

Average Size of Unstretched Remote-Spanners

Philippe Jacquet, Laurent Viennot

► **To cite this version:**

Philippe Jacquet, Laurent Viennot. Average Size of Unstretched Remote-Spanners. 5th ACM Workshop on Analytic Algorithmics and Combinatorics (ANALCO), 2009, New York, United States. SIAM, pp.23–33, 2009, <http://www.siam.org/proceedings/analco/2009/anl09_004_jacquetp.pdf>. <inria-00471727>

HAL Id: inria-00471727

<https://hal.inria.fr/inria-00471727>

Submitted on 8 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Average size of unstretched Remote-Spanners

Philippe Jacquet*

Laurent Viennot[†]

November 25, 2008

Abstract

Motivated by the optimization of link state routing in ad hoc networks, and the concept of multipoint relays, we introduce the notion of remote-spanner. Given an unweighted graph G , a remote spanner is a set of links H such that for any pair of nodes (u, v) there exists a shortest path in G for which all links in the path that are not adjacent to u belong to H . The remote spanner is a kind of minimal topology information beyond its neighborhood that any node would need in order to compute its shortest paths in a distributed way. This can be extended to k -connected graphs by considering minimum length sum over k disjoint paths as distance.

In this paper, we give distributed algorithms for computing remote-spanners in order to obtain sparse remote-spanners with various properties. We provide a polynomial distributed algorithm that computes a k -connecting unstretched remote-spanner whose number of edges is at a factor $2(1 + \log \Delta)$ from optimal where Δ is the maximum degree of a node. Interestingly, its expected compression ratio in number of edges is $O(\frac{k}{n} \log n)$ in Erdős-Rényi graph model and $O((\frac{k}{n})^{\frac{2}{3}})$ in the unit disk graph model with a uniform Poisson distribution of nodes.

1 Introduction

This paper concerns the characterization and the distributed computation of sparse remote-spanners. Given an unweighted graph G , a sub-graph H with vertex set $V(H) = V(G)$ is an *unstretched-remote-spanner* if it keeps distances in G from any node u when it is completed with all neighboring links of u . More precisely, for any two nonadjacent nodes u, v , $d_{H_u}(u, v) = d_G(u, v)$, where H_u is the sub-graph with edge set $E(H) \cup \{uv \mid v \in N(u)\}$ and d_{H_u} is the distance in H_u . (Note that $d_{H_u}(u, v) = 1 = d_G(u, v)$ when u and v are adjacent).

We introduce this notion based on the functioning of link state routing protocols used in practical networks where each router generally knows its list of neighbors. Link state routing was proposed by McQuillan et al. [22] as a replacement for distance vector routing. It was then standardized as OSPF protocol [23, 24]. With a very high level description, link state routing basically consists in two periodic procedures. First, each router sends regularly probing messages on its network interfaces to discover its neighbors. Second, it regularly floods the network with link state advertisement messages containing its list of neighbors. (Additional messages may be sent when the list of neighbors changes). Each node then

*INRIA, Rocquencourt, Polytechnique, Paris-Sud University, France. INRIA Project-team “HIPERCOM”. Supported by Systematic project ”MOBISIC” philippe.jacquet@inria.fr

[†]INRIA, Paris Diderot University, France. Project-Team “GANG” between INRIA and LIAFA. Supported by ANR project “ALADDIN” and CRC “MARDI” with France-Telecom Orange. laurent.viennot@inria.fr

knows its list of neighbors and the whole network topology. This can be very costly in a large and dense network, a case that can be encountered in ad hoc networks where wireless connections may provide many neighbors to each node. To optimize link state routing in such situations, it was proposed more recently to alleviate the cost of link state advertisements by flooding only a subset of links [21, 8]. This was standardized by IETF as the OLSR routing protocol [7].

This principle can be indeed applied to any link state routing protocol: broadcast only a subset of links to all nodes, thus defining a sub-graph H . As each node u regularly discovers its neighbors, it can augment this graph with its neighboring links, to obtain a sub-graph H_u with edge set $E(H) \cup \{uv \mid v \in N(u)\}$. It then computes its routing tables according to distances in H_u : it forwards packets with destination v to a closest neighbor u' to v in H_u . u' then forwards similarly the packet and so on. This results in a classical greedy routing scheme. As the path from u' to v in H_u is included in H , it is known by u' , implying $d_{H_{u'}}(u', v) \leq d_{H_u}(u, v) - 1$. We thus see that this greedy routing from u to v results in a route of length at most $d_G(u, v)$. The notion of unstretched-remote-spanner thus formalizes the required properties on the broadcasted sub-graph H to ensure that greedy routing performs without route optimality loss.

This formalization is inspired by the regular notion of graph spanner introduced by Peleg et al. [26, 25]. An unstretched-spanner is a sub-graph H such that for all nodes u, v , $d_H(u, v) = d_G(u, v)$. Conversely to remote-spanners, neighboring links outside the spanner are not considered. Spanners are key ingredients of various distributed applications, e.g., synchronizers [2, 26], compact routing [1, 27, 29], covers [3], dominating sets [12], distance oracles [4, 30], emulators and distance preservers [5, 10], broadcasting [17], near-shortest path algorithms [13, 14, 15, 16]. Recent reviews of the literature on spanners can be found in [28, 32]. We are working on an extension of remote-spanner where an (α, β) stretch is allowed. and we believe that a part of this work can be investigated in the context where a node knows its neighboring links in addition to the spanner as this information is usually accessible in practical routing context.

We introduce the concept of multi-connectivity in remote spanner. Preserving multi-connectivity has practical interest for improving reliability of the network and to allow multi-path routing. We says that two nonadjacent nodes u and v are k -connected in G if there exists k pairwise disjoint paths from u to v (i.e. having no internal node in common). An unstretched remote-spanner H is said to be k -connecting if for all nonadjacent nodes u, v and all positive integer $k' \leq k$, u and v are k' -connected in H_u if they are k' -connected in G . Additionally, we require that the stretch of the length sum of these paths is one, i.e. $d_{H_u}^{k'}(u, v) = d_G^{k'}(u, v)$ where $d_G^{k'}(u, v)$ is minimum length sum of k' disjoint paths in a sub-graph K .

Conversely to spanners, we show that remote-spanners have local characterizations that yield simple distributed algorithms for computing them. Additionally, unstretched-remote-spanners can be sparse whereas an unstretched spanner must obviously include all edges. Finally, we are able to compute remote-spanners providing both multi-connectivity and controlled stretch.

1.1 Our results

We show that any unstretched-remote-spanner is an union of stars called *multipoint relay star* dominating nearby nodes. More precisely, we define an multipoint relay star T_u for a node u

as a subset of adjacent links of u neighborhood of node u , such that for all v at distance 2 from u , there exists $(u, x) \in T_u$ and x is neighbor of v . In other words, $T_u \cup \{u\}$ dominates the ball $B(u, 2)$ of radius 2 centered at u . Notice that the set of end nodes of the multipoint relay star is a multipoint relay set according to the terminology in [7]. We say that a sub-graph H induces multipoint relays if it contains a multipoint relay set for each node in the graph, *i.e.* for every node u there exist a multipoint relay star T for u with $E(T) \subseteq E(H)$. We provide in Section 2 a distributed algorithm for computing the multipoint relay stars of the unstretched-remote-spanner for a given node whose number of edges is within a factor $1 + \log \Delta$ from optimal where Δ denote the maximal degree of a node.

In Section 3, we study k -connecting remote-spanners. We provide a distributed algorithm computing a k -connecting unstretched-remote-spanner in a polynomial time and with optimal number of edges up to a factor $2(1 + \log \Delta)$.

Interestingly, we also show in section 4 that the k -connecting unstretched remote spanner shows an expected number of edges in $O(kn \log n)$ in Erdős-Rényi graph model and in $O(k^{\frac{2}{3}} n^{\frac{4}{3}})$ in the unit disk graph model with a uniform Poisson distribution of nodes (compared to $\Omega(n^2)$ if the remote-spanner is equal to the graph itself). In the context of wireless network, Erdős-Rényi graph model is pertinent about indoor wireless networking, whereas unit disk graph model is pertinent about outdoor wireless networking.

1.2 Related work

Our unstretched-remote-spanner generalizes the notions of multipoint relays introduced in ad hoc networks [21, 7] for optimizing flooding and shortest path routing and extended in [33, 21, 7] for providing small connected dominating sets. It was already known that they provide shortest path routes, *i.e.* their union forms an unstretched-remote-spanner. However, it was not noticed that they are also necessary: any unstretched-remote-spanner must induce multipoint relays. As multipoint relays are also used for optimizing flooding, this definition was extended to obtain better reliability of flooding with the k -coverage feature [7, 9]. It was never proved that this extension indeed ensures k -connectivity.

The concept of remote-spanner is introduced in [19] as an extension of regular spanner. This work mainly focuses on remote-spanners with stretch, *i.e.* preserving distances up to some multiplicative and additive factors. More precisely, a remote-spanner H of a graph G has stretch (α, β) if for each pair of nodes u and v the distance between u and v in H_u , the graph H augmented by all the edges between u and its neighbors in G , is at most α times the distance between u and v in G plus β . This definition is extended to k -connected graphs by considering minimum length sum over k disjoint paths as distance. It is then said that an (α, β) -remote-spanner is k -connecting. Distributed algorithms for computing in constant time $(1 + \varepsilon, 1 - 2\varepsilon)$ -remote-spanners and 2-connecting $(2, -1)$ -remote-spanners are given. The resulting spanner has a linear number of edges if the input graph is the unit ball graph of a doubling metric. Distributed algorithms for computing k -connecting $(1, 0)$ -remote-spanners with optimal number of edges up to a logarithmic factor are also proposed in [19]. However, no bound is given on the minimum number of edges of such a remote-spanner.

2 1-connected Unstretched-remote-spanners

We prove in Section 3 that a sub-graph is an unstretched-remote-spanner iff it induces $(2, 0)$ -dominating trees. However, we postpone the study of such remote-spanners to Section 3 in

the more general setting of k -connecting unstretched-remote-spanners.

Algorithm 1 proposes an algorithm called $\text{DOMTREEGDY}(u)$ for computing an multipoint relay star for node u . It consists in solving greedily a set-cover problem for dominating the nodes at distance 2 from u .

```

 $T := (\{u\}, \emptyset)$ 
 $M_2 := \emptyset, S := B(u, 2) \setminus B(u, 1)$ 
while  $S \neq \emptyset$  do
  Pick  $x \in N(u) \setminus M_2$  such that  $|B(x, 1) \cap S|$  is maximal.
   $M_2 := M_2 \cup \{x\}$ 
  Add to  $T$  the link  $u$  to  $x$  in  $G$ .
   $S := S \setminus B(x, 1)$ 

```

Algorithm 1: Algorithm $\text{DOMTREEGDY}(u)$ for a node u . The tree T is the multipoint relay star computed for u .

Proposition 1 *Algorithm $\text{DOMTREEGDY}(u)$ computes a multipoint relay star for node u with minimal number of edges up to a factor $1 + \log \Delta$.*

The proof of this proposition adapted from [6] about the classical greedy heuristic for set-cover is written *in extenso* in [19].

3 Remote-spanners providing multi-connectivity

We consider only simple paths, i.e. a node appears at most once in a path. We say that two paths are *disjoint* if they do not have any internal node in common. Several paths are disjoint if they are pairwise disjoint. We define the k -connecting distance $d_G^k(s, t)$ between two nodes s and t as the minimum length sum obtained over all sets of k disjoint paths from s to t . (We set $d_G^k(s, t) = \infty$ if there do not exist k disjoint paths from s to t). We thus have $d_G^1(s, t) = d_G(s, t)$. We similarly define $d_H^k(s, t)$ for any sub-graph H .

Recall that an unstretched-remote-spanner H is said to be k -connecting if it satisfies $d_{H_s}^{k'}(s, t) = d_G^{k'}(s, t)$ for all nonadjacent nodes s and t and all positive integer $k' \leq k$. This definition is equivalent to the unstretched-remote-spanner definition for $k = 1$. Recall also that a k -connecting multipoint relay star T for a node u is a tree sub-graph rooted at node u dominating every node v at distance 2 from u in the following sense: v has k neighbors in $B_T(u, 1)$, or $uw \in E(T)$ for all $w \in N(u) \cap N(v)$. This definition is equivalent to the multipoint relay star definition for $k = 1$.

3.1 k -connecting unstretched-remote-spanners

We now characterize k -connecting unstretched-remote-spanners as sub-graphs inducing multipoint relay stars. It is clearly a necessary condition: if H is a k -connecting unstretched-remote-spanner, consider two nodes u and v such that $d_G(u, v) = 2$. If u and v have k' common neighbors with $1 \leq k' \leq k$, then the stretch condition implies that $d_{H_v}^{k'}(v, u) \leq 2k'$. As minimal path length between u and v is 2, u and v must thus have at least k' common neighbors in H_v . H must thus contain a k -connecting multipoint relay star for u . Indeed, we can obtain the following characterization.

Proposition 2 *A sub-graph is a k -connecting unstretched-remote-spanner iff it induces k -connecting multipoint relay stars.*

Consider a sub-graph H inducing multipoint relay stars and two nonadjacent nodes s and t such that $d_G^{k'}(s, t) < \infty$ for some k' with $1 \leq k' \leq k$.

Lemma 1 *Among all tuples $P_1, \dots, P_{k'}$ of k' disjoint paths from s to t , consider one with minimal length. If P_1 lies outside H by $i \geq 2$, then there exists a path P'_1 from s to t with same length as P_1 such that $P'_1, P_2, \dots, P_{k'}$ are disjoint and P'_1 lies outside H by $i - 1$ edges.*

The proof of this lemma can be found in [19].

By iteratively applying Lemma 1, we obtain that there exist k' disjoint paths with minimal length, all of them lying 1-outside H . This implies that $d_{H^c}^{k'}(s, t) = d_G^{k'}(s, t)$ and H is thus a k -connecting unstretched remote-spanner. Proposition 2 follows.

3.2 Computing k -connecting remote-spanners

Algorithm 2 is similar to $\text{DOMTREEGDY}(u)$ except that it keeps adding nodes in the dominating tree for node u until nodes at distance 2 are dominated k times. For $k = 1$, it is equivalent to $\text{DOMTREEGDY}(u)$. The greedy heuristic in classical generalization of the set-cover problem still performs within a factor $1 + \log \Delta$ from optimal [11, 31].

```

 $T := (\{u\}, \emptyset)$ 
 $M := \emptyset, S := B(u, 2) \setminus B(u, 1), X := N(u)$ 
while  $S \neq \emptyset$  do
  Pick  $x \in X \setminus M$  such that  $|B(x, 1) \cap S|$  is maximal.
   $M := M \cup \{x\}$ 
  Add edge  $ux$  to  $T$ .
   $S := S \setminus \{v \in S \mid N(v) \cap N(u) \subseteq M \text{ or } |N(v) \cap M| \geq k\}$ 

```

Algorithm 2: Algorithm $\text{DOMTREEGDY}_k(u)$ for a node u . The tree T is the dominating tree computed for u .

As long as there remains a node v in S which is initially the set of nodes at distance 2 from u , it is not dominated k times and it has a neighbor x which is not in M , the set of nodes added as leaves of T . It is thus always possible pick some x at the beginning of the while loop until S is empty. We can thus state the following proposition.

Proposition 3 *Algorithm $\text{DOMTREEGDY}_k(u)$ computes a k -connecting multipoint relay star for node u with minimal number of edges up to a factor $1 + \log \Delta$.*

Theorem 1 *A k -connecting unstretched-remote-spanner with number of edges within a factor $2(1 + \log \Delta)$ from optimal can be computed in time $O(1)$.*

The proof is in [19].

4 Analysis of remote-spanner in random network models

4.1 Erdős-Rényi random graph

In this model we suppose that a network with n nodes, each pair of node being independently connected by an edge with probability p . Therefore the average node degree $\delta = pn$. The following result are adapted from [20]

Proposition 4 *In the Erdős-Rényi random graph with n node average degree $\delta = pn$ the average size of the optimal multipoint relay star is smaller than $\frac{1}{p} \log n + \frac{1}{p}$ when $n \rightarrow \infty$ for $n \rightarrow \infty$ with p fixed.*

Proof. Let consider a node u and let assume that one randomly adds edges in T_u until all nodes at distance 2 of u are connected. Let assume that the sequence of selected neighbors is : $v_1, v_2, \dots, v_\ell, \dots$. Let w be a node distinct of u , we have

$$P(N(w) \cap \{u, v_1, \dots, v_\ell\} = \emptyset) = (1 - p)^{\ell+1} .$$

Thus

$$\begin{aligned} P(|T_u| > \ell) &= P(\exists w : d_G(u, w) = 2 : N(w) \cap \{v_1, \dots, v_\ell\} = \emptyset) \\ &\leq P(\exists w : N(w) \cap \{u, v_1, \dots, v_\ell\} = \emptyset) \\ &\leq \sum_{w \neq u} P(N(w) \cap \{u, v_1, \dots, v_\ell\} = \emptyset) = (n - \ell - 1)(1 - p)^{\ell+1} \end{aligned}$$

Consequently

$$\begin{aligned} E(|T_u|) &= \sum_{\ell} P(|T_u| > \ell) \\ &\leq \sum_{\ell} \min\{1, n(1 - p)^{\ell+1}\} \\ &\leq \lfloor \frac{\log n}{\log \frac{1}{1-p}} \rfloor + \frac{1}{p} \end{aligned}$$

□

Proposition 5 *In the Erdős-Rényi random graph with n node average degree $\delta = pn$ the average size of the optimal k -connected multipoint relay star is smaller than $O(\frac{k}{p} \log n)$ when $n \rightarrow \infty$ for $n \rightarrow \infty$ with p fixed.*

Proof. We build T_u as in the previous proof excepted that we go ahead adding new edge until every node w at distance 2 from u are covered at least k times. The probability that a node w is not neighbor of u and is not covered by less than k times by the nodes v_1, v_2, \dots, v_ℓ is $(1 - p) \sum_{i < k} (1 - p)^{\ell-i} p^i \binom{\ell}{i}$. This quantity is smaller than $(1 - p) \exp(-\frac{k}{p} \ell + 1)$, using Chernov bound.

Therefore

$$\begin{aligned} E(|T_u|) &\leq \sum_{\ell} \min\{1, n \exp(-\frac{k}{p} \ell + 1)\} \\ &\leq \frac{k}{p} (\log n + 1) + \frac{e}{1 - e^{-\frac{k}{p}}} \end{aligned}$$

□

We define as the *topology compression ratio* of the remote spanner the ratio of the number of links in the remote spanner to the total number of links in the graph.

Theorem 2 *The topology compression average ratio of a k -connecting unstretched remote spanner in an Erdős-Rényi random graph with n nodes and average degree $\delta = pn$ is smaller than $O(\frac{k}{p} \frac{\log n}{\delta}) = O(\frac{k}{n} \log n)$.*

4.2 Unit disk random graph

We consider that nodes are dispatched in a large square of area R with a Poisson uniform distribution. The number of nodes is then a random variable N with average value $E(N) = n$. We denote $\nu = \frac{n}{R}$ the per unit area node density. Recall that the unit disk graph is the set of connected nodes within distance at most 1 in the plane.

Proposition 6 *In the unit disk graph with uniform Poisson density ν , the average size of the optimal multipoint relay star is smaller than $3\pi\nu^{\frac{1}{3}} + O(\nu^{-\frac{1}{3}})$ edges when $\nu \rightarrow \infty$.*

Proof. Let u be a node far from square border. We proof consists in building a multipoint relay star T_u with average number of edges at most $3\pi\nu^{\frac{1}{3}} + O(\nu^{-\frac{1}{3}})$ edges. We claim that this bound remains valid for nodes close to the border.

Let $\varepsilon < 1$ be a non negative number. We denote $C(u, \varepsilon)$ the region of points which are at cartesian distance between $1 - \varepsilon$ and 1 to node u , and $K(u, \varepsilon) \subseteq N(u)$ the set of nodes which are located in $C(u, \varepsilon)$. The main idea is to add in T_u an edge uv for every $v \in K(u, \varepsilon)$. Note that the area of the $C(u, \varepsilon)$ is $(2\varepsilon - \varepsilon^2)\pi$ with tight upper-bound $2\pi\varepsilon$ when $\varepsilon \rightarrow 0$.

We denote $D(u, r)$ the disk of radius r centered on u in the plane and we call *crown nodes* the nodes which are in $D(u, 2) \setminus D(u, 1)$. Let v be a crown node at distance $2 - x$ from u for some $0 < x < 1$. Let $A_\varepsilon(x)$ denote the area of $C(u, \varepsilon) \cap D(v, 1)$. We denote $N_\varepsilon(x) = |N(v) \cap K(u, \varepsilon)|$ the number of nodes already in T_u that dominate such a node v . The random variable $N_\varepsilon(x)$ is Poisson of parameter $\nu A_\varepsilon(x)$.

We call ε -missed node any crown node v such that the two following conditions hold: (i) $N(v) \cap K(u, \varepsilon) = \emptyset$, and (ii) $N(v) \cap N(u) \neq \emptyset$. To obtain a multipoint relay star for u , we must add edges in T_u so that any ε -missed node v is dominated by one neighbor of u in T_u . We thus complete T_u by adding at most one edge for each ε -missed node v (one added edge may cover more than one of these nodes).

Notice that ε -missed nodes cannot be at cartesian distance to node u greater than $2 - \varepsilon$ since all their common neighbors with u would then be in $C(u, \varepsilon)$.

The average number of nodes inside $K(u, \varepsilon)$ is smaller than $2\pi\nu\varepsilon$. Let $J_\varepsilon(\nu)$ be the average number of ε -missed nodes. We have

$$\begin{aligned} J_\varepsilon(\nu) &= 2\pi\nu \int_\varepsilon^1 (2-x) P(N_\varepsilon(x) = 0) dx \\ &= 2\pi\nu \int_\varepsilon^1 (2-x) \exp(-\nu A_\varepsilon(x)) dx \end{aligned} \tag{1}$$

The average size of T_u is therefore bounded by

$$E(|T_u|) \leq 2\pi\nu\varepsilon + J_\varepsilon(\nu).$$

The main aim is to give an upper bound of $J_\varepsilon(\nu)$.

Let $A(x)$ be the size of the area $D(u, 1) \cap D(v, 1)$ when the cartesian distance between u and v is exactly $2 - x$ for $0 < x < 1$. Since the disk $D(u, 1 - \varepsilon)$ that makes the internal border of $C(u, \varepsilon)$ is included in $D(w, 1)$ for any node w at cartesian distance ε to node u , we have $A_\varepsilon(x) > A(x) - A(x - \varepsilon)$ when $x > \varepsilon$.

Via pure geometric consideration we have the identity

$$A(x) = 2\theta - \sin(2\theta)$$

where $\theta = \arccos(1 - x)$. Notice that when $x \rightarrow 0$, we have $A(x) = \frac{4}{3}\theta^3 + O(\theta^5)$ and therefore $A(x) = \frac{2^{5/2}}{3}x^{\frac{3}{2}} + O(x^{\frac{5}{2}})$. Similarly $A_\varepsilon(x) > 2^{3/2}\varepsilon(x - \varepsilon)^{\frac{1}{2}} + O(\varepsilon(x - \varepsilon)^{\frac{3}{2}})$ for $x > \varepsilon$. In fact we have $A_\varepsilon(x) > \rho\varepsilon(x - \varepsilon)^{\frac{1}{2}}$ for some $\rho > 0$ uniformly for all $x \in [\varepsilon, 1]$

Therefore

$$J_\varepsilon(\nu) < 4\pi\nu \int_0^{1-\varepsilon} \exp(-2^{3/2}\nu\varepsilon x^{1/2}(1 + O(x)))dx$$

By the change of variable $y = \varepsilon x^{\frac{1}{2}}$ and using the uniform domination of the integrand by function $\exp(-\rho\varepsilon\nu x^{\frac{1}{2}})$ we obtain

$$\begin{aligned} J_\varepsilon(\nu) &< \frac{4\pi\nu}{(\varepsilon\nu)^2} \int_0^\infty \exp(-2^{3/2}y^{\frac{1}{2}})dy + O\left(\frac{\nu}{(\varepsilon\nu)^4}\right) \\ &= \frac{\pi\nu}{(\varepsilon\nu)^2} + O\left(\frac{\nu}{(\varepsilon\nu)^4}\right) \end{aligned}$$

Therefore we have the inequality (assuming $\varepsilon\nu \rightarrow \infty$).

$$E(|T_u|) \leq 2\pi\varepsilon\nu + \frac{\pi\nu}{(\varepsilon\nu)^2} + O\left(\frac{\nu}{(\varepsilon\nu)^4}\right)$$

The minimum of the right-hand side (omitting the $O()$ term) is attained when $\varepsilon = \nu^{-\frac{2}{3}}$. \square

Proposition 7 *In the unit disk graph with uniform Poisson density ν , the average number of edges in an optimal k -connecting multipoint relay star is $O(k^{\frac{2}{3}}\nu^{\frac{1}{3}})$.*

Proof. Let u be a node far from square border. The proof consists in building a k -connecting multipoint relay star T_u with average number of edges at most $ck^{2/3}\nu^{1/3}$ edges, for some constant c and for a sufficiently high density ν . We claim that this bound remains valid for nodes close to the border. Using Proposition 2, we will thus obtain a k -connecting multipoint relay star with average size $O(k^{2/3}\nu^{1/3})$.

We still use the main idea of adding in T_u an edge uv for every $v \in K(u, \varepsilon)$. We call (ε, k) -missed node any crown node v such that (i) $|N(v) \cap K(u, \varepsilon)| < k$, and (ii) $|N(v) \cap N(u)| > |N(v) \cap K(u, \varepsilon)|$. To obtain a k -connecting multipoint relay star for u , we must add edges in T_u so that any (ε, k) -missed node v is dominated by $\min\{k, |N(v) \cap N(u)|\}$ neighbors of u in T_u . We thus complete T_u by adding $\min\{k, |N(v) \cap N(u)|\} - |N(v) \cap K(u, \varepsilon)|$ edges for each (ε, k) -missed node v .

We still have all (ε, k) -missed nodes at cartesian distance to node u smaller than $2 - \varepsilon$. We choose ε to ensure $\nu A_\varepsilon(\varepsilon) = \alpha k$ for some $\alpha > 1$ (e.g. $\alpha = 2$). To estimate $P(N_\varepsilon(x) < k)$ we use Chernov's bound that is for Poisson random variable $N_\varepsilon(x)$:

$$P(N_\varepsilon(x) < k) \leq \exp(-\nu A_\varepsilon(x)\beta)$$

with $\beta = (1 - \frac{1}{\alpha}(1 - \log(\alpha)))$.

We have $A_\varepsilon(x) > \alpha'\varepsilon(x - \varepsilon)^{\frac{1}{2}}$ uniformly for all $x \in [\varepsilon, 1]$, for some $\alpha' > 0$. Let $J_{(\varepsilon, k)}(\nu)$ be the average number of (ε, k) -missed nodes we have

$$\begin{aligned} J_{(\varepsilon, k)}(\nu) &= 2\pi\nu \int_\varepsilon^1 (2 - x)P(N_\varepsilon(x) < k)dx \\ &\leq 4\pi\nu \int_0^\infty \exp(-\beta\nu\alpha'\varepsilon x^{\frac{1}{2}})dx \\ &= \frac{\pi\nu}{(\beta\alpha'\varepsilon\nu)^2} + O\left(\frac{\nu}{(\beta\alpha'\varepsilon\nu)^4}\right) \end{aligned}$$

The issue is to count the nodes needed in $N(u) \setminus K(u, \varepsilon)$ in order to appropriately dominate ε -missed nodes. This means that we have $|V(T_u) \setminus K(u, \varepsilon)| \leq kJ_{(\varepsilon, k)}(\nu)$.

$$\begin{aligned} |E(T_u)| &\leq 2\pi\varepsilon\nu + kJ_k(\nu) \\ &\leq \pi\left(\frac{3}{2}\alpha\right)^{2/3}k^{\frac{2}{3}}\nu^{\frac{1}{3}} + \left(2\left(\frac{2}{3\alpha}\right)^{2/3}\right)^2\frac{1}{\beta^2\alpha}k^{-\frac{1}{3}}\nu^{\frac{1}{3}} \\ &\quad + O\left(k^{-\frac{5}{3}}\nu^{-\frac{1}{3}}\right), \end{aligned}$$

since $\nu A(\varepsilon) = \alpha k$ and therefore $\varepsilon \approx \frac{3^{2/3}}{2^{5/3}}(\alpha k)^{2/3}\nu^{-\frac{2}{3}}$. □

Theorem 3 *The topology compression average ratio of a k -connecting unstretched remote spanner in a unit disk graph with n nodes and average degree δ is smaller than $O\left(\left(\frac{k}{\delta}\right)^{\frac{2}{3}}\right)$.*

5 Concluding remarks

We have introduced the notion of unstretched-remote-spanner which is well suited for grasping the trade-offs when optimizing the subset of links advertised in a link state routing protocol. Most strikingly, we proposed distributed construction of sparse remote-spanners which also enjoy multi-connectivity. An interesting followup resides in constructing sparse k -connecting $(1 + \varepsilon, O(1))$ -remote-spanners for any $\varepsilon > 0$ and $k > 1$. Additionally, it seems possible to extend our results to edge-connectivity where we consider paths that are edge-disjoint rather than internal-node disjoint.

References

- [1] Ittai Abraham, Cyril Gavoille, and Dahlia Malkhi. On space-stretch trade-offs: Upper bounds. In *18th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 207–216, July 2006.
- [2] Baruch Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804–823, 1985.
- [3] Baruch Awerbuch, Bonnie Berger, Lenore Jennifer Cowen, and David Peleg. Fast distributed network decompositions and covers. *Journal of Parallel and Distributed Computing*, 39(2):105–114, 1996.
- [4] Surender Baswana and Telikepalli Kavitha. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In *47th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 591–602, October 2006.
- [5] Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. In *14th Symp. on Discrete Algorithms (SODA)*, pages 414–423. ACM-SIAM, January 2003.
- [6] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operation Research*, 4(3):233–235, 1979.

- [7] T. Clausen, P. Jacquet (editors), C. Adjih, A. Laouiti, P. Minet, P. Muhlethaler, A. Qayyum, and L. Viennot. Optimized link state routing protocol (OLSR). RFC 3626, October 2003. Network Working Group.
- [8] T. Clausen, P. Jacquet, A. Laouiti, P. Muhlethaler, a. Qayyum, and L. Viennot. Optimized link state routing protocol. In *IEEE INMIC*, 2001.
- [9] T. Clausen, P. Jacquet, and L. Viennot. Investigating the impact of partial topology in proactive manet routing protocols. In *The 5th Int. Symp. on Wireless Personal Multi-media Communications (WPMC)*, october 2002.
- [10] Don Coppersmith and Michael Elkin. Sparse source-wise and pair-wise distance preservers. In *16th Symp. on Discrete Algorithms (SODA)*, pages 660–669. ACM-SIAM, January 2005.
- [11] Gregory Dobson. Worst-case analysis of greedy heuristics for integer programming with nonnegative data. *Mathematics of Operations Research*, 7(4):515–531, 1982.
- [12] Devdatt Dubhashi, Alessandro Mai, Alessandro Panconesi, Jaikumar Radhakrishnan, and Aravind Srinivasan. Fast distributed algorithms for (weakly) connected dominating sets and linear-size skeletons. *Journal of Computer and System Sciences*, 71:467–479, 2005.
- [13] Michael Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms*, 1(2):283–323, October 2005.
- [14] Michael Elkin. A near-optimal fully dynamic distributed algorithm for maintaining sparse spanners. In *26th Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 195–204, August 2007.
- [15] Michael Elkin and David Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM Journal on Computing*, 33(3):608–631, 2004.
- [16] Michael Elkin and Jian Zhang. Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models. In *23rd Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 160–168, July 2004.
- [17] Arthur M. Farley, Andrzej Proskurowski, Daniel Zappala, and Kurt Windisch. Spanners and message distribution in networks. *Discrete Applied Mathematics*, 137(2):159–171, 2004.
- [18] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- [19] P. Jacquet and L. Viennot. Remote-spanners: What to know beyond neighbors. Technical Report RR-6679, INRIA, 2008.
- [20] Philippe Jacquet, Anis Laouiti, Pascale Minet, and Laurent Viennot. Performance of multipoint relaying in ad hoc mobile routing protocols. In *NETWORKING 2002*. Springer, 2002.

- [21] Philippe Jacquet, Pascale Minet, Paul Muhlethaler, and Nicolas Rivierre. Increasing reliability in cable-free radio LANs: Low level forwarding in HIPERLAN. *Wireless Personal Communications*, 4(1):65–80, January 1997.
- [22] J. McQuillan, I. Richer, and E. Rosen. The new routing algorithm for the arpanet. *IEEE Transactions on Communications*, 28(5):711–719, 1980.
- [23] John Moy. Open shortest path first (OSPF) specification. RFC 1131, October 1989. Network Working Group.
- [24] John Moy. OSPF version 2. RFC 2328, April 1998. Network Working Group.
- [25] D. Peleg and A. A. Schaffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989.
- [26] David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM Journal on Computing*, 18(4):740–747, 1989.
- [27] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36(3):510–530, July 1989.
- [28] Seth Pettie. Low distortion spanners. In *34th Int. Colloquium on Automata, Languages and Programming (ICALP)*, pages 78–89, July 2007.
- [29] Mikkel Thorup and Uri Zwick. Compact routing schemes. In *13th Annual ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 1–10, July 2001.
- [30] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM*, 52(1):1–24, January 2005.
- [31] L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.
- [32] David P. Woodruff. Lower bounds for additive spanners, emulators, and more. In *47th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 389–398, October 2006.
- [33] Jie Wu, Wei Lou, and Fei Dai. Extended multipoint relays to determine connected dominating sets in manets. *IEEE Trans. Comput.*, 55(3):334–347, 2006.