

Local Construction of Near-Optimal Power Spanners for Wireless Ad-Hoc Networks*

Iyad A. Kanj, Ljubomir Perković, and Ge Xia

Abstract—We present a local distributed algorithm that, given a wireless ad-hoc network modeled as a unit disk graph U in the plane, constructs a planar power spanner of U whose degree is bounded by k and whose stretch factor is bounded by $1 + (2 \sin \frac{\pi}{k})^p$, where $k \geq 10$ is an integer parameter and $p \in [2, 5]$ is the power exponent constant. For the same degree bound k , the stretch factor of our algorithm significantly improves the previous best bounds by Song et al. We show that this bound is near-optimal by proving that the slightly smaller stretch factor of $1 + (2 \sin \frac{\pi}{k+1})^p$ is unattainable for the same degree bound k . In contrast to previous algorithms for the problem, the presented algorithm is local. As a consequence, the algorithm is highly scalable and robust. Finally, while the algorithm is efficient and easy to implement in practice, it relies on deep insights on the geometry of unit disk graphs and novel techniques that are of independent interest.

Index Terms—Spanners, unit disk graphs, Gabriel graphs, Yao graphs, local distributed algorithms.



1 INTRODUCTION

A wireless ad-hoc network is commonly modeled as a *unit disk graph* in the two dimensional Euclidian plane. The points of the unit disk graph correspond to the mobile wireless devices and its edges connect pairs of points whose corresponding devices are in each other's transmission range. The network is usually assumed to be homogenous in the sense that all the devices have the same transmission range equal to one unit. Each edge is associated with a power or energy cost to support the corresponding link in the network. The cost is usually assumed to be the Euclidian distance between the endpoints raised to some power p , which is a constant in the interval $[2, 5]$.

Two neighboring points communicate by sending a message through their connecting edge, while distant points communicate through messages relayed by intermediate neighbors. The communication cost between two distant points is the sum of the costs of the edges on the path formed by the intermediate points. A *smallest cost path* between any pair of points is a path connecting the pair of points that has the smallest energy cost. Energy consumption is a critical issue for (battery-powered) mobile devices, and the primary goal becomes to construct a backbone topology for the network, useful for routing and other purposes, that is energy efficient.

*A preliminary version of this work appeared in proceedings of DIAL-M POMC 2007.

- Iyad A. Kanj is with the School of Computing, DePaul University, Chicago, Illinois 60604-2301. Email: ikanj@cs.depaul.edu. Supported by a DePaul University Competitive Research Grant.
- Ljubomir Perković is with the School of Computing, DePaul University Chicago, Illinois 60604-2301. Email: lperkovic@cs.depaul.edu.
- Ge Xia is with the Department of Computer Science, Lafayette College, Easton, Pennsylvania 18042. Email: gexia@cs.lafayette.edu. Supported by a Lafayette College Research grant.

There are many desirable requirements on this backbone topology. We list some of them below.

- Bounded degree: Because of interference and contention issues, a major requirement on the network topology is that each device maintains links to only a constant number of devices in its transmission range. This will also allow the devices to attenuate their transmission power to levels required to reach the selected devices only.
- Planarity: The network topology should be amenable to guaranteed and efficient routing. The folklore “right hand rule” (in face routing), discussed in [1], is one of many routing rules that require the network to be planar.
- Energy efficiency: For each pair of points, the backbone topology should contain a path connecting the two points whose cost is close or equal to the cost of a smallest cost path connecting the pair in the original network.
- Scalability/Robustness: The network topology should react to the growth/change in the topology of the underlying network in a graceful and controlled manner. Compared to that of other types of communication networks, the topology of wireless ad hoc networks is highly volatile, changing rapidly and unpredictably due to the movement of points and a possible adverse environment. For example, when points enter/exit the network, or move in the plane, the backbone topology should be maintained without widespread and drastic changes.

The requirements on the algorithm constructing the backbone topology is for it to be distributed, simple, and *local*. The notion of locality has been defined by Linial [2], Peleg [3], and Wattenhofer [4]. Intuitively, a distributed algorithm is said to be *q-local* if the compu-

tation at each point of the network depends solely on the initial states/information of the points at distance at most q from the point (i.e., within q hops from the point). More formally, a distributed algorithm is q -local if it can be simulated to run in at most q synchronous communication rounds for some integer $q > 0$ [2], [3], [4]. A distributed algorithm is called *local* if it is q -local for some integer constant q . This definition captures the notion that a local distributed algorithm should be free from centralized control. Obviously, a local distributed algorithm is much more desirable for a wireless ad-hoc network because it allows each point to perform its operations independently and simultaneously.

The corresponding topology control problem consists therefore of finding a low cost constant degree planar subgraph of the unit disk graph using a local distributed algorithm. This subgraph is referred to as a *power spanner* of the unit disk graph, and the low cost requirement can be quantified as follows: a subgraph is a power spanner with stretch factor ρ if the cost of the smallest cost path in the subgraph between any pair of points is at most ρ times the cost of the smallest cost path in the original graph itself. The problem of computing efficient topologies for wireless ad-hoc networks was extensively studied in various settings (see for instance [5], [6], [7], [8], [9], [10], [11], [12], [13]). Moreover, Scheideler gives an excellent survey on power spanners and their use in wireless ad-hoc networks [14].

For the problem of constructing power spanners of unit disk graphs, Wattenhofer et al. [13] derived algorithms with arbitrarily small stretch factor but unbounded degree. To bound the degree, their stretch factor needs to be at least 2. Song, Wang, Li, and Frieder proposed distributed “localized” algorithms [5] with guarantees on the maximum degree and stretch factors. Given a parameter $k > 6$, their first algorithm (referred to as OrdYaoGG) computes a power spanner of maximum degree $k + 5$ and maximum stretch factor $\rho = 1/(1 - (2 \sin \frac{\pi}{k})^p)$. In their second algorithm (referred to as SYaoGG), given a parameter $k \geq 9$, they obtain a bound of k for the maximum degree and $\rho = (\sqrt{2})^p/(1 - (2\sqrt{2} \sin \frac{\pi}{k})^p)$ for the stretch factor.

Even though the afore mentioned algorithms were “localized” in the sense that each point only communicates with its neighbors, a close examination shows that these algorithms process the points in some implicit order and require information propagation in the network (see Remark 2.3 for more details). In particular, in these algorithms the computation of a point may depend on the initial state of points that are far away (more than a constant number of hops away). Therefore, these algorithms are not local distributed algorithms according to the definition given by [2], [3], [4] (described above). In fact, a localized but not local distributed algorithm may not differ much from a centralized algorithm since it can simulate a centralized algorithm by collecting information through propagation and processing it at a certain point/node in the network.

TABLE 1
Degree bounds (Δ) and stretch factors (ρ) of the algorithms with power exponent $p = 2$.

$k =$	10		11		12		13		14	
	Δ	ρ	Δ	ρ	Δ	ρ	Δ	ρ	Δ	ρ
Lower Bounds	10	1.32	11	1.27	12	1.23	13	1.20	14	1.18
KPX	10	1.39	11	1.32	12	1.27	13	1.23	14	1.20
OrdYaoGG	15	1.62	16	1.47	17	1.37	18	1.30	19	1.38
SYaoGG	10	8.48	11	5.48	12	4.31	13	3.70	14	3.32

In this paper, we present a local distributed algorithm (referred to as KPX) that constructs a power spanner of a unit disk graph. The backbone/spanner and the algorithm have the following properties.

- Bounded degree: For any parameter $k \geq 10$, the backbone constructed by the algorithm has maximum degree $\Delta = k$.
- Planarity: The backbone constructed is a planar graph.
- Energy efficiency: The stretch factor of the spanner is bounded by $\rho = 1 + (2 \sin \frac{\pi}{k})^p$. For the same degree bound, this bound on the stretch factor significantly improves the previous best bounds by Song et al. [5] (see Table 1 for a comparison between these bounds). Furthermore, we show that this stretch factor is tight by proving that constructing a power spanner of a unit disk graph with degree bound of k and a stretch factor smaller than $\rho = 1 + (2 \sin \frac{\pi}{k+1})^p$ is not possible.
- Locality: The computation performed by any point in the unit disk graph is solely dependent on the coordinates of the point itself and the coordinates of its neighbors. As a matter of fact, the distributed algorithm for constructing the backbone can be implemented using 2 synchronous communication rounds, and hence is a 2-local distributed algorithm.
- Simplicity: Although our proofs rely on sophisticated analysis of the geometry of unit disk graphs, the algorithm is very simple and uniform. In particular, the storage space used by each point is proportional to its degree in the graph, and the total number of exchanged messages during the whole computation is proportional to the number of edges in the unit disk graph.
- Scalability/Robustness: Due to its local nature, our algorithm is highly scalable and robust. The communication complexity grows as a linear function in the size of the network. When the topology of the underlying network changes, the optimal power spanner can be maintained without affecting the points beyond the vicinity of the changes because the computation performed by any point is solely dependent on the coordinates of the point itself and those of its neighbors.

Moreover, for a unit disk graph of n points and m edges, the algorithm exchanges no more than $O(m)$

messages and has a local processing time of $O(\Delta \lg \Delta) = O(n \lg n)$ at a point of degree Δ .

The paper is organized as follows. Section 2 reviews the necessary definitions and background. In section 3 we define the notion of a *generalized Gabriel graph* of a unit disk graph and prove some structural properties about generalized Gabriel graphs. In Section 4 we define the notion of a *canonical path* between a pair of points in a generalized Gabriel graph. In Section 5 we present the algorithm. We conclude the paper in Section 7 by comparing our results to the previous results in the literature.

2 PRELIMINARIES

A wireless network consists of a set of n points in the two dimensional Euclidian plane. Each point has a transmission range of one unit; in other words, two points A and B can transmit to each other if their Euclidian distance, denoted by $|AB|$, is at most 1 unit. It is assumed that each point knows its coordinates through a Global Position System (GPS). A *unit disk graph* U is therefore defined on the n points as follows: for every two points A and B , AB is an edge in U if and only if $|AB| \leq 1$. The edge AB is embedded in the plane as the straight line segment AB . The unit disk graph U is assumed to be connected. The power required to support a link/edge AB in U is commonly assumed to be $|AB|^p$, where p is a constant in the interval $[2, 5]$. Two far apart points A and B communicate through intermediate points that form a simple path $A = M_0, M_1, \dots, M_r = B$ in U . The energy cost of this path is:

$$\sum_{j=0}^{r-1} |M_j M_{j+1}|^p.$$

Among all paths between A and B , a path in U with the smallest energy cost is defined to be a *smallest cost path* and we denote its cost as $c_U(A, B)$. A subgraph H of U is a *power spanner* if there is a constant ρ such that for every two points $A, B \in U$ we have: $c_H(A, B) \leq \rho c_U(A, B)$. The constant ρ is called the *stretch factor* of H . The following lemma is from [5]:

Lemma 2.1 ([5]): A subgraph H of graph U has stretch factor ρ if and only if for every edge $AB \in U$, $c_H(A, B) \leq \rho c_U(A, B) = \rho |AB|^p$.

In this paper we present an algorithm that constructs a bounded degree planar power spanner of U with a very small power stretch factor ρ . Much of the previous—and our current—work on bounded degree planar power spanners is based on the concepts of Gabriel and Yao subgraphs. We review these concepts next.

The *Gabriel subgraph* G of a unit disk graph U (embedded in the plane) is obtained by removing every edge $AB \in U$ such that there is a point $M \in U$, $M \notin \{A, B\}$, with $|MA|^2 + |MB|^2 \leq |AB|^2$, i.e., M is contained in the closed disk of diameter AB ([15]). The following properties were shown in [5].

Proposition 2.2 ([5]): Let U be a unit disk graph and let G be the Gabriel subgraph of U .

1. G is planar.
2. G is connected.
3. The power stretch factor of G is 1.
4. If AB is an edge in U , then $AB \in G$ if and only if for every point M in U the angle $\angle AMB$ in the interior of triangle $\triangle AMB$ is acute.

The Yao subgraph [16] with integer parameter $k > 6$ of a plane graph is constructed by repeating the following step for every point M : k equally separated rays out of M are arbitrarily defined, and k closed cones of size $2\pi/k$ are thus created; then, in each cone, the shortest edge MN inside the cone (if any) is chosen and added to the Yao subgraph.¹

Remark 2.3: Song et al. [5] applied a Yao subgraph construction to a Gabriel graph G . In order to bound the maximum degree, they first (implicitly) oriented the edges of the Gabriel graph G (using the classical acyclic orientation of a planar graph) so that every point in G has in-degree at most 5. Then, they applied the above Yao step to every point of G but to the *outgoing edges only*. The subgraph G' thus obtained has then maximum degree $k+5$. This graph orientation requires the points in the network to be processed in a certain global order, and hence their algorithm is *not* a local distributed algorithm. Furthermore, the graph orientation results in the “+5” issue in the degree bound.

In this paper, we overcome the above hurdles by developing a simple, uniform, and local distributed algorithm through a set of novel techniques. We start first by introducing the notion of a *generalized Gabriel graph*.

3 GENERALIZED GABRIEL GRAPHS

Let U be a unit disk graph and let $p \in [2, 5]$ be a constant.

Definition 3.1: The *generalized Gabriel graph* of U with parameter p , denoted G^p , is defined to be the subgraph of U having the same point set as U , and such that an edge $AB \in U$ is also an edge in G^p if and only if there does not exist a point $M \in U$, $M \notin \{A, B\}$, satisfying $|MA|^p + |MB|^p \leq |AB|^p$.

Note that the Gabriel graph of a unit disk graph U is the generalized Gabriel graph of U with parameter $p = 2$.

Let A and B be two points in U . Define the region $D^p(A, B)$ to be the set of all points M in the plane satisfying $|MA|^p + |MB|^p \leq |AB|^p$. Note that $D^2(A, B)$ is the closed disk of diameter AB .

Lemma 3.2: Let A, B, M , and M' be points in U such that M and M' are on the same side of the straight line AB . Suppose that the angles $\angle BAM$, $\angle ABM$, $\angle BAM'$, and $\angle ABM'$ are not obtuse. If $|MA|/|MB| = |M'A|/|M'B|$ and $\angle AM'B \geq \angle AMB$ then $M' \in \triangle AMB$.

Proof: Let $\alpha = \angle BAM$, $\beta = \angle ABM$, $\alpha' = \angle BAM'$, and $\beta' = \angle ABM'$. Consider the triangle $\triangle AMB$. From

1. The requirement $k > 6$ is to ensure connectivity.

elementary trigonometry we have:

$$|MA|/|MB| = \sin \beta / \sin \alpha. \quad (1)$$

Similarly, by considering the triangle $\triangle AM'B$ we have:

$$|M'A|/|M'B| = \sin \beta' / \sin \alpha'. \quad (2)$$

From the hypothesis, we have $|MA|/|MB| = |M'A|/|M'B|$. Combining this equality with Equalities (1) and (2) above we obtain:

$$\sin \beta / \sin \alpha = \sin \beta' / \sin \alpha'. \quad (3)$$

Since the angles $\alpha, \alpha', \beta, \beta'$ are not obtuse, it follows from Equality (3) that either $\alpha' \leq \alpha$ and $\beta' \leq \beta$, or $\alpha' \geq \alpha$ and $\beta' \geq \beta$. By the hypothesis, $\angle AM'B \geq \angle AMB$ and hence the proposition $\alpha' \geq \alpha$ and $\beta' \geq \beta$ can be ruled out. Therefore, we must have $\alpha' \leq \alpha$ and $\beta' \leq \beta$. Since both M and M' lie on the same side of the line AB , it follows that M' is inside (or on) the triangle $\triangle AMB$ as claimed. \square

Lemma 3.3: Let A and B be two points in U . The following are true.

1. $D^p(A, B)$ is a closed convex set.
2. $D^p(A, B)$ is symmetric with respect to the line AB , and every point M in $D^p(A, B)$ lies between the two lines/strips passing through A and B and perpendicular to AB . In particular, every point M in $D^p(A, B)$ satisfies $\angle ABM < \pi/2$ and $\angle BAM < \pi/2$. Moreover, both A and B lie on the boundary of $D^p(A, B)$.
3. $D^p(A, B)$ contains $D^2(A, B)$.
4. Every point $M \in D^p(A, B)$ satisfies $\angle AMB \geq \arccos(1 - 2^{-3/5})$.
5. Let M_1 and M_2 be two points in U such that $|M_1A|/|M_1B| \geq |M_2A|/|M_2B| \geq 1$, and $\angle AM_2B \geq \angle AM_1B$. If $M_1 \in D^p(A, B)$ then $M_2 \in D^p(A, B)$.

Proof: Part 1: $D^p(A, B)$ is closed because it contains all its boundary points. To show that $D^p(A, B)$ is convex, assume without loss of generality that $A = (-a/2, 0)$ and $B = (a/2, 0)$, where $a = |AB|$ is a positive constant. This can be assumed by affecting the appropriate rotation and translation, which are convexity preserving. Then the set of points $D^p(A, B)$ can be described by $D^p(A, B) = \{(x, y) \mid ((x + a/2)^2 + y^2)^{p/2} + ((x - a/2)^2 + y^2)^{p/2} \leq a^p\}$. Let L be any straight line with equation $y = bx + c$ and define the function:

$$f(x) = ((x+a/2)^2 + (bx+c)^2)^{p/2} + ((x-a/2)^2 + (bx+c)^2)^{p/2}.$$

The function f measures the value $|MA|^p + |MB|^p$ for an arbitrary point $M = (x, y)$ on L . It can be verified by the interested reader that $f''(x)$ is non-negative, and hence f is a convex function. This implies that for any two points C and D , and any point M on the line segment CD , the value of f at M does not exceed both values of f at C and D . In particular, if C and D are two points in $D^p(A, B)$, then every point on the line segment CD is also in $D^p(A, B)$, and $D^p(A, B)$ is a convex set.

Part 2: If a point M satisfies $|MA|^p + |MB|^p \leq |AB|^p$ then its symmetry M' with respect to the line AB satisfies $|M'A|^p + |M'B|^p \leq |AB|^p$ as well, and hence is in $D^p(A, B)$. This shows that $D^p(A, B)$ is symmetric with respect to AB . Now if a point M lies outside the two lines passing through A and B and perpendicular to AB , then the distance between M and either A or B is larger than $|AB|$, and hence M cannot satisfy $|MA|^p + |MB|^p \leq |AB|^p$, and $M \notin D^p(A, B)$. This also implies that for a point $M \in D^p(A, B)$ we have $\angle ABM < \pi/2$ and $\angle BAM < \pi/2$. The statement that both A and B are on the boundary of $D^p(A, B)$ follows from the fact that they both satisfy the equation of the boundary curve of $D^p(A, B)$, namely $|MA|^p + |MB|^p = |AB|^p$.

Part 3: Consider an arbitrary point $M \in D^2(A, B)$. Then $|MA|^2 + |MB|^2 \leq |AB|^2$, and we have $|MA|, |MB| \leq |AB|$. For $p \geq 2$, we have $|MA|^p + |MB|^p \leq |MA|^2|AB|^{p-2} + |MB|^2|AB|^{p-2} \leq (|MA|^2 + |MB|^2)|AB|^{p-2} \leq |AB|^p$. This implies that $M \in D^p(A, B)$, and $D^2(A, B) \subseteq D^p(A, B)$.

Part 4: Let M_0 be a point in $D^p(A, B)$ that minimizes $\angle AM_0B$. Suppose that $|M_0A| = x$, $|M_0B| = y$, $\angle AM_0B = \alpha$ (the internal angle in the triangle $\triangle AM_0B$), and $|AB| = a$. Considering the triangle $\triangle AM_0B$ we can write:

$$\cos(\alpha) = (x^2 + y^2 - a^2)/2xy = x/2y + (y^2 - a^2)/2xy. \quad (4)$$

Since $M_0 \in D^p(A, B)$, we have $x^p + y^p \leq a^p$. Since $x, y \geq 0$, it follows that $x, y \leq a$, and the term $(y^2 - a^2)$ in Equation (4) is negative. This shows that, for a fixed value y , $\cos(\alpha)$ increases with x . Noting that $0 \leq \alpha \leq \pi$, it follows from all the previous facts that the largest value for $\cos(\alpha)$, and hence the smallest value for α , is attained when $x = (a^p - y^p)^{1/p}$. In this case we have:

$$\cos(\alpha) = ((a^p - y^p)^{\frac{2}{p}} + y^2 - a^2)/2(a^p - y^p)^{\frac{1}{p}}y. \quad (5)$$

By studying the variation of $\cos(\alpha)$ as a function of y , we can show that $\cos(\alpha)$ is maximum when $y = a/2^{\frac{1}{p}} = x$, in which case we have:

$$\cos(\alpha) = 1 - 2^{\frac{2}{p}-1} \leq 1 - 2^{-3/5}, \quad (6)$$

for $p \in [2, 5]$, and hence $\alpha \geq \arccos(1 - 2^{-3/5})$. It follows that for every point $M \in D^p(A, B)$ we have $\angle AMB \geq \arccos(1 - 2^{-3/5})$.

Part 5: Define the function h on the interval $[1, \infty)$ as follows. For a value $\sigma \in [1, \infty)$, let M be a point on the boundary of $D^p(A, B)$ such that $|MA|/|MB| = \sigma$, and define $h(\sigma) = \angle AMB$. We first need to show that h is a well-defined function.

Let σ be a given fixed value in $[1, \infty)$. Let M be a point on the boundary of $D^p(A, B)$ and suppose that $|MA| = x$, $|MB| = y$, and $|AB| = a$. Then we have $x^p + y^p = a^p$. If we let $x/y = \sigma$ and solve for x and y in the equation $x^p + y^p = a^p$, we get $x = a\sigma/(1 + \sigma^p)^{1/p}$ and $y = a/(1 + \sigma^p)^{1/p}$. Therefore, there are two points M_0 and M'_0 , symmetrical with respect to AB , which are the intersection of the circle centered at A (resp. B) and

of radius $a\sigma/(1 + \sigma^p)^{1/p}$ (resp. $a/(1 + \sigma^p)^{1/p}$) with the boundary of $D^p(A, B)$. (It can be easily verified that the intersection is precisely a set of two points.) The value $h(\sigma)$ is hence $h(\sigma) = \angle AM_0B = \angle AM'_0B$ (since M_0 and M'_0 are symmetrical with respect to AB), which exists and is unique. It follows that h is a well-defined function.

For $\sigma \in [1, \infty)$, consider a point M on the boundary of $D^p(A, B)$ such that $|MA| = x$, $|MB| = y$, and $\sigma = x/y$. Note that, by the definition of h , $h(\sigma) = \angle AMB$. By letting $\alpha = \angle AMB$, and by a similar analysis to that in part 4 above, we can study the variation of the function $\cos(\alpha)$ given in Equation (5) as a function of σ , by replacing y by its value $y = a/(1 + \sigma^p)^{1/p}$ in terms of σ . We can show that, for $\sigma \in [1, \infty)$, the function $\cos \alpha$ decreases with σ , and hence $h(\sigma) = \alpha = \angle AMB$ increases with σ .

Now let M_1 and M_2 be two points of U satisfying the conditions in part 5 of the lemma. Let M'_1 be the point on the boundary of $D^p(A, B)$ such that $\angle AM'_1B = h(|M_1A|/|M_1B|)$, and such that M_1 and M'_1 are on the same side of AB . By the definition of h , we have $|M_1A|/|M_1B| = |M'_1A|/|M'_1B|$. Since M_1 and M'_1 are in $D^p(A, B)$, by part 1 of this lemma, we have $\angle AM_1B < \pi/2$, $\angle BM_1A < \pi/2$, $\angle AM'_1B < \pi/2$, and $\angle BM'_1A < \pi/2$. If $\angle AM'_1B > \angle AM_1B$, then by Lemma 3.2 M'_1 lies strictly inside the triangle $\triangle AM_1B$. Since $M_1 \in D^p(A, B)$ and $D^p(A, B)$ is convex, all points in the triangle $\triangle AM_1B$ are in $D^p(A, B)$. But M'_1 is a boundary point lying strictly inside $\triangle AM_1B$, a contradiction. It follows that $\angle AM_1B \geq \angle AM'_1B = h(|M_1A|/|M_1B|)$.

Since $\angle AM_2B \geq \angle AM_1B$ and $|M_1A|/|M_1B| \geq |M_2A|/|M_2B|$ by the hypothesis, and since h is increasing, it follows that $\angle AM_2B \geq \angle AM_1B \geq h(|M_1A|/|M_1B|) \geq h(|M_2A|/|M_2B|)$. Therefore, $\angle AM_2B \geq h(|M_2A|/|M_2B|)$. Now by letting M'_2 be the point on the boundary of $D^p(A, B)$ such that $\angle AM'_2B = h(|M_2A|/|M_2B|)$ and such that both M_2 and M'_2 are on the same side of AB , we can conclude by a similar token to the above, that M_2 lies inside $\triangle AM'_2B$. By the convexity of $D^p(A, B)$, $M_2 \in D^p(A, B)$ as claimed. \square

The following proposition on the structural properties of generalized Gabriel graphs is parallel to Proposition 2.2.

Proposition 3.4: Let U be a unit disk graph, $p \in [2, 5]$ a constant, and G^p the generalized Gabriel graph of U with parameter p . Then the following are true.

1. G^p is planar.
2. G^p is a connected.
3. The power stretch factor of G^p is 1.
4. If an edge $AB \in U$ is not an edge in G^p , then there exists a point $M \in U$ such that $\angle AMB \geq \alpha_0$, where $\alpha_0 = \arccos(1 - 2^{-3/5})$; on the other hand, if AB is an edge in U and there exists a point $M \in U$ such that $\angle AMB \geq \pi/2$, then AB is not an edge in G^p .

Proof: The proof of parts 1, 2, and 3 are similar to those for the case of Gabriel graphs, which can be found in the literature (for example, see [17], [8]).

If an edge $AB \in U$ is not an edge in G^p , then there exists a point $M \in U$, $M \neq A, B$, such that $|MA|^p + |MB|^p \leq |AB|^p$, and hence $M \in D^p(A, B)$. By part 4 of Lemma 3.3, $\angle AMB \geq \alpha_0 = \arccos(1 - 2^{-3/5})$. On the other hand, if AB is an edge in U and there exists a point $M \in U$ such that $\angle AMB \geq \pi/2$, then $M \in D^2(A, B)$. By part 3 of Lemma 3.3, $M \in D^p(A, B)$ as well, and $|MA|^p + |MB|^p \leq |AB|^p$. By the definition of generalized Gabriel graphs, AB is not an edge in G^p . \square

Remark 3.5: The results in this section describe the structure and the properties of the regions $D^p(A, B)$, where $A, B \in U$, and $p \in [2, 5]$, as well as the properties of generalized Gabriel graphs. The region $D^p(A, B)$ is a closed convex set of points which is symmetric with respect to AB and which encloses the disk of diameter AB . Unfortunately, a point $M \in U$ cannot be classified with respect to $D^p(A, B)$ solely based on the angle $\angle AMB$ as is the case with the disk of diameter AB (i.e., $D^2(A, B)$). This leads to a major difference between Gabriel graphs and generalized Gabriel graphs with parameter $p > 2$. The existence of an edge $AB \in U$ in the Gabriel graph G of U can be precisely characterized as follows: An edge $AB \in U$ is in G if and only if there does not exist any point $M \in U$ such that $\angle AMB \geq \pi/2$. In contrast, for the case of generalized Gabriel graphs such nice characterization does not exist, as one may have noticed from part (iv) of Proposition 3.4. This is mainly due to the fact that Gabriel graphs correspond to generalized Gabriel graphs with parameter value $p = 2$, and hence the curve describing the set of points M satisfying $|MA|^2 + |MB|^2 = |AB|^2$ is a circle, which could be described precisely as the set of points M such that $\angle AMB = \pi/2$ (of course, in addition to the two points A and B). For parameter values $p > 2$, this is no longer the case. The set of points M satisfying $|MA|^p + |MB|^p = |AB|^p$ for $p > 2$, can no longer be precisely described based on the angle $\angle AMB$. Therefore, dealing with generalized Gabriel graphs becomes much more complicated than dealing with Gabriel graphs, and generalized Gabriel graphs do not possess all the nice properties of Gabriel graphs.

We close this section by showing how the generalized Gabriel graph of a unit disk graph can be constructed efficiently by a local distributed algorithm.

Theorem 3.6: Let U be a unit disk graph on n points and m edges, and let $p \in [2, 5]$ be a constant. The generalized Gabriel graph G^p of U can be constructed by a local distributed algorithm exchanging $O(m)$ messages and with a local processing time of $O(\Delta \lg \Delta) = O(n \lg n)$ at a point with degree Δ .

Proof: Suppose that the parameter p is fixed. We say that a point $M \in U$ kills an edge XY , or alternatively XY is killed by M , if $|MX|^p + |MY|^p \leq |XY|^p$, that is, if M is within the region $D^p(X, Y)$. As a first step, the algorithm starts by constructing the Gabriel graph G^2 of U . For that, each point A in U does the following. It sorts its neighbors in a non-decreasing order of their polar angles (with respect to itself). Let $\langle A_1, \dots, A_k \rangle$ be

the resulting list. Point A also creates a corresponding ordered list of edges $\mathcal{L} = \langle AA_1, \dots, AA_k \rangle$. Note that the edges in \mathcal{L} appear in a counterclockwise order around A starting with the edge AA_1 , which has the smallest polar angle.

Point A starts by considering point A_1 . Then it scans \mathcal{L} starting at edge AA_2 removing from \mathcal{L} every edge that is killed by A_1 . Then A does the same going backward starting with edge AA_k . The first traversal of \mathcal{L} going forward corresponds to a counterclockwise traversal of the edges in \mathcal{L} starting at the first edge after AA_1 in \mathcal{L} , while the second traversal going backward corresponds to a clockwise traversal of the edges in \mathcal{L} starting with the first edge before AA_1 . Each traversal stops when it encounters the first edge that is not killed by A_1 . After that, A considers point A_2 and so on. In general, at step i , A considers point A_i and continues its traversals (from where they were left) of the remaining edges in \mathcal{L} , both in clockwise and counterclockwise order, removing from \mathcal{L} every edge killed by A_i . The process stops when the last vertex A_k is considered and processed. After that, point A chooses the endpoints of the remaining edges in \mathcal{L} to be its neighbors in G^2 .

To prove that the above algorithms is correct, it suffices to show that each point A chooses an edge if and only if its a Gabriel incident edge. Since A eliminates edges only if they are killed by some neighboring point, no edge eliminated by A can be a Gabriel edge. We show the converse next.

Let AA_r be an edge in U that is not a Gabriel edge. We need to show that AA_r will be removed from \mathcal{L} at a certain point. First note that, by the definition of Gabriel graphs, for any edge in U that is not a Gabriel edge, there must exist a point in U that kills it. Moreover, any point M that kills a Gabriel edge XY must be a common neighbor (in U) of both X and Y . Since AA_r is not a Gabriel edge, there must exist a neighbor of A that kills it. Let A_i be a neighbor of A that kills AA_r such that the smaller of the two angles (clockwise and counterclockwise) between AA_i and AA_r is minimized over all neighbors of A that kill AA_r . Let α be the angular sector between AA_i and AA_r , and assume, without loss of generality, that α is counterclockwise. Since A_i kills AA_r , and by our choice of A_i , the measure of the angular sector α is a number in $[0, \pi)$. If $\alpha = 0$, then we assume that A_i is chosen so that $|AA_i|$ is the shortest among all points A_i satisfying $\angle AA_i, AA_r = \alpha$ and A_i kills AA_r . Now we claim that A_i kills every edge counterclockwise from AA_i lying within the angular sector α . In effect, let AA_j be such an edge. If $\angle AA_i, AA_j = 0$, then by the choice of A_i , we have $|AA_i| \leq |AA_j|$ and A_i lies on the segment AA_j . Therefore, $A_i \in D^2(A, A_j)$ and A_i kills AA_j . Now suppose that $\angle AA_i, AA_j > 0$. Since $\angle AA_j, AA_r < \alpha$, and by the choice of A_i , A_j does not kill AA_r . Therefore, A_j must be exterior to the triangle $\triangle AA_i A_r$; otherwise, by convexity of $D^2(A, A_r)$ (part 1 of Lemma 3.3), A_j kills AA_r (since A_i kills AA_r). Consequently, and because $\alpha < \pi$, $\angle AA_i A_j > \angle AA_i A_r \geq \pi/2$.

The latter inequality is true because A_i kills AA_r and $D^2(A, A_r)$ is the disk of diameter AA_r . It follows that A_i kills AA_j .

Therefore, when the algorithm considers the point A_i and scans \mathcal{L} in a counterclockwise order, every edge in \mathcal{L} between AA_i and AA_r inclusive, is killed by A_i and will be removed from \mathcal{L} .

This shows that the above algorithm constructs the Gabriel graph of U correctly. Scanning the list \mathcal{L} of a point A and killing edges takes time $O(\Delta)$, where Δ is the degree of A . This is true since once an edge is removed from \mathcal{L} , it will never be reconsidered. So the amortized time over all the points $\langle A_1, \dots, A_k \rangle$ considered by A is $O(\Delta)$. Therefore, the time spent by each point is dominated by the sorting phase, which takes $O(\Delta \lg \Delta)$.

Now we describe how the generalized Gabriel graph G^p can be constructed for an arbitrary value of the parameter $p \in [2, 5]$. First, the Gabriel graph is constructed as described above and each point A keeps a list of the Gabriel edges incident on it, sorted in a counterclockwise order around the A . Let this list be $\langle A_1, \dots, A_k \rangle$. (Note that since G^p is subgraph of G^2 , all the edges in G^p are present in G^2 .) Then point A repeats exactly the same procedure described above, starting with point A_1 , killing consecutive edges both in clockwise and anticlockwise order around A_1 , until this no longer can be done, then proceeding to point A_2 , and so on. The only difference here is that a point M kills an edge XY if and only if $|MX|^p + |MY|^p \leq |XY|^p$.

To show that the algorithm computes the generalized Gabriel graph correctly, we need to show that every edge AA_r that is not a generalized Gabriel edge must be removed from \mathcal{L} at a certain point. We proceed as above; the only difference here is that each edge in \mathcal{L} is a Gabriel edge. Let AA_r be a Gabriel edge that is not a generalized Gabriel edge. Then there is a neighbor of A that kills it. Let A_i be a neighbor with the properties described above, namely a neighbor minimizing the smaller of the two angles between AA_i and AA_r . Let α be the angular sector between AA_i and AA_r , and assume that α is counterclockwise. Similar to the above, our task amounts to showing that A_i kills every edge within the angular sector α . Let AA_j be such an edge. Then A_j must be exterior to the triangle $\triangle AA_i A_r$ by the choice of A_i and the convexity of $D^p(A, A_r)$ (part 1 of Lemma 3.3). We need to show that A_i kills AA_j , or equivalently, $|AA_i|^p + |A_i A_j|^p \leq |AA_j|^p$.

Since A_i kills AA_r we have:

$$|AA_i|^p + |A_i A_r|^p \leq |AA_r|^p. \quad (7)$$

From Inequality (7) we know that AA_r is the longest edge in the triangle $\triangle AA_i A_r$ and hence $\angle A_i A_r \leq \pi/2$.

Since AA_j is a Gabriel edge, by part 4 of Proposition 3.4 (applied with $p = 2$) we know that $\angle AA_i A_j$ and $\angle AA_r A_j$ are acute. By considering the quadrilateral $AA_i A_j A_r$, it follows from the above that $\angle A_i A_j A_r \geq$

$\pi/2$. By part 4 of Proposition 3.4, $A_i A_r$ is not an edge in G^p and we have:

$$|A_i A_j|^p + |A_j A_r|^p \leq |A_i A_r|^p. \quad (8)$$

By the choice of A_i , A_j does not kill AA_r and we have:

$$|AA_j|^p + |A_j A_r|^p > |AA_r|^p. \quad (9)$$

Combining Inequalities (7), (8), and (9), we derive:

$$|AA_i|^p + |A_i A_j|^p \leq |AA_j|^p,$$

and A_i kills AA_j as claimed.

This shows that the algorithm constructs G^p correctly. Observe that the information needed by a point to construct its incident edges in G^p is local: the point only needs to know its coordinates and the coordinates of its neighbors. Therefore, the above algorithm is a 1-local distributed algorithm. Noting that the local time spent by a point in constructing G^p from G^2 is dominated by the time spent by the point in constructing G^2 , and that the only messages exchanged by the algorithm are the messages in which points notify their neighbors of their coordinates (and hence the total number of messages is $O(m)$), the proof is complete. \square

4 CANONICAL PATHS

We assume in this section that G^p is the generalized Gabriel graph of U with parameter $p \in [2, 5]$. From part 3 in Proposition 3.4, G^p has a stretch factor of 1 but, as in the case with Gabriel graphs, the degree of G^p is not bounded. To construct a backbone of U satisfying the desired properties described in Section 1, we will need to bound the degree of every point in G^p by having every point in G^p choose a bounded number of incident edges, thus constructing a subgraph G' of G^p in which the degree of every point is bounded. Of course, G' will no longer have a stretch factor of 1. In fact, the stretch factor of G' cannot be smaller than $1 + (2 \sin \frac{\pi}{k+1})^p$ when the degree bound is k , as we will prove in the next section. Thus the choice of these selected edges from every point has to be done in a careful manner so that the stretch factor is still close to the above lower bound.

The proof of the following lemma is exactly the same as that of Lemma 6 in [5].

Lemma 4.1 (Lemma 6, [5]): Let A, B, C be three points in a generalized Gabriel graph G^p such that CB and CA are edges in G^p . Suppose that $|CA| \leq |CB|$ and $\angle BCA \leq \alpha$ for some $\alpha \in (0, \pi/2)$. Then $|CA|^p + |AB|^p \leq (1 + (2 \sin \frac{\alpha}{2})^p) |CB|^p$.

Our goal in this paper is to obtain a bound of $1 + (2 \sin \frac{\pi}{k})^p$ on the stretch factor. By Lemma 4.1, we can achieve this by guaranteeing the following property in G' : For any edge $CB \in G^p$ that is not chosen in G' there exists a chosen edge CA in G' such that $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$ (this condition will be relaxed a little bit, as we will explain in the next section), and such that either $AB \in G'$ or there exists a path between A and B in G' whose cost is not higher than $|AB|^p$. We will

call such a path a *canonical path* and define it using the notion of a *canonical point*.

Let A and B be two points in U such that $AB \in U$ but $AB \notin G^p$.

Definition 4.2: We define the *canonical point* for the pair of points (A, B) in U to be a point $M \in U$, $M \notin \{A, B\}$, satisfying $|MA|^p + |MB|^p \leq |AB|^p$ and minimizing the area of the triangle $\triangle AMB$.

By the definition of generalized Gabriel graphs such a point must exist.

Proposition 4.3: Let A and B be two points in U such that the edge $AB \in U$ but $AB \notin G^p$, and let M be the canonical point for the pair (A, B) . Then the following are true.

1. There is no point of U inside the triangle $\triangle AMB$.
2. Both MA and MB are edges in U .
3. $\angle AMB \geq \alpha_0$, where $\alpha_0 = \arccos(1 - 2^{-3/5})$.
4. $\angle ABM < \pi/2$ and $\angle BAM < \pi/2$.

Proof: By the convexity of $D^p(A, B)$, if there is a point M' inside $\triangle AMB$, then $|M'A|^p + |M'B|^p \leq |AB|^p$ and the area of $\triangle AM'B$ is smaller than that of $\triangle AMB$. This is a contradiction to the minimality of the area of the triangle $\triangle AMB$ by the definition of the canonical point M of (A, B) . Part 2 follows from the fact that $|MA|^p + |MB|^p \leq |AB|^p$, and hence both MA and MB are not longer than AB , which is an edge in U . Part 3 directly follows from part 4 of Lemma 3.3. Part 4 follows from the fact that AB is the longest edge in the triangle $\triangle AMB$. \square

Lemma 4.4: Let XY and XZ be two edges in G^p . Then no point of G^p lies inside the triangle $\triangle XYZ$.

Proof: Proceed by contradiction. If M is a point inside $\triangle XYZ$ then the angles $\angle XMY$, $\angle XMZ$, and $\angle YMZ$ add up to 2π , and none of them is greater than π . By part 4 of Proposition 3.4, the two angles $\angle XMY$ and $\angle XMZ$ are acute forcing the angle $\angle YMZ$ to be larger than π . \square

Definition 4.5: Let CA and CB be edges in G^p . We now define the *canonical path* P_{AB} between $X = A$ and $Y = B$ to be a path in G^p constructed recursively as follows: if $XY \in G^p$ then edge XY is returned, otherwise the concatenation of the canonical paths from X to M and from M to Y is returned, where M is the canonical point for (X, Y) .

Theorem 4.6: The cost of the canonical path between A and B is at most $|AB|^p$.

Proof: This follows inductively from the definition of a canonical path. \square

Let CA and CB be edges in G^p such that $AB \in U$, and let P_{AB} be the canonical path between A and B . From the definition of P_{AB} , the recursive construction of P_{AB} can be viewed as a sequence of steps $\Gamma_{AB} = \{\gamma_1, \gamma_2, \dots\}$, where step γ_i defines a canonical point M_i for a pair of points A_i, B_i , and the first step defines the canonical point M_0 for the pair $A_0 = A$ and $B_0 = B$. Let $F(C, A, B)$ be the union of $\triangle CAB$ and all the triangles defined in




Fig. 1. The figure illustrates a region $F(C, A, B)$ where P_{AB} is highlighted.

Γ_{AB} , i.e.,

$$F(C, A, B) = \triangle CAB \cup \bigcup_i \triangle A_i M_i B_i.$$

Note that $F(C, A, B)$ is a continuous region in the plane (see Fig. 1 for an illustration). If M is a point on the canonical path, we define the *angle at M* to be the angle inside the region $F(C, A, B)$ and formed by edges incident to M on the boundary of $F(C, A, B)$. We have the following proposition on the structure of $F(C, A, B)$.

Remark 4.7: If the edges CA and CB in the triangle $\triangle ABC$ are in U , and if $\angle BCA \leq 2\pi/k$, where $k \geq 10$ is an integer constant, then the edge AB is also in U being smaller than the larger edge between CA and CB , which are both in U .

Proposition 4.8: Let CA and CB be edges in G^p such that $\angle BCA = \alpha \leq 2\pi/k$, where $k \geq 10$ is an integer constant. Then the following are true. (See Figure 3 for illustration.)

1. There are no points of U in the interior of $F(C, A, B)$.
2. Every edge in G^p interior to $F(C, A, B)$ must have C as one of its endpoints.
3. The angles $\angle CAB$ and $\angle CBA$ are at least $\pi/2 - \alpha \geq \pi/2 - 2\pi/k$.
4. Each angle at any canonical path point is at least α_0 , where $\alpha_0 = \arccos(1 - 2^{-3/5})$.
5. If there is an edge $CM \in G^p$ inside $F(C, A, B)$ from C to a point M on the canonical path such that CM is not exterior to $F(C, A, B)$, then the angle at M is at least $\pi - \alpha \geq \pi - 2\pi/k$, and if L and N are the points adjacent to M on the canonical path P_{AB} , then each of the angles $\angle CML$ and $\angle CMN$ is at least $\pi/2 - \alpha \geq \pi/2 - 2\pi/k$.

Proof: For part 1, note that by Lemma 4.4 there are no points of U inside triangle $\triangle ABC$. It follows by an inductive argument from part 1 of Proposition 4.3 that there are no points in the interior of $F(C, A, B)$.

For part 2, note first that if $F(C, A, B)$ contains an edge in G^p that does not connect C to a point of the canonical path, then this edge must connect two points $P, Q \neq C$ that are nonadjacent on the boundary of $F(C, A, B)$. We will show that this leads to a contradiction.

If the sequence of steps Γ_{AB} defining the canonical path P_{AB} is empty, then P_{AB} consists of the edge AB




Fig. 2. PQ intersects one (left) or two (right) edges of $\triangle ZVW$.

and the statement is vacuously true. Suppose now that Γ_{AB} is nonempty.

Since $PQ \in F(C, A, B)$, there must exist a first step in the sequence Γ_{AB} with a corresponding pair of points (Z, V) whose canonical point is W , and such that PQ intersects triangle $\triangle ZVW$. Since (Z, V) is the first pair in this sequence of steps with this property, PQ does not intersect ZV . Therefore, PQ either intersects ZW only, VW only, or both ZW and VW . We will first consider the case when PQ intersects precisely one edge, say VW .

Since W is the canonical point of (Z, V) , by part 1 of Proposition 4.3, no point of G^p is inside $\triangle ZVW$. By assumption, PQ intersects VW , this forces one endpoint of PQ to be Z . Suppose now that $P = Z$. (See the left side of Fig. 2 for an illustration).

Since $PQ \in G^p$, by part 4 of Proposition 3.4, we have $\angle ZWQ < \pi/2$ and $\angle ZVQ < \pi/2$. Since W is the canonical point of (Z, V) , by part 4 of Proposition 4.3, we have $\angle WZV < \pi/2$. Since the sum of the four angles inside the quadrilateral $ZVQW$ is 2π , it follows that $\angle VQW > \pi/2$, and hence by part 4 of Proposition 3.4, we have:

$$|VQ|^p + |WQ|^p \leq |WV|^p. \quad (10)$$

Also, since $PQ \in G^p$ we have:

$$|PQ|^p < |WQ|^p + |WP|^p. \quad (11)$$

Because $P = Z$, and from Inequalities (10) and (11) we have:

$$\begin{aligned} |VQ|^p + |ZQ|^p &< |WZ|^p + |WV|^p \\ &< |ZV|^p. \end{aligned} \quad (12)$$

On the other hand, since $\angle ZWQ < \pi/2$ and $\angle WZV < \pi/2$, the point Q is closer to the line ZV than W , and hence the area of $\triangle ZQV$ is smaller than the area of $\triangle ZWV$, a contradiction to the minimality of the area of $\triangle ZWV$ by the definition of a canonical point.

Now let us consider the case when PQ intersects both ZW and VW (see the right side of Fig. 2). Since the angles $\angle WZV$ and $\angle WVZ$ are acute, both P and Q must be on the same side of ZV ; otherwise, either $\angle PZQ$ or $\angle PVQ$ is not acute and $PQ \notin G^p$. With proper renaming of P and Q , and Z and V , we can assume that the intersection of PQ and ZW is closer to P than Q , and

that $|PW| \geq |WQ|$. We leave the verification of this statement to the interested reader.

Similar to the above, we can verify that the areas of $\triangle ZPV$ and $\triangle ZQV$ are smaller than that of $\triangle ZWV$. We will prove next that either $|VP|^p + |ZP|^p \leq |ZV|^p$ or $|VQ|^p + |ZQ|^p \leq |ZV|^p$ is true, and hence derive a contradiction to the minimality of the area of $\triangle ZWV$.

We will first prove that both $\angle ZPW$ and $\angle VQW$ are larger than $\pi/2$. If $\angle ZQW \geq \pi/2$, then $|ZQ| < |ZW|$ and because $\angle VQW > \angle ZQW$, we also have $|VQ| < |VW|$. This would imply that $|VQ|^p + |ZQ|^p < |VW|^p + |ZW|^p < |ZV|^p$, a contradiction to the minimality of the area of $\triangle ZWV$. Thus we derive that $\angle ZQW < \pi/2$. We also have $\angle PWQ < \pi/2$ and $\angle PZQ < \pi/2$ for $PQ \in G^p$. It follows that $\angle ZPW > \pi/2$ because the sum of the four angles inside the quadrilateral $ZVQW$ is 2π . Symmetrically, $\angle VQW > \pi/2$, and by part 4 of Proposition 3.4, this implies that:

$$|WQ|^p + |VQ|^p \leq |WV|^p. \quad (13)$$

Now consider the triangles $\triangle PWQ$ and $\triangle ZWQ$. Since $\angle ZPW > \pi/2$ and $|PW| \geq |WQ|$ (by assumption), we have $|ZW|/|WQ| \geq |PW|/|WQ| \geq 1$. We also have $0 \leq \angle ZWQ \leq \angle PWQ \leq \pi/2$. By part 5 of Lemma 3.3, if $|ZW|^p + |WQ|^p \leq |ZQ|^p$, then $|PW|^p + |WQ|^p \leq |PQ|^p$, a contradiction to the fact that $PQ \in G^p$. Hence:

$$|ZQ|^p < |ZW|^p + |WQ|^p. \quad (14)$$

Combining inequalities (13) and (14), we have:

$$|VQ|^p + |ZQ|^p < |ZW|^p + |WV|^p < |ZV|^p,$$

contradicting the minimality of the area of $\triangle ZWV$. This completes the proof of part 2.

Part 3 follows from the fact that $\angle CAB, \angle CBA < \pi/2$ in the triangle $\triangle ABC$ (by part 4 of Proposition 3.4 because $CA, CB \in G^p$).

For part 4, the fact that each angle at a canonical path point is at least α_0 follows from the definition the canonical path by an inductive argument.

To prove part 5, suppose that there is an edge $CM \in G^p$ joining C to a point M on the canonical path between A and B such that CM is not exterior to $F(C, A, B)$. Since all the boundary edges of $F(C, A, B)$ are edges in G^p , and since $CM \in G^p$, then by the planarity of G^p (Proposition 3.4), the edge CM must lie completely within the face $F(C, A, B)$. Let N and L be the two neighbors of M on the boundary of $F(C, A, B)$, and suppose, without loss of generality, that A lies on the path joining C to N on the boundary of $F(C, A, B)$, and B on that joining C to L . Observe that the line segment joining N to C must lie entirely within $F(C, A, B)$, otherwise, the point A would be interior to the triangle $\triangle CNM$ contradicting Lemma 4.4—since both edges CM and MN of $\triangle CNM$ are in G^p . (Note that by the planarity of G^p , the edges CA and MN of G^p do not cross, and hence if A is not inside $\triangle CNA$ then it cannot be inside the angular sector $\angle NCM$.) Similarly, the line segment

joining L to C must lie entirely within $F(C, A, B)$. It follows from this fact that $\angle LCN \leq \alpha$. Now in the triangle $\triangle MNC$, we have $\angle CNM < \pi/2$ by Proposition 3.4 (since $CM \in G^p$), and hence $\angle CMN > \pi/2 - \alpha$. Similarly, we have $\angle CLM < \pi/2$, and $\angle CML > \pi/2 - \alpha$. Now in the quadrilateral $CNML$ we have $\angle NML = 2\pi - \angle LCN - \angle CNM - \angle CLM \geq 2\pi - \alpha - \pi/2 - \pi/2 = \pi - \alpha$. \square




Fig. 3. Illustration for Proposition 4.8.

In order to guarantee that if the edge $AB \notin G'$ (recall that G' is the desired backbone, which is a subgraph of G^p) then all the edges on the canonical path between A and B are in G' , we introduce certain structures called *sectors* to distinguish the edges on the canonical path.

Definition 4.9: Let A be a point in U and let $k \geq 10$ be an integer constant. Define an $S_{\geq 2}$ sector around A to be a maximal sequence of $\ell \geq 2$ consecutive edges (called $S_{\geq 2}$ sector edges) incident on A such that the angle between the first and the last is at least $(\ell - 1)\alpha_0$, where $\alpha_0 = \arccos(1 - 2^{-3/5}) \geq 0.389\pi$, and the angle between every two consecutive edges in the sequence is at least $\pi/2 - 2\pi/k$. Define an S_1 sector around A to be a pair of consecutive edges incident on A such that (1) the angle between them is at least $\pi/2 - 2\pi/k$, (2) one of the edges is strictly shorter than the other, and (3) the shorter edge (called an S_1 sector edge) does not belong to any $S_{\geq 2}$ sector. If an edge is an $S_{\geq 2}$ sector edge or an S_1 sector edge, we call it a *sector edge*.

The following lemma guarantees that edges on canonical paths are sector edges for both their endpoints.

Lemma 4.10: Let CA and CB be edges in G^p such that $\angle BCA \leq \alpha$, where $\alpha = 2\pi/k$ and $k \geq 10$ is an integer constant. Then each edge WV on the canonical path P_{AB} between A and B is a sector edge for both W and V .

Proof: Consider an arbitrary edge WV on the canonical path between A and B . We will show that WV is either an $S_{\geq 2}$ sector edge or an S_1 sector edge of W . By symmetry, the same proof will carry for V . We distinguish the following cases.

Case 1: W is A or B . First, suppose that $W = B$. Since $CB \in G^p$, by part 3 of Proposition 4.8, $\angle CBA \geq \pi/2 - 2\pi/k$. If the angle $\angle CBV(\angle CWV)$ is less than $\angle CBA$, then by the planarity of G^p , V must lie inside the triangle $\triangle ABC$, a contradiction to Lemma 4.4. Therefore, $\angle CWV \geq \angle CBA \geq \pi/2 - 2\pi/k$. If WV does not belong to an $S_{\geq 2}$ sector of W , WC and WV will form an S_1

sector because, by Proposition 4.8, there is no edge inside $F(C, A, B)$ connecting W to points other than C , and hence WC and WV are consecutive edges around W . Moreover, since $\angle CAB$ and $\angle CBA$ are acute (because CA and CB are in G^p) and $\angle BCA \leq 2\pi/k$, it is easy to verify that $|AB| < |BC|$ when $k \geq 10$. Now since WV is on the canonical path from A to B , $|WV| \leq |AB| < |BC|$, and hence WV is the shorter edge of the S_1 sector formed by WC (BC) and WV . This proves that WV is either an $S_{\geq 2}$ sector edge or an S_1 sector edge around W . The case where $W = A$ follows by symmetry.

Case 2: Now we can assume that W is a point on P_{AB} and $W \neq A, B$. By Proposition 4.8, the only possible edge in G^p inside $F(C, A, B)$ incident to W is CW . We distinguish two subcases based on whether or not there exists an edge $CW \in G^p$ interior to $F(C, A, B)$. Let WV' be the consecutive edge to WV on P_{AB} .

Subcase 2.1: CW is not an edge in G^p . By part 4 of Proposition 4.8, the angle at W is at least α_0 . Since the only edge in G^p inside $F(C, A, B)$ that can connect to W is CW , and CW does not exist, WV and WV' are two consecutive edges with $\angle VWV' \geq \alpha_0$, and hence must be part of an $S_{\geq 2}$ sector around W by the definition of an $S_{\geq 2}$ sector.

Subcase 2.2: CW is an edge in G^p . By part 5 of Proposition 4.8, the angle at W is at least $\pi - 2\pi/k$, and $\angle CWV, \angle CWV' \geq \pi/2 - 2\pi/k$. This shows that the consecutive edges WV, WC , and WV' are part of an $S_{\geq 2}$ sector, and hence WV is an $S_{\geq 2}$ sector edge around W . \square

Figure 3 illustrates an example of a canonical path between A and B , and the S_1 and $S_{\geq 2}$ sectors on this path. Let us denote this canonical path from A to B by $(A = M_0, M_1, M_2, M_3, M_4, M_5 = B)$. Assuming that the conditions in the statement of Lemma 4.10 are true (namely that CA and CB are in G^p , $\angle BCA \leq \alpha$, where $\alpha = 2\pi/k$ and $k \geq 10$), then by Lemma 4.10 it follows that $\langle AC, AM_1 \rangle$ and $\langle BC, BM_4 \rangle$ are S_1 sectors, and that $\langle M_1A, M_1M_2 \rangle$, $\langle M_2M_1, M_2C, M_2M_3 \rangle$, $\langle M_3M_2, M_3M_4 \rangle$, and $\langle M_4M_3, M_4B \rangle$ are all $S_{\geq 2}$ sectors. Therefore, each edge on this path is a sector edge for both its endpoints.

5 THE ALGORITHM

The algorithm constructs a bounded degree planar power spanner of degree bounded by k and a stretch factor bounded by $1 + (2 \sin(\pi/k))^p$, for parameter values $k \geq 10$ and $p \in [2, 5]$. We assume that the integer parameter $k \geq 10$ and the power constant $p \in [2, 5]$ are given.

5.1 Constructing the Generalized Gabriel Graph

We start by constructing the generalized Gabriel subgraph G^p of U . By Theorem 3.6, this step can be done by a 1-local distributed algorithm: each point decides which edges to keep based on its coordinates and the coordinates of its neighbors.

5.2 Computing the Approximation Angle

We would like to ensure that: (1) all edges of the S_1 and $S_{\geq 2}$ sectors around any point in G^p are selected (this ensures that, for any two points A and B in G^p , the edges on the canonical path between A and B in G^p will be selected); (2) for any edge $CB \in G^p$ incident on a point C , there is a selected edge CA incident on C with $\angle CB, CA \leq 2\pi/k$; and (3) for any point in G^p , at most k edges incident on the point are selected.

The above conditions, however, cannot be satisfied for values of $k \in \{10, 11, 12, 13, 14\}$. To see why this is the case, suppose for instance that $k = 10$, and suppose that there are five equally spaced $S_{\geq 2}$ sectors around a point C , each consisting of two consecutive edges making an angle of precisely $\alpha_0 = \arccos(1 - 2^{-3/5})$. Suppose further that each of the five angular sectors separating any two of the five $S_{\geq 2}$ contains edges that are shorter than the $S_{\geq 2}$ edges. Then one readily sees that picking the ten $S_{\geq 2}$ edges leaves us with no edges to pick. This makes it impossible to satisfy the above condition since for those shorter edges between the sectors that were not picked, the condition is not satisfied. Observe here that, even though for any edge CB lying in a region between two $S_{\geq 2}$ sectors there is a selected edge that makes an angle of at most $2\pi/10$ with CB , namely one of the boundary edges of the $S_{\geq 2}$ sectors, this fact does not help satisfying the desired condition because the boundary edge of the $S_{\geq 2}$ sector may be longer than CB .

To fix this problem, we first draw the following observation. The reason why we needed to satisfy the condition that for every edge CB there is a selected edge CA such that $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$, is that we wanted to guarantee that $|CA|^p + |AB|^p \leq (1 + 2^p \sin^p(\pi/k))|CB|^p$ (by Lemma 4.1), where $(1 + 2^p \sin^p(\pi/k))$ is the desired stretch factor. It turns out that even if the two edges CA and CB of G^p are such that $|CA| > |CB|$, the inequality $|CA|^p + |AB|^p \leq (1 + 2^p \sin^p(\pi/k))|CB|^p$ is still satisfied as long as $\angle BCA$ is smaller than a small ‘‘approximation’’ angle α_{apx} that depends on k . Certainly this approximation angle will be smaller than $2\pi/k$. We will show next how a lower bound on the angle α_{apx} can be computed. With this in mind, we can now make use of the selected boundary edges of the S_1 and $S_{\geq 2}$ sectors to cover an additional angular sector of size α_{apx} each.

Proposition 5.1: Let CA and CB be two edges in G^p . Let $\alpha_{apx} = \pi/k$, where $k \geq 10$ is an integer constant. If $\angle BCA \leq \alpha_{apx}$ then

$$|CA|^p + |AB|^p \leq (1 + (2 \sin \frac{\pi}{k})^p)|CB|^p.$$

Proof: Since $CA \in G^p$ we have $|CA|^p < |CB|^p + |AB|^p$, and hence

$$|CA|^p + |AB|^p < |CB|^p + 2|AB|^p. \quad (15)$$

Since CA and CB are edges in G^p , by part 4 of Proposition 3.4, the angles $\angle CAB$ and $\angle CBA$ are acute.

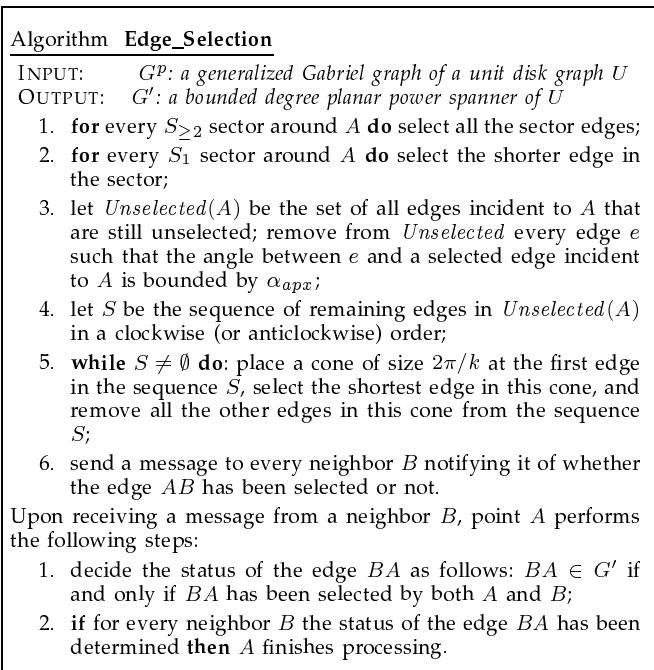


Fig. 4. Construction of the spanner G' .

This implies that

$$|AB|/|CB| \leq \tan(\alpha_{apx}) \leq \tan \frac{\pi}{k} \leq \sqrt{2} \sin \frac{\pi}{k}. \quad (16)$$

The last inequality is true because $\cos \frac{\pi}{k} \geq 1/\sqrt{2}$ for $k \geq 10$. Combining inequalities (15) and (16), we have

$$|CA|^p + |AB|^p \leq (1 + 2(\sqrt{2} \sin \frac{\pi}{k})^p) |CB|^p \quad (17)$$

$$\leq (1 + (2 \sin \frac{\pi}{k})^p) |CB|^p, \quad (18)$$

for $p \in [2, 5]$. \square

5.3 The Algorithm `Edge_Selection`

This algorithm works for $k \geq 10$. Every point A performs the following algorithm until the status of each of its incident edges has been defined, and at that point, A finishes processing. We assume that the integer parameter $k \geq 10$ and the power constant $p \in [2, 5]$ are given. The algorithm is given in Fig 4.

The presented algorithm is a 1-local distributed algorithm (assuming that the coordinates of the neighbors of the points have been computed during the computation of the generalized Gabriel graph). Moreover, no more than $O(m)$ messages are exchanged in the network since only $O(1)$ messages are exchanged along every edge. By Theorem 3.6, constructing G^p can be done by spending no more than $O(\Delta \lg \Delta) = O(n \lg n)$ local time at a point of degree Δ . After the construction of G^p , each point maintains a list of its incident edges sorted in a counterclockwise order around the point. With this in mind, it is easy to see that each of the above steps in the algorithm can be executed in $O(\Delta) = O(n)$ time by a point of degree Δ . It follows that the total local time spent by a point of degree Δ is $O(\Delta \lg \Delta) = O(n \lg n)$.

Let G' be the subgraph of G^p whose points are the point of G^p , and whose edges are the edges of G^p chosen by the above algorithm.

Theorem 5.2: The following are true.

1. The degree of G' is bounded by k .
2. If CB is an edge in G^p but is not an edge in G' , then there exists an edge $CA \in G'$ such that either (1) $\angle BCA \leq \alpha_{apx}$ or (2) $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$, and in both cases, all the edges on the canonical path between A and B are in G' .
3. The stretch factor of G' is bounded by $1 + (2 \sin \frac{\pi}{k})^p$.

Proof: Part 1: Consider an arbitrary point A and the edges incident to it in G' . By the `Edge_Selection` algorithm, each edge incident to A in G' must be selected by both of its endpoints, and in particular must be selected by A . To prove that the degree of A in G' is bounded by k , it suffices to show that A selects no more than k incident edges.

The proof is based on a mapping that maps each edge selected by A to an angular sector around A . We require that different edges be mapped to different disjoint sectors (i.e., non-overlapping sectors). We say that an edge AB covers an angle of measure δ , denoted by $\psi(AB) = \delta$, if it is mapped to an angular sector of measure δ . Because the sectors are disjoint, the sum of the angles covered by the edges selected by A is at most 2π .

Note that A selects three types of edges: the edges selected by A in step 1 of the algorithm, which are the $S_{\geq 2}$ sector edges; the edges selected in step 2 of the algorithm, which are the S_1 sector edges; and those edges selected in step 5 of the algorithm, which will be referred to as the *cone edges*.

If A selects only cone edges and no sector edges, then the number of cones around A is at most $\lceil 2\pi/(2\pi/k) \rceil = k$, and hence A selects at most k edges. Therefore we can assume that A selects at least one sector edge.

An $S_{\geq 2}$ sector consists of a maximal sequence of $\ell \geq 2$ consecutive $S_{\geq 2}$ sector edges such that the angle between the first and the last is at least $(\ell - 1)\alpha_0$, and the angle between any two consecutive edges in this sequence is at least $\pi/2 - 2\pi/k$. We divide the angular sector of size $(\ell - 1)\alpha_0$ into ℓ equal parts and map each $S_{\geq 2}$ sector edge to one part. Thus, each $S_{\geq 2}$ sector edge is mapped to an angle of least $(\ell - 1)\alpha_0/\ell \geq \arccos(1 - 2^{-3/5})/2 \geq 0.194\pi$. Moreover, since the $S_{\geq 2}$ sectors do not share edges, it is clear that the area mapped to by any $S_{\geq 2}$ sector edge does not overlap with the area mapped to by another distinct $S_{\geq 2}$ sector edge, regardless of whether the two edges are part of the same $S_{\geq 2}$ sector or not.

We map each S_1 sector edge to an area starting from the S_1 sector edge and spanning an angle of $2\pi/k$.

This completes our mapping for the sector edges.

Before we map the cone edges, we create an area around sectors called a "buffer area" as follows.

By Proposition 5.1, each sector edge covers an additional angle of $\alpha_{apx} = \pi/k$. Each such angle forms a buffer area. In addition, since an S_1 sector spans an angle

of $\pi/2 - 2\pi/k$, we also designate the difference between the sector angle and the area that the sector edge is mapped to as a buffer area. This buffer area measures at least $\pi/2 - 2\pi/k - 2\pi/k \geq \pi/2 - 4\pi/k \geq \pi/k$ for $k \geq 10$. Now each sector (S_1 or $S_{\geq 2}$) has two buffer areas, each measuring at least π/k , surrounding the sector on either side. Note that the buffer areas for different sectors may overlap.

Now to map the selected cone edges, consider a maximal sequence S_c of $\ell' \geq 1$ consecutive cone edges. The angle spanned by the cones corresponding to S_c (including the possible gaps between them) is at least $2(\ell' - 1)\pi/k$ because each cone in the sequence covers an angle of $2\pi/k$, except (possibly) for the last one. Since A selects at least one sector edge, there are $S_{\geq 2}$ or S_1 sectors at both ends of S_c , and hence there are buffer areas at both ends of S_c . Therefore, the angle covered by S_c plus these buffer areas measures at least $2(\ell' - 1)\pi/k + 2\pi/k = 2\ell'\pi/k$. Moreover, this area is disjoint from all the angular sectors that have been mapped to by the sector edges. Now we evenly divide this area into ℓ' cone edges, each cone edge covers an angle of at least $(2\ell'\pi/k)/\ell' = 2\pi/k$. Note that the areas that different maximal sequences of selected cone edges are mapped to are disjoint.

This completes the mapping of all the selected edges.

Now $\psi(AB)$, the measure of an angle covered by a selected edge AB is bounded by:

$$\psi(AB) \geq \begin{cases} 0.194\pi, & \text{when } AB \text{ is an } S_{\geq 2} \text{ sector edge;} \\ 2\pi/k, & \text{when } AB \text{ is an } S_1 \text{ sector edge;} \\ 2\pi/k, & \text{when } AB \text{ is a cone edge.} \end{cases}$$

Noting that all these areas are mutually disjoint, and that the area around point A measures 2π , it follows that the number of edges selected by A is at most $\lfloor 2\pi / \min(\psi) \rfloor \leq k$, when $k \geq 10$. This completes the proof of part 1.

Part 2: If $CB \in G^p$ is not in G' , then either C or B did not select CB . Without loss of generality, assume C did not select CB . Then CB must have been removed by C in either step 3 or step 5 of the **Edge_Selection** algorithm.

If CB is removed by C in step 3, then there exists an edge $CA \in G^p$ such that C selects CA and $\angle BCA \leq \alpha_{apx} = \pi/k$. Next we will show that CA is an $S_{\geq 2}$ sector edge around A and hence will also be selected by A . First note that the angle $\angle CBA < \pi/2$ by part 4 of Proposition 3.4. Since $\angle BCA \leq \pi/k$, $\angle CAB \geq \pi/2 - \pi/k \geq 2\pi/5 \geq \alpha_0$. By Lemma 4.4, there is no point inside the triangle $\triangle ABC$. Hence the angle between CA and a consecutive edge around A is at least α_0 . This implies that AC belongs to an $S_{\geq 2}$ sector of A and will be selected by A . Since CA is selected by both of its endpoints, $CA \in G'$.

On the other hand, if CB is removed by C in step 5, then there exists an edge $CA \in G^p$ such that CA is a cone edge selected by C , and hence $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$. Therefore $\angle CAB \geq (\pi - 2\pi/k)/2 \geq 2\pi/5$.




Fig. 5. A figure illustrating the lower bound on the stretching factor.

By the same argument as above, CA will also be selected by A , and hence is present in G' . This proves the first half of part 2.

By Lemma 4.10, every edge UV on the canonical path between A and B is a sector edge for both of its endpoints U and V , and hence is guaranteed to be present in G' . This completes the proof of part 2.

Part 3: If $CB \in U$ is not in G' , we will show that there exists a path from C to B in G' whose cost is at most $(1 + (2 \sin \frac{\pi}{k})^p)|CB|^p$. Since the power stretch factor of G^p is 1 by Proposition 3.4, we only need to consider the case where $CB \in G^p$ is not in G' . By part 2, in this case there exists an edge $CA \in G'$ such that either $\angle BCA \leq \alpha_{apx}$ or $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$.

If $\angle BCA \leq \alpha_{apx}$, by Proposition 5.1, we have $|CA|^p + |AB|^p \leq (1 + (2 \sin \frac{\pi}{k})^p)|CB|^p$. If $|CA| \leq |CB|$ and $\angle BCA \leq 2\pi/k$, by Lemma 4.1, we also have $|CA|^p + |AB|^p \leq (1 + (2 \sin \frac{\pi}{k})^p)|CB|^p$.

Furthermore, by Theorem 4.6, the cost of canonical path from A to B is at most $|AB|^p$, and by part 2 all the edges on the canonical path between A and B are in G' . This proves that if an edge $CB \in G^p$ is not in G' , then there is a path from C to B in G' whose cost is at most $(1 + (2 \sin \frac{\pi}{k})^p)|CB|^p$. The power stretch factor of G' is $1 + (2 \sin \frac{\pi}{k})^p$ as claimed. \square

5.4 Optimality

The stretch factor of $1 + (2 \sin \frac{\pi}{k})^p$ obtained above is almost optimal in the sense that constructing a power spanner of a unit disk graph with degree bound of k and a stretch factor smaller than $\rho = 1 + (2 \sin \frac{\pi}{k+1})^p$ is not possible. Consider the following example (see Figure 5). Point A_0 has $k+1$ neighbors, evenly distributed on a circle centered at A_0 . The angle between any two consecutive edges incident to A_0 is $2\pi/(k+1)$. In order to bound the degree of A_0 by k , at least one edge incident to A_0 has to be removed. Suppose A_0A_1 is removed, then the cost of sending a message from A_0 to A_1 is at least $|A_0A_2|^p + |A_2A_1|^p \geq (1 + (2 \sin \frac{\pi}{k+1})^p)|A_0A_1|^p$. This implies that a power spanner of a unit disk graph with degree bound of k has a stretch factor at least $\rho = 1 + (2 \sin \frac{\pi}{k+1})^p$.

6 ROBUSTNESS

In this section we will formally show that the algorithm presented in this paper is highly robust to topological changes. When the topology of the underlying network changes because of the introduction of a new point, or the deletion or motion of an existing point, the power spanner constructed by the algorithm **Edge_Selection** can be maintained efficiently without affecting the points beyond the vicinity of the change. Intuitively, this should be true due to the local distributed nature of the algorithm.

Since the motion of a point is equivalent to its deletion and reinsertion at a different location, we only discuss how to update the power spanner upon the insertion of a new point and the deletion of an existing point. We first discuss how such a topological change affects the generalized Gabriel graph.

Let U be a unit disk graph and let G^p be its generalized Gabriel graph. If a point A in G^p is deleted, then this can only cause the appearance of some edges in G^p , namely those edges of U killed by the point A . Note that both endpoints of any such edge must be neighbors of A . On the other hand, if a point A is inserted into G^p , then this insertion can only cause the disappearance of some edges from G^p , namely those that are killed by A . Again, both endpoints of any such edge must be neighbors of A . Therefore, we have the following theorems.

Theorem 6.1: Let U be a unit disk graph, G^p the generalized Gabriel graph of U , and A a point in U . If U' is the unit disk graph resulting from U by the removal of A , then the generalized Gabriel graph of U' can be computed from G^p by having every neighbor of A in U recompute its set of neighbors in the generalized Gabriel graph of U' .

Theorem 6.2: Let U be a unit disk graph and let G^p be its generalized Gabriel graph. If U' is the unit disk graph resulting from the insertion of a new point A in U , then the generalized Gabriel graph of U' can be computed from G^p by having every neighbor of A in U' , in addition to point A itself, compute its set of neighbors in the generalized Gabriel graph of U' .

Theorem 6.1 and Theorem 6.2 show that the change in the generalized Gabriel graph due to the insertion/deletion of a point is very local, only affecting the neighbors of the point.

Now we discuss how the insertion and deletion of a point in U affect the power spanner G' of U constructed by the algorithm **Edge_Selection**.

Let G^p be the generalized Gabriel graph of U . Suppose that a point $A \in U$ is deleted from U and let U' be the resulting unit disk graph. Let H^p be the generalized Gabriel graph of U' . From the above discussion, we know that H^p can be constructed from G^p by having each neighbor of A in U compute its neighbors in H^p . Let $N(A)$ be the set of neighbors of A in U . Note that in the algorithm **Edge_Selection**, every point selects its edges in the spanner G' from its set of incident edges in

G^p . Then adjacent points exchange messages along their common edge deciding whether to keep the edge or not. Only edges that are selected by both endpoints are kept.

For a point M in U' , its sets of incident edges in U and U' are different if and only if $M \in N(A)$. Let $N(N(A))$ be the set of neighbors in U' of those points in $N(A)$. For any point $M \in U' - (N(A) \cup N(N(A)))$, and for any neighbor M' of M in U' , the set of neighbors of M and M' are unchanged by the deletion of A from U , and hence their sets of selected edges in the algorithm **Edge_Selection** are unaffected by the deletion of A . Therefore, to update the spanner G' upon the deletion of A , we reapply the algorithm **Edge_selection** to H^p , but only to points in $N(A)$, to determine for each point its new set of selected edges for the new power spanner. Then each point in $N(A)$ exchanges messages with its neighbors in U' to agree on the common selected edges.

We conclude that the topological change in the power spanner upon the deletion of a point from U is local, and affects points that are only within two hops from the deleted point. The case is similar when a new point is inserted into U . We have the following theorem.

Theorem 6.3: Let U be a unit disk graph, G^p its generalized Gabriel graph, and G' its power spanner constructed by the algorithm **Edge_Selection**. If U' is the unit disk graph resulting from the insertion (resp. deletion) of a point A in (resp. from) U , and if H^p is the generalized Gabriel graph of U' , then the power spanner of U' can be computed from G' by applying the algorithm **Edge_Selection** to points in H^p that are within two hops from A in U .

The above theorem shows that the algorithm **Edge_Selection** is highly robust to topological change.

We also illustrate the robustness of the algorithm empirically. Figure 6 illustrates the changes incurred upon the insertion of a new point. The figure on the left shows a power spanner computed by the algorithm, and a newly inserted point in red color. The figure on the right shows the edges that will be added in red color, and those that will be deleted in blue dashed color. The changes are restricted to the two-hop vicinity of the inserted point.




Fig. 6. Illustrations for the robustness of the algorithm.

7 EXPERIMENTS AND DISCUSSIONS

We have shown that theoretically, our algorithm computes a power spanner whose stretch factor and maximum degree combination is better than what can be obtained using previous algorithms. In order to understand how our algorithm compares to the previous ones *empirically*, we performed experiments in which we compared power spanners obtained by three different algorithms: SYaoGG and OrdYaoGG from [5], and KPX from the present paper.

In our experiments, we placed n nodes on a 1500×1500 grid, uniformly at random, and we set the transmission range of each node to 600. We then ran two sets of simulations.

In the first set, we used $n = 30, 60, 90, \dots, 270, 300$, power exponents 2 and 4, and values of $k = 10, 12, 14, 16$. For each simulation, we fixed values of n , $power$, and k , and we ran all three algorithms on 100 different, randomly generated graphs of size n .

For power exponent 2, our experiments did not show significant differences between the performance of our algorithm and that of the other two. This is understandable because the full strength of the techniques in our algorithm (KPX)—such as the notion of generalized Gabriel graphs for example—comes out when the power exponent is greater than 2. This is confirmed by our experiments when we set the power exponent to 4. As can be seen from the top two graphs in Figure 7, the average stretch factor and the average maximum degree obtained by the algorithm KPX—over the 100 graphs—are both substantially smaller than those obtained by the other two algorithms. The same is true for the maximum stretch factor and the maximum maximum degree—over the 100 graphs—as shown in the bottom two graphs of Figure 7. The graphs in Figure 7 correspond to our simulations when $k = 10$, but similar behavior holds for larger k .

In the second set of simulations, we investigated the relationship between the input k , the *actual* maximum degree and the stretch factor obtained by the three algorithms. For that purpose, we used $n = 300$, power exponent 4, and values of $k = 10, 11, \dots, 16$. We ran all three algorithms on 500 different randomly generated graphs, and recorded the stretch factor and the maximum degree of the obtained spanners.

In Figure 8 we show the maximum stretch factor obtained by each algorithm—over the 500 input graphs—for each value of k . We also plot the theoretical upper bound $(1 + (2 \sin \frac{\pi}{k})^4)$ on the stretch factor of our algorithm. The maximum stretch factor returned by our algorithm is almost always substantially below the theoretical upper bound, which suggests that our algorithm will in practice do even better than what the theoretical upper bound suggests. Our algorithm also does better than OrdYaoGG and SYaoGG. While the algorithm OrdYaoGG obtains stretch factors that approach the ones obtained by our algorithm, it is important to remember

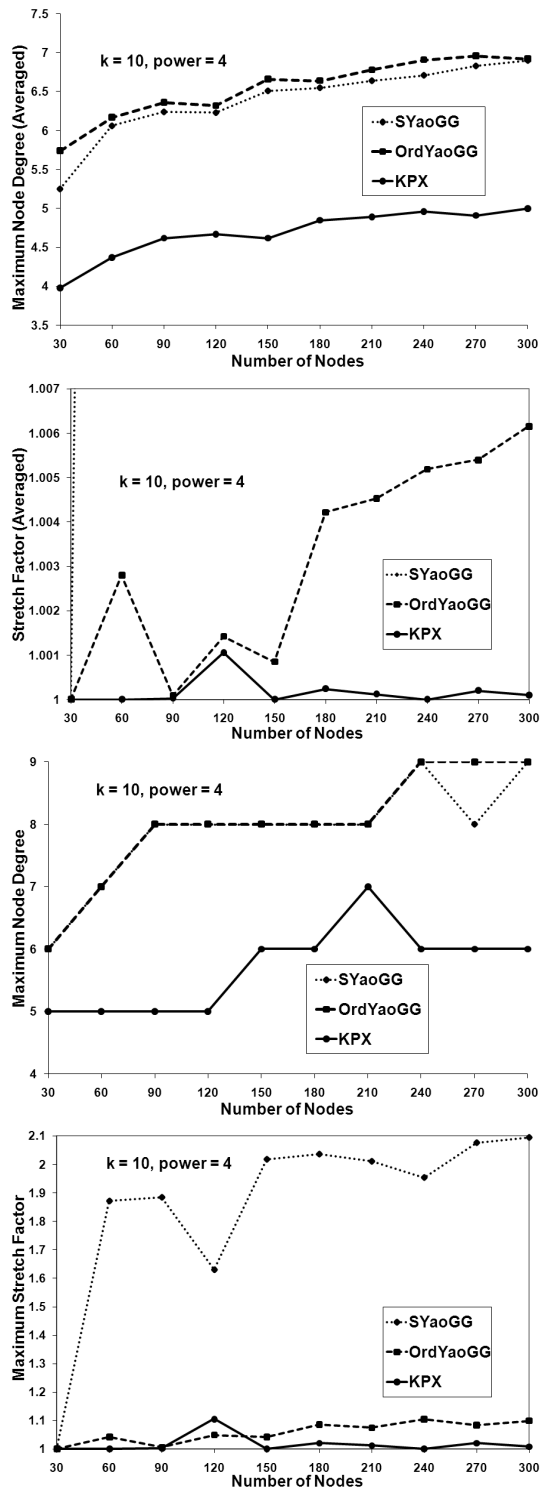


Fig. 7. The average maximum degree, the average stretch factor, the maximum maximum node degree, and the maximum stretch factor.

that algorithm OrdYaoGG is *not* local.

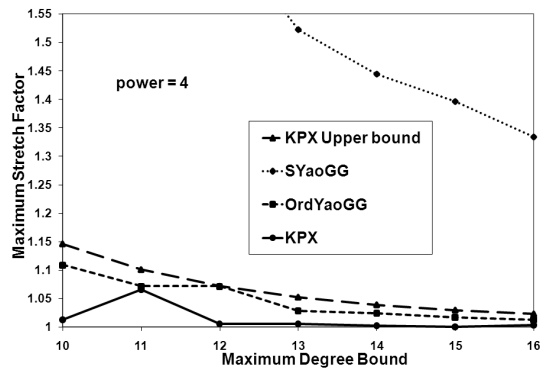


Fig. 8. Stretch factors as functions of input k .

When the three algorithms are run with a specific input value k , they are guaranteed to return a spanner of maximum degree k . In our experiments, and as can be seen in Figure 7, however, the *actual* maximum degree of the obtained spanner is always less than the input bound k . We use Table 2 to analyze the actual maximum degrees and stretch factors obtained by the three algorithms. In it we show the number of spanners obtained—in our simulations over 500 random input graphs—with given maximum degree and stretch factor range. Note that for more than half of the graphs, algorithm SYaoGG constructs spanners with degree 7 or more and stretch factor 1.1 or more. The algorithm OrdYaoGG improves on the stretch factor but not on the degree bound. Our local distributed algorithm amazingly computes spanners with maximum degree 5 and stretch factor 1 more than 85% of the time!

TABLE 2

The number of spanners with given maximum degree and stretch factor range returned by the SYaoGG (top), OrdYaoGG and KPX (bottom) algorithms.

$\rho \setminus \Delta$	4	5	6	7	8
$= 1.0000$	0	0	13	34	4
< 1.0001	0	0	0	0	0
< 1.001	0	0	0	1	1
< 1.01	0	0	3	11	0
< 1.1	0	0	18	71	9
≥ 1.1	0	0	74	228	33

$\rho \setminus \Delta$	4	5	6	7	8
$= 1.0000$	0	0	65	281	51
< 1.0001	0	0	0	0	0
< 1.001	0	0	1	5	0
< 1.01	0	0	7	26	4
< 1.1	0	0	12	36	11
≥ 1.1	0	0	0	1	0

$\rho \setminus \Delta$	4	5	6	7	8
$= 1.0000$	27	427	36	0	0
< 1.0001	0	0	0	0	0
< 1.001	0	0	0	0	0
< 1.01	1	5	1	0	0
< 1.1	0	2	1	0	0
≥ 1.1	0	0	0	0	0

It is clear for both sets of simulations that algorithm KPX is superior both theoretically and experimentally.

ACKNOWLEDGMENTS

The authors would like to thank the associate editor and the referees whose comments improved the paper.

Iyad A. Kanj obtained his Ph.D. from Texas A&M University in 2001. Since then he has been with DePaul University where he is currently an associate professor of Computer Science. His research interests are in the areas of parameterized complexity, graph theory and algorithms, combinatorial optimization, computational geometry, distributed algorithms and wireless computing.

Ljubomir Perković (Ph.D.'98) is an associate professor of Computer Science at DePaul University, Chicago. He obtained his Ph.D. in Algorithms, Combinatorics and Optimization at the School of Computer Science at Carnegie Mellon University. His research interests include distributed and wireless computing, graph theory and algorithms, and probabilistic analysis of algorithms.

Ge Xia received the Ph.D. degree in Computer Science from Texas A&M University in 2005. Since then, he has been with the Department of Computer Science at Lafayette College, where he is currently an assistant professor. His research interests include computational complexity and optimization, graph theory and algorithms, wireless networks, and computational biology.

REFERENCES

- [1] P. Bose and P. Morin, "Online routing in triangulations," *SIAM J. Comput.*, vol. 33, no. 4, pp. 937–951, 2004.
- [2] N. Linial, "Locality in distributed graph algorithms," *SIAM J. Comput.*, vol. 21, no. 1, pp. 193–201, 1992.
- [3] D. Peleg, *Distributed computing: A Locality-Sensitive Approach*. SIAM Monographs on Discrete Mathematics and Applications, 2000.
- [4] R. Wattenhofer, "Sensor networks: distributed algorithms reloaded - or revolutions?" in *SIROCCO*, 2006, pp. 24–28.
- [5] W.-Z. Song, X.-Y. L. Yu Wang, and O. Frieder, "Localized algorithms for energy efficient topology in wireless ad hoc networks," *Mobile Networks and Applications*, vol. 10, no. 6, pp. 911–923, 2005.
- [6] Y. Wang and X.-Y. Li, "Localized construction of bounded degree and planar spanner for wireless ad hoc networks," ser. Joint Workshop on Found. of Mobile Comput., 2003, pp. 59 – 68.
- [7] L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer, "Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks," in *PODC*, 2001, pp. 264–273.
- [8] X.-Y. Li, P.-J. Wan, and Y. Wang, "Power efficient and sparse spanner for wireless ad hoc networks," in *ICCC*, 2001, pp. 564 – 567.
- [9] R. Rajaraman, "Topology control and routing in ad hoc networks: a survey," *SIGACT News*, vol. 33, no. 2, pp. 60–73, 2002.
- [10] X.-Y. Li, P.-J. Wan, Y. Wang, and O. Frieder, "Sparse power efficient topology for wireless networks," in *HICSS*, 2002, pp. 296–.
- [11] R. Ramanathan and R. Hain, "Topology control of multihop wireless networks using transmit power adjustment," in *INFOCOM*, 2000, pp. 404–413.
- [12] Y. Wang, X.-Y. Li, and O. Frieder, "Distributed spanner with bounded degree for wireless networks," *International Journal of Foundations of Computer Science*, vol. 14, no. 2, pp. 183–200, 2003.
- [13] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, "Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks," in *INFOCOM*, 2001, pp. 1388–1397.
- [14] C. Scheideler, "Overlay networks for wireless ad hoc networks," 2005, <http://www14.in.tum.de/personen/scheideler/papers>.
- [15] K. R. Gabriel and R. R. Sokal, "A new statistical approach to geographic variation analysis," *Syst. Zoology*, vol. 18, no. 3, pp. 259–278, 1969.
- [16] A. C.-C. Yao, "On constructing minimum spanning trees in k-dimensional spaces and related problems," *SIAM J. on Comput.*, vol. 11, no. 4, pp. 721–736, 1982.
- [17] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia, "Routing with guaranteed delivery in ad hoc wireless networks," *Wireless Networks*, vol. 7, no. 6, pp. 609–616, 2001.