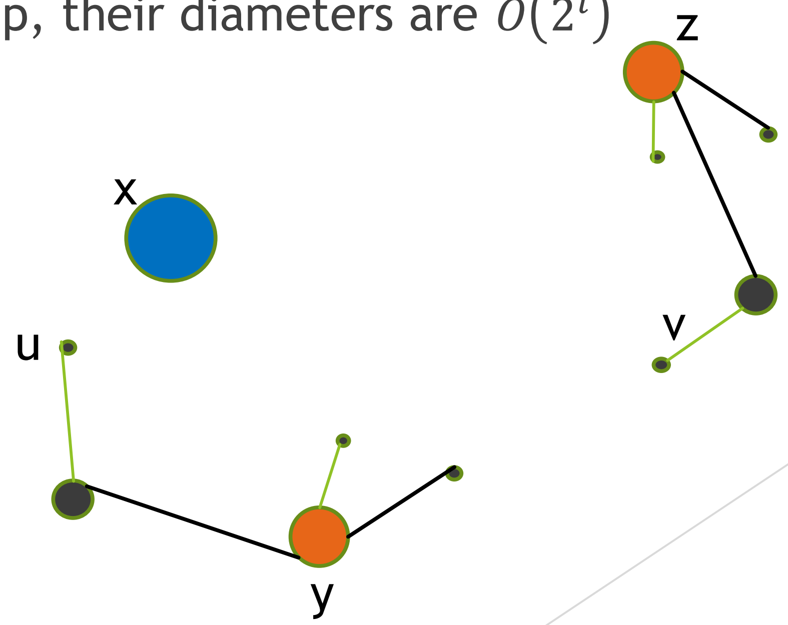

Bounding the Stretch

- Suppose u,v are such that $d(u,v) \approx 2^i/\varepsilon$.
- There is a net point $x \in N_{ij}$ such that $d(u,x) \leq 2^i$ (for some $1 \leq j \leq t$).
- Suppose u in y's cluster, and v in z's cluster just before level i.
- Since both are clusters of level at most i-p, their diameters are $O(2^i)$
- Both y,z are within $3 \cdot 2^i/\varepsilon$, so x will cluster them

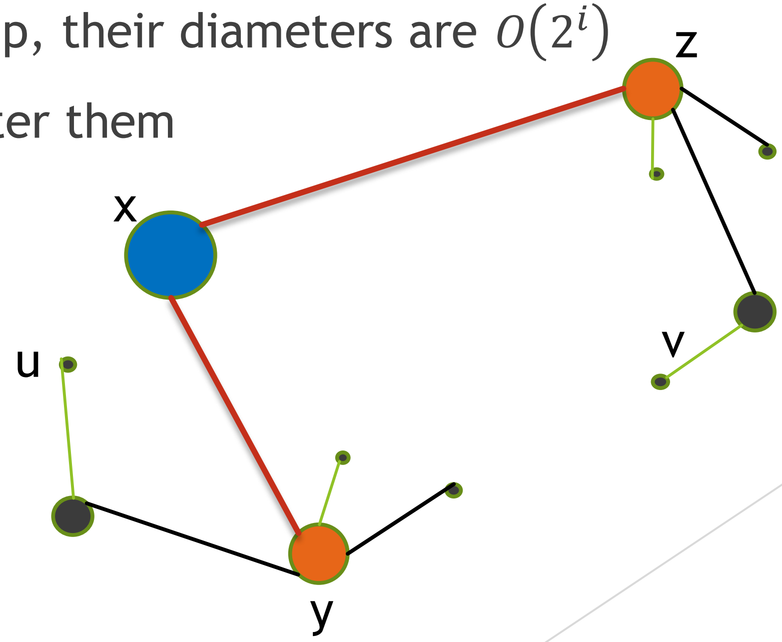

Bounding the Stretch

- Suppose u,v are such that $d(u,v) \approx 2^i/\varepsilon$.
- There is a net point $x \in N_{ij}$ such that $d(u,x) \le 2^i$ (for some $1 \le j \le t$).
- Suppose u in y's cluster, and v in z's cluster just before level i.
- Since both are clusters of level at most i-p, their diameters are $O(2^i)$
- Both y,z are within $3 \cdot 2^i/\varepsilon$, so x will cluster them

$$
\begin{aligned}
d_T(u,v) &= d_T(u,y) + d_T(y,x) + d_T(x,z) + d_T(z,v) \\
&\le d(y,x) + d(x,z) + O(2^i) \\
&\le d(u,v) + O(2^i) \\
&= (1+\varepsilon) \cdot d(u,v)
\end{aligned}
$$

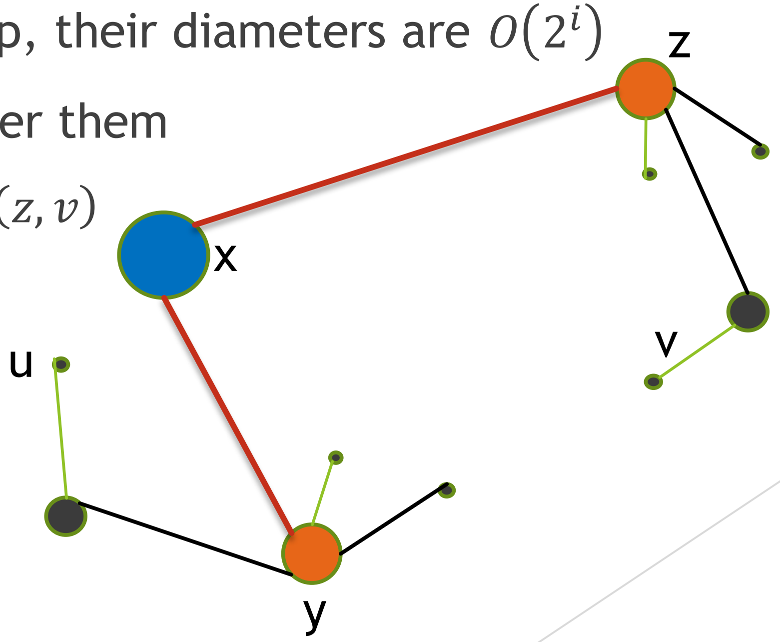

Open Questions

- Tree cover with k trees for general metrics:
  - Is there a $\Omega(n^{1/k})$ lower bound on the distortion?
  - Or can we get logarithmic distortion? Maybe $O(\log_k n)$.
- Tree covers for planar graphs with O(1) trees?
- Obtain *spanning* tree covers for doubling graphs.