Switching to Directional Antennas with Constant Increase in Radius and Hop Distance

Prosenjit Bose[1], Paz Carmi[2], Mirela Damian[3],
Robin Flatland[4], Matthew J. Katz[5], and Anil Maheshwari[6]

[1] Carleton University, Ottawa, Canada
jit@scs.carleton.ca
[2] Ben-Gurion University, Beer-Sheva, Israel
carmip@cs.bgu.ac.il
[3] Villanova University, Villanova, P.A., USA
mirela.damian@villanova.edu
[4] Siena College, Loudonville, N.Y., USA
flatland@siena.edu
[5] Ben-Gurion University, Beer-Sheva, Israel
matya@cs.bgu.ac.il
[6] Carleton University, Ottawa, Canada
anil@scs.carleton.ca

Abstract. For any angle $\alpha < 2\pi$, we show that any connected communication graph that is induced by a set P of n transceivers using omni-directional antennas of radius 1, can be replaced by a strongly connected communication graph, in which each transceiver in P is equipped with a directional antenna of angle α and radius r_{dir}, for some constant $r_{\text{dir}} = r_{\text{dir}}(\alpha)$. Moreover, the new communication graph is a c-spanner of the original graph, for some constant $c = c(\alpha)$, with respect to number of hops.

Keywords: directional antennas, wireless networks, communication graph, hop spanner.

1 Introduction

Motivation and Problem Definition: Antennas used in wireless networks have traditionally been omni-directional. Such an antenna broadcasts in all directions, and its broadcast region can be represented geometrically by a disk centered at the transceiver. Recently, attention has been given to directional antennas which broadcast over a limited angle. The broadcast region of a directional antenna can be represented geometrically as a (closed) circular sector with an angular aperture α and radius r_{dir}. An antenna orientation is specified by a counter-clockwise angle θ measured from the positive x-axis to the sector's bisector. Directional antennas have the advantage of requiring less power compared to omni-directional antennas of the same radius, or by using the same power they can reach farther. In addition, narrower broadcast regions reduce interference and provide an added measure of security from eavesdroppers.

F. Dehne, J. Iacono, and J.-R. Sack (Eds.): WADS 2011, LNCS 6844, pp. 134–146, 2011.

The *direction assignment problem* is the task of finding orientations for a set of directional antennas such that the induced communication graph has certain desired properties. Let P be a set of n points in the plane representing n transceivers each equipped with an antenna. The induced communication graph has a vertex for each point and an edge directed from a to b if and only if b's point is contained in the broadcast region of a's antenna. Let $DG(r)$ be the induced communication graph when the points in P are equipped with omni-directional antennas of radius r. We will assume r is sufficiently large to ensure that $DG(r)$ is connected. It is easy to see that to achieve connectivity, r must be at least as long as the longest edge in a Euclidean minimum spanning tree of P. Consider now the same point set but with each point equipped instead with a directional antenna of angle α. Our goal is to determine a small radius $r_{dir} = r_{dir}(r, \alpha)$ for the directional antennas and to assign orientations to them, such that the resulting communication graph, G_{dir}, is (i) strongly connected, and (ii) for any two points $p, q \in P$, the number of hops (i.e., edges) in a minimum-hop path from p to q in G_{dir} is bounded by some constant $c = c(\alpha)$ times the number of hops in a minimum-hop path from p to q in $DG(r)$. In other words, condition (ii) requires that G_{dir} is a c-spanner of $DG(r)$ with respect to number of hops. Without loss of generality, we will assume going forward that $r = 1$.

New results: For $\alpha \leq \pi/3$, we show (in Section 2) that by fixing the radius of the antennas to $4\sqrt{2}(3.5k-6)$, where $k = \lceil 2\pi/\alpha \rceil$, one can assign orientations to the antennas such that G_{dir} is strongly connected and a $(\lceil 8 \log k \rceil - 1)$-spanner of $DG(1)$ with respect to the number of hops. In Section 3, we show that for the special case of $\alpha = \pi/3$ a radius of $36\sqrt{2}$ is sufficient to obtain a hop 10-spanner. This result immediately yields a hop 10-spanner using a radius of $4\sqrt{2}(3+k)$, for any $\alpha > \pi/3$. We also consider a path version of the problem in Section 4 where $DG(1)$ is assumed to be a path, and prove that a radius of $\lceil (3k+3)/2 \rceil$ is sufficient to obtain a hop $(2k+4)$-spanner, for $\alpha \leq \pi/3$. For $\alpha > \pi/3$, the result holds (trivially) using the same radius and hops as for $\alpha = \pi/3$. We note that although our directed networks have the advantages of reduced interference and added security, they do not achieve a power savings over the omni-directional network $DG(1)$ since the power savings of the smaller broadcast angle is offset by the larger radius $r_{\mathrm{dir}} > 1$. We leave it as an open problem to find directed networks with r_{dir} small enough to achieve a power savings.

Related Work: Caragiannis et al. [5] consider the problem of orienting directional antennas to form a strongly connected communication graph using minimal r, but they do not attempt to minimize the hops. For any $\alpha \geq 0$, they present an algorithm that constructs a communication graph containing a Hamiltonian tour and achieves a 3-approximation of the optimal radius. However, the number of hops between two nodes at unit distance or less can be linear. In [7], Damian and Flatland minimize both the radius and hops (as we do here) but for antennas with $\alpha \geq \pi/2$. Their approach depends fundamentally on finding orientations for small, proximate groups of antennas such that together the antennas completely cover an encompassing circular region and at the same time form a strongly connected sub-network amongst themselves. This, however, is not generally achievable for $\alpha < \pi/2$ — consider for example the case when the antennas all lie on a line — and thus their approach does not generalize to smaller angles. To our knowledge, our work here is the first to consider minimizing both the radius and hops for small angles $\alpha < \pi/2$.

In other related work, Nijnatten [10] considers a variant of the problem in which each antenna may have a different radius and the goal is to minimize the overall power consumption of the network. Ben-Moshe et al. [2] consider antennas with $\alpha = \pi/2$ but restrict the orientations to one of the four standard quadrant directions. Carmi et al. [6] show that for any set of points equipped with $\pi/3$ antennas, one can direct the antennas so that the resulting undirected communication graph is connected, assuming the antennas have a radius equal to the diameter of the point set. Ackerman et al. [1] later presented a simpler proof for the same result. Bhattacharya et al. [3] and Dobrev et al. [8] consider transceivers with multiple directional antennas; in the former they focus on minimizing the sum of the antenna angles for a fixed r, and in the latter they show that with k antennas per node, strong connectivity can be achieved using a radius that is at most $2\sin(\frac{\pi}{k+1})$ times the optimal, for any $\alpha \geq 0$. For a survey of recent results on directional antennas, see [9].

2 $\alpha \leq \pi/3$

Recall that $DG(1)$ is the communication graph induced by omni-directional antennas of radius 1 positioned at the points in P. Without loss of generality, we assume the point set is normalized so that the longest edge in a minimum spanning tree of P is at most 1, and thus $DG(1)$ is connected. Replacing each of the omni-directional antennas with a directional antenna of angle $\alpha \leq \pi/3$, we describe here how to determine orientations and a small radius $r_{\text{dir}} = r_{\text{dir}}(\alpha)$ for the antennas such that (i) the resulting communication graph, G_{dir}, is strongly connected, and (ii) for any two points $p, q \in P$, the number of hops in a minimum-hop path from p to q in G_{dir} is bounded by some constant $c = c(\alpha)$ times the number of hops in a minimum-hop path from p to q in $DG(1)$.

Lemma 1. *Let Q be a set of $m \geq 3$ points in the plane. Then, there exist three points $a, b, c \in Q$, such that $\angle abc \leq \pi/m$.*

Proof. Let $m' \leq m$ be the number of vertices in $\mathrm{CH}(Q)$, the convex hull of Q. Then, the sum of the angles at the vertices of $\mathrm{CH}(Q)$ is $(m'-2)\pi$, and there exists a vertex v whose corresponding angle is at most $(m'-2)\pi/m'$. Connect v to each of the points in Q. We obtain $(m-2)$ wedges, and the angle of at least one of them is bounded by $\frac{(m'-2)\pi/m'}{m-2} \leq \frac{\pi}{m}$, since $\frac{m'-2}{m'} \leq \frac{m-2}{m}$. □

For simplicity of exposition, assume for now α is such that $k = \lceil 2\pi/\alpha \rceil$ is a power of 2. We will later eliminate this restriction (in Theorem 1). We first describe our main building block.

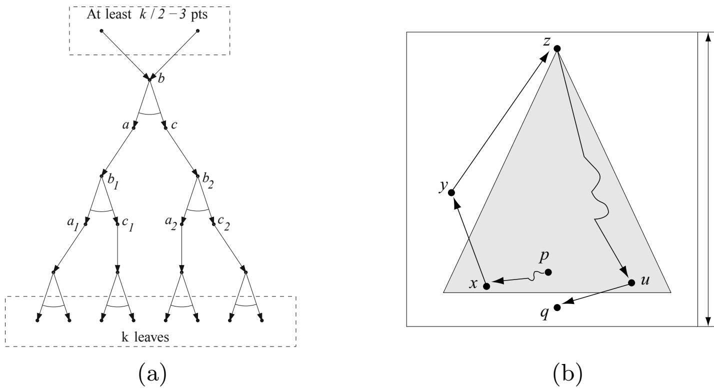

Fig. 1. (a) The communication structure of Theorem 1. (b) A path from p to q.

Tree Construction. We start with the assumption that there are sufficient points available for the construction described here. Later, we will determine a minimum number of points that are necessary to carry out this construction. Let Q be an arbitrary set of points. We construct a rooted binary tree whose nodes are points in Q and whose edges are directed towards the leaves (see Figure 1a). The tree has k leaves, and the angle formed at each internal node is at most α. We construct this tree as follows. Use Lemma 1 to pick three points $a, b, c \in Q$, such that $\angle abc \leq \alpha$. Make b the root of the tree and a and c the left and right children of b, respectively; remove points a, b, c from Q. Now, use Lemma 1 once again to obtain three new points $a_1, b_1, c_1 \in Q$, such that $\angle a_1b_1c_1 \leq \alpha$. Make b_1 the single child of a, and make a_1 and c_1 the left and right children of b_1, respectively; remove points a_1, b_1, c_1 from Q. In the next application of Lemma 1, we obtain three new points $a_2, b_2, c_2 \in Q$, such that $\angle a_2b_2c_2 \leq \alpha$, and make b_2 the single child of c, etc. We continue applying Lemma 1 until we obtain a balanced tree consisting of $3k - 3$ nodes, of which k are leaves, $k - 1$ have two children each, and $k - 2$ have one child each. See Figure 1a.

We now analyze the minimum number of points in Q necessary for this construction. Note that, in the last application of Lemma 1, we need $\pi/|Q| \leq \alpha$, to be able to select three points in Q forming an angle of at most α. So at least $\pi/\alpha \leq \pi/(2\pi/k) = k/2$ points are necessary in the last iteration. Therefore, $|Q|$ must be at least $(3k-3)+(k/2-3) = 3.5k-6$, since $(3k-3)$ points are selected (as established above), and at least $(k/2-3)$ points remain in Q after the last iteration. So the minimum size of Q is $\ell = 3.5k - 6$.

From this point on, whenever we use the term *tree*, we refer to the rooted tree constructed using this method. We refer to those points that are in the tree as *nodes*, to distinguish them from the points not in the tree.

Antenna Orientations. Once we have built our tree, we assign an orientation to each of the antennas at the points in Q according to the following rules:

(A1) At each node p in the tree, orient p's antenna to induce the directed edge(s) outgoing from p. This is always possible, because the angle spanned by the outgoing edge(s) at each node does not exceed α.
(A2) At each point $p \in Q$ that is not in the tree, orient p's antenna to point to the root of the tree (refer to Figure 1a).
(A3) At the k tree leaves, assign antenna orientations so that collectively they cover the whole plane (assuming infinite range). By the result by Bose et al. [4], this is always possible with k antennas of angle at least $2\pi/k \leq \alpha$.

The following lemma summarizes the properties of the resulting communication structure. All paths in this structure are *directed*.

Lemma 2. *Let Q be a set of points, each representing a transceiver equipped with a directional antenna of angle $\alpha \leq \pi/3$ and range diam(Q). Assume that $k = \lceil 2\pi/\alpha \rceil$ is a power of 2 and that $|Q| \geq 3.5k - 6$. Then, one can assign an orientation to each of the antennas at the points in Q, such that the resulting directed graph contains a rooted tree T with the following properties: (i) for any node q of T other than the root, there exists a path from q to a leaf of T consisting of at most $(2\log k - 2)$ hops, and (ii) for any point $q \in Q$, there exists a path from the root of T to q consisting of at most $(2 \log k)$ hops.*

Proof. Property (i) is immediate, since the number of hops from the root of T to a leaf equals the height of T, which is $(2\log k - 1)$; since q is assumed to be a non-root node, the worst case occurs when q is a child of the root, which yields $(2\log k - 2)$ hops. Rule (A3) for orienting the antennas at the leaf nodes in T guarantees the existence of a leaf node p that covers q. This hop, summed up with the $(2\log k - 1)$ from the root to p, yields the claimed bound of $(2\log k)$ hops from the root to q.

We now describe how to direct the antennas of the transceivers in the input set P. Recall that we are assuming that $r = 1$. We may assume that $|P| \geq \ell$, because otherwise, the distance between any two points in P is bounded by $\ell - 2$, and we can set $r_{\mathrm{dir}} = \ell - 2$ and form a directed cycle with at most $\ell - 2$ hops between any two points. Lay a grid $\mathcal{G}$ over P such that the length of each side of a cell is 2ℓ. Let $\mathcal{C}$ be a cell of $\mathcal{G}$, and define the *block of* $\mathcal{C}$ as the 3×3 portion of $\mathcal{G}$ that is centered at $\mathcal{C}$. Each of the 8 cells surrounding $\mathcal{C}$ is a *neighbor of* $\mathcal{C}$. A cell of $\mathcal{G}$ is considered *full* if it contains at least ℓ points of P. It is considered *non-full* if it contains at least one point of P, but less than ℓ points of P.

Proposition 1. *Let $\mathcal{C}$ be a cell of $\mathcal{G}$. Then, any path in $DG(1)$ that begins at a point in $\mathcal{C}$ and exits the block of $\mathcal{C}$, must pass through a full cell in $\mathcal{C}$'s block (not including $\mathcal{C}$ itself, which may or may not be full). In particular, if there are points of P outside $\mathcal{C}$'s block, then at least one of the 8 neighbors of $\mathcal{C}$ is full.*

Proof. This follows immediately from the fact that the maximum distance between any two adjacent points along this path is 1, and the edge length of a cell is 2ℓ. See Figure 2(a). □

Let F be the set of full cells of $\mathcal{G}$. We distinguish between two cases: $F = \emptyset$ and $F \neq \emptyset$. If $F = \emptyset$, then it is easy to see (using considerations similar to those used in the proof of Proposition 1) that P can be enclosed in a $2\ell \times 2\ell$ square. In other words, one can enclose P within a single grid cell, by shifting the grid if necessary. In this case, we simply set $r_{\text{dir}} = 2\sqrt{2}\ell$ and apply the construction used by Lemma 2 to the set P, with one simple modification:

(A3a) For each leaf node p of the tree, if p's antenna covers only tree nodes, then rotate it so that it covers the root of the tree.

This alteration guarantees that p can reach any point in P, via the tree root. Note that Lemma 2 still holds after reorienting some antennas, according to the rule (A3a) above. Let G_{dir} be the induced communication graph.

Lemma 3. *G_{dir} is a directed, strongly connected, $(4\log k)$-spanner of $DG(1)$, with respect to number of hops.*

Proof. Pick an arbitrary edge $pq \in DG(1)$. We show that G_{dir} contains a path from p to q with at most $(4\log k)$ hops. Refer to Figure 1b. If p is a tree node other than the root, then by Lemma 2(i), there is a path in G_{dir} from p to a tree leaf x, with at most $(2\log k - 2)$ hops. Rule (A3a) for reorienting the antennas at the leaves (if necessary), guarantees that x points either to the root z of the tree, or to a non-tree point $y \in P$, which in turn points to z. By Lemma 2(ii), there is a path from z to q in G_{dir} with at most $2\log k$ hops. Concatenating these paths together yields a path from p to q in G_{dir} with at most $(2\log k - 2) + 2 + (2\log k) = (4\log k)$ nodes. The remaining cases when p is the root of the tree, or a non-tree point in P, are subsumed by the case when p is a tree node discussed above. □

Assume now that $F \neq \emptyset$, i.e., there exists at least one full cell. Before discussing how to handle full cells, we introduce a few definitions. For any full cell $\mathcal{C}^+ \in F$, let $T(\mathcal{C}^+)$ denote the tree associated with $\mathcal{C}^+$. We say that the antenna at a leaf node of $T(\mathcal{C}^+)$ is *useful*, if one of the following three conditions holds: (i) it covers the root of either $T(\mathcal{C}^+)$, or the tree associated with another full cell in $\mathcal{C}^+$'s block; (ii) it covers a point in $\mathcal{C}^+$ that is not a node in $T(\mathcal{C}^+)$; (iii) it covers a point in a non-full cell in $\mathcal{C}^+$'s block, whose antenna covers the root of $T(\mathcal{C}^+)$. These situations are depicted in Figure 2(b)-(d). Define the *hop-distance between a point $p \in P$ and a full cell $\mathcal{C}^+ \in F$* as the number of hops in a minimum-hop path in $DG(1)$ between p and a point in $\mathcal{C}^+$.

We are now ready to define the orientations of the antennas at the points of P:

(A4) For each full cell $\mathcal{C}^+ \in F$, apply the construction of Lemma 2 to its corresponding set of points. Then, for each leaf x of $T(\mathcal{C}^+)$, if the antenna at x is not useful, reorient it so that it covers the root of $T(\mathcal{C}^+)$.

(A5) For each non-full cell $\mathcal{C}^-$ and for each point $p \in \mathcal{C}^-$, direct the antenna at p to the root of $T(\mathcal{C}^+)$, where $\mathcal{C}^+$ is the (hop-distance) closest full cell to p. Ties are broken arbitrarily.

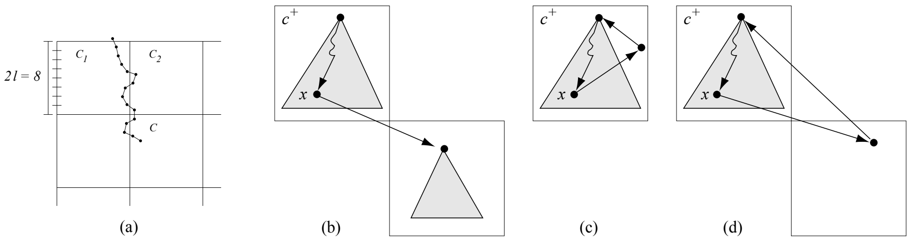

Fig. 2. (a) At least one of $\mathcal{C}$'s neighbors ($\mathcal{C}_1$) is full, (b-d) Situations in which leaf node x is useful

Note that Lemma 2 still holds, when applied to the point set restricted to a single full cell $\mathcal{C}^+$, and the full cell that determines the direction of the antenna at a point p that lies in a non-full cell $\mathcal{C}^-$, is a neighbor of $\mathcal{C}^-$ (this follows from Proposition 1). Lemma 4 below identifies two key properties of the leaf nodes in a full cell $\mathcal{C}^+$.

Lemma 4. *For any full cell $\mathcal{C}^+$ and any leaf node q of $T(\mathcal{C}^+)$, the following properties hold: (i) there is a path from q to the root of $T(\mathcal{C}^+)$ consisting of at most $(2\log k + 1)$ hops, and (ii) there is a path from q to the root of the tree associated with any full cell in $\mathcal{C}^+$'s block, via the root of $T(\mathcal{C}^+)$, with at most $(4\log k + 1)$ hops.*

Proof. Our reorientation rules force q to be useful, meaning that q either points to the root of a tree T in $\mathcal{C}^+$'s block (see Figure 2(b)), or can reach the root of $T(\mathcal{C}^+)$ in two hops (see Figures 2(c),(d)). Recall that the original orientation of the antennas at the k leaves of T guaranteed coverage of the entire plane, at infinite range (Rule A3). In particular, one of these antennas covers the root of $T(\mathcal{C}^+)$, proving itself useful (and therefore not subject to reorientation). These two hops (from q to the root of T, and from the leaf of T to the root of $T(\mathcal{C}^+)$), summed up with the $(2\log k - 1)$ hops from the root to a leaf of T, yield $(2\log k + 1)$ hops. So Property (i) is settled. To settle Property (ii), pick an arbitrary full cell $\mathcal{C}$ in $\mathcal{C}^+$'s block. We extend the path from the root of $T(\mathcal{C}^+)$ to the root of $T(\mathcal{C})$, in a similar way: one of the leaves of $T(\mathcal{C}^+)$ must cover the root of $T(\mathcal{C})$, and we can reach this leaf from the root of $T(\mathcal{C}^+)$ in $(2\log k - 1)$ hops. Summing these up, we obtain $(2\log k + 1) + (2\log k) = (4\log k + 1)$ hops. □

Finally, we set the radius r_{dir} of all antennas to $4\sqrt{2}\ell$, so that a point p can reach any other point in a neighboring cell (assuming the antenna at p is directed accordingly). Let G_{dir} be the resulting communication graph.

Lemma 5. *G_{dir} is a directed, strongly connected, $(8\log k - 1)$-spanner of $DG(1)$, with respect to number of hops.*

Proof. We prove that, for any edge pq of $DG(1)$, there exists a path from p to q in G_{dir} consisting of at most $(8\log k - 1)$ hops. Let $\mathcal{C}_p$ be the cell containing p, and $\mathcal{C}_q$ the cell containing q. Then, either $\mathcal{C}_p = \mathcal{C}_q$, or $\mathcal{C}_p$ and $\mathcal{C}_q$ are neighboring cells, because $|pq| \leq 1$. We discuss the cases where $\mathcal{C}_p$ is full or non-full.

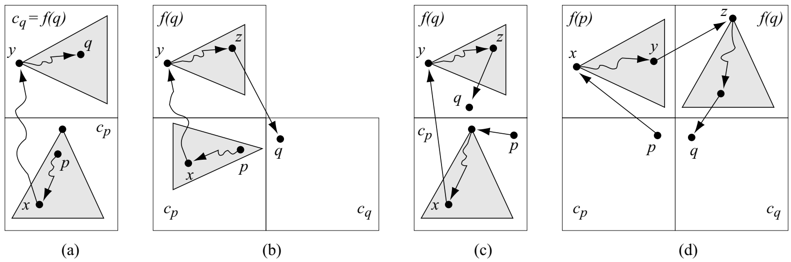

Fig. 3. Path from p to q in G_{dir}. (a) q is a node of $T(f(q))$. (b) q is not a node of $T(f(q))$. (c) p is not a node of $T(\mathcal{C}_p)$. (d) $\mathcal{C}_p$ is not full.

$\mathcal{C}_p$ is full. Associate with q a full cell, $f(q)$, as follows. If $\mathcal{C}_q$ is full, then $f(q) = \mathcal{C}_q$. Otherwise, $f(q)$ is the full cell that determines the direction of q's antenna. Since pq is an edge of $DG(1)$, the hop-distance between q and $\mathcal{C}_p$ is at most 1, so the hop-distance between q and $f(q)$ is at most 1. It follows that $f(q)$ is either $\mathcal{C}_p$, or a neighbor of $\mathcal{C}_p$. We now show how to reach q from p via the root of $T(f(q))$. We begin with the worst case scenario, in which p is a node in $T(\mathcal{C}_p)$ other than the root. In this case, from p we can reach an arbitrary leaf node x of $T(\mathcal{C}_p)$, in at most $(2\log k - 2)$ hops (by Lemma 2(i)). Next, we follow the path established in the proof of Lemma 4(ii) to reach the root y of $T(f(q))$, in at most $(4\log k + 1)$ hops. (Notice that $f(q)$ is in $\mathcal{C}_p$'s block, thus enabling us to use Lemma 4.) Finally, from y we follow the path in $T(f(q))$ that leads directly to q, if q is a node in $T(f(q))$ (see Figure 3a); otherwise, we follow the path that leads to the leaf z of $T(f(q))$ that covers q (see Figure 3b). Note that z always exists, since $\mathcal{C}_q$ and $f(q)$ are neighboring cells and q points to y. This latter path is longer, and has $(2\log k - 1) + 1$ hops. Summing up the number of hops along the entire path from p to q, we obtain a total of $(2\log k - 2) + (4\log k + 1) + (2\log k) = (8\log k - 1)$ hops. The cases in which p is the root of $T(\mathcal{C}_p)$ or a point not in $T(\mathcal{C}_p)$ are similar.

$\mathcal{C}_p$ is non-full. Let $f(p)$ be the full cell that determines the direction of p's antenna. Since pq is an edge of $DG(1)$, the difference between the hop-distance between p and $f(p)$ and the hop-distance between q and $f(q)$ is at most 1. Moreover, from the definition of $f(p)$ and $f(q)$, it follows that the corresponding paths (in $DG(1)$) from p to $f(p)$ and from q to $f(q)$ do not pass through a full cell (except at their final point). (If $f(q)$ is full, then the latter path is simply the degenerate path q.) This implies that $f(q)$ is either $f(p)$ or a neighbor of $f(p)$. We now show how to reach q from p via the root of $T(f(q))$. From p, we can reach the root of $T(f(p))$ in one hop. Next we follow the path in $T(f(p))$ from x to that leaf node y that covers the root z of $T(f(q))$; this path has $(2\log k - 1)$ hops. One hop takes us from y to z, then $(2\log k - 1)$ more hops take us from z to that leaf node of $T(f(q))$ that covers q. See Figure 3d. The total number of hops along the entire path from p to q is therefore at most $1 + (2\log k) + (2\log k) = (4\log k + 1)$. □

The following theorem summarizes the main result of this section.

Theorem 1. *Let P be a set of n points, each representing a transceiver equipped with a directional antenna of radius $\alpha \le \pi/3$, and assume that $DG(1)$ is connected. Then, one can assign a direction to each of the n antennas, such that by fixing their transmission range to $4\sqrt{2}(3.5k-6)$, where $k = \lceil 2\pi/\alpha \rceil$, the resulting communication graph is strongly connected. Moreover, it is a $(\lceil 8 \log k \rceil - 1)$-spanner of $DG(1)$, with respect to number of hops.*

Proof. The assumption that k is a power of two is eliminated here, so the tree with k leaves, used as the core communication structure within each full cell, is not necessarily balanced. More precisely, the bottom level of the tree may be incomplete. In this case, $(\lceil 2 \log k \rceil - 1)$ is an upper bound on the height of the tree. The result of this theorem follows immediately from Lemma 5. □

In concluding this section, we observe (without proof, due to space constraints) that our method for orienting antennas takes linear time in the number of points.

3 $\alpha = \pi/3$

We now show that we can obtain better constants for the special case $\alpha = \pi/3$. In our solution, we use the following result by Ackerman et al. [1]:

Proposition 2. *[1] Let Q be a set of points in the plane, equipped with antennas of angle $\pi/3$ and range diam(Q). There exist three points $x, y, z \in Q$, whose antennas can be oriented such that: (i) y covers x and z and both x and z cover y, and (ii) every point in Q is covered by at least one of $\{x, y, z\}$.*

Note that Proposition 2 does not claim that the three antennas can cover the entire plane; in fact, this would not be possible with three antennas only. The result is tied to a fixed point set Q: the three points are carefully selected from among the points on the convex hull of Q, so that collectively they cover Q.

For a point set Q of at least nine points, we construct a rooted tree structure T as follows. Select points x, y, z as in Proposition 2 and orient their antennas accordingly. Make y the root of the tree and make x and z its children. Select any six additional points in $Q \backslash \{x, y, z\}$, and for each, make it the child of a point in $\{x, y, z\}$ whose antenna covers it. Orient the antennas of these six points so that collectively they cover the entire plane when their radiuses are infinity. These six leaves serve the same function as the k leaves of the tree structure of Section 2. We note that x or z may also be a leaf of T; for example, x is a leaf if the six additional points are only covered by z's antenna and thus are all made children of z. But to remain consistent with Section 2, in what follows when we refer to the leaves of T, we are only refering to the six additional points, not x or z. To complete the construction, for each point in Q not in the tree, orient its antenna to point to y.

This tree structure is analogous to the tree of Section 2 (with $\alpha = \pi/3$), but it has better hop spanning properties because it has smaller height (at most 2). It is easy to verify properties analogous to those in Lemma 2 from Section 2 for the tree here. Specifically, for any non-root node, there exists a path from it to a leaf consisting of at most 3 hops (property (i)). For an example requiring 3 hops, consider the path from x to a leaf when all six leaves are children of z. For property (ii), 3 hops are sufficient to go from the root y to any point in Q, since the height of the tree is at most 2 and the leaves cover all points in Q.

Orienting the antennas of the input set P is now done the same as in Section 2, but using the tree structure above in each full cell. Since 9 points are sufficient to build the tree, we set $l = 9$. The grid cells are of size 18×18, and a cell is *full* if it encloses 9 or more points. By concatenating the same paths as in Lemma 4 but using the smaller trees described here, we immediately get the following two analogous properties: for any full cell $\mathcal{C}^+$ and any leaf node q of $T(\mathcal{C}^+)$, (i) there is a path from q to the root of $T(\mathcal{C}^+)$ consisting of at most 4 hops, and (ii) there is a path from q to the root of the tree associated with any full cell in $\mathcal{C}^+$'s block consisting of at most 7 hops. In fact, for the tree structure in this section, it is easy to show that these two properties also hold for the non-leaf nodes $\{x, y, z\}$ of $T(\mathcal{C}^+)$. This is true because both x and z's antennas cover the root y. So from a non-leaf node, it is at most 1 hop to the root y, at most 2 hops down $T(\mathcal{C}^+)$ to the leaf that covers the root of the desired tree in $\mathcal{C}^+$'s block, and then 1 hop to that neighboring tree's root.

By following the same arguments in Lemma 5 but using the smaller trees, we can show that the induced communication graph, G_{dir}, is a strongly connected 10-spanner of $DG(1)$, with respect to hops. As in Lemma 5, the worst case hop count in going from point p to q (where pq is an edge in $DG(1)$) occurs when $\mathcal{C}_p$ is full and p is a node in $T(\mathcal{C}_p)$ that is not the root. In this case, the path goes from p to the root of the neighboring tree $T(f(q))$ (worst case 7 hops, when p is a leaf), down $T(f(q))$ to the leaf that covers q (2 hops), and then to q (1 hop). We now set each antenna range to $36\sqrt{2}$, so that any two points that lie in neighboring cells can reach each other. Let G_{dir} be the resulting communication graph. The following theorem summarizes the main result of this section.

Theorem 2. *Let P be a set of n points, each representing a transceiver equipped with a directional antenna of angle $\alpha = \pi/3$, and assume that $DG(1)$ is connected. Then, one can assign a direction to each of the n antennas, and a transmission range of $36\sqrt{2}$, such that the resulting communication graph is a directed, strongly connected, 10-spanner of $DG(1)$, with respect to number of hops.*

We observe that this result also applies to angles $\alpha > \pi/3$ when using a radius $r_{dir} = 4\sqrt{2}(3 + \lceil 2\pi/\alpha \rceil)$, noting that in this case $\ell = 3 + \lceil 2\pi/\alpha \rceil$. The hop count remains the same.

4 Points along a Path

In this section we consider the special case in which the graph $DG(1)$ is a path. Let k be the smallest integer such that $2\pi/k \leq \alpha$, and let $m = \lceil |P|/k \rceil$. Partition $DG(1)$ into a sequence of m subpaths, the first $m-1$ of which have exactly $k+3$ points each. Let P_i denote the sequence of points along the i^{th} subpath. Thus each P_i, for $i < m$, contains $k+3$ points (P_m may have fewer points).

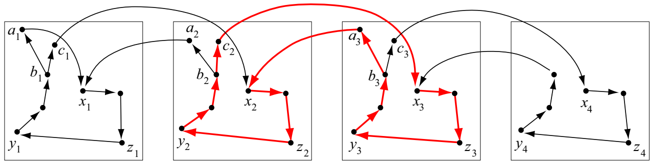

Fig. 4. The highlighted path shows that the hop bound of Theorem 3 is tight

Consider a sequence P_i with $k+3$ points. Define the *entry* point of P_i to be the midpoint x_i of P_i. Let y_i and z_i be the start and end point of the sequence P_i, respectively. From the point set $Q_i = P_i \setminus \{x_i, y_i, z_i\}$, select three points a_i, b_i and c_i, such that $\angle a_i b_i c_i \leq \alpha$. By Lemma 1, such points always exist, because $|Q_i| = k$. This enables us to orient b_i's antenna so that it covers both a_i and c_i. At each point p in the sequence $R_i = P_i \setminus \{a_i, b_i, c_i\}$, orient p's antenna so that it covers the next point in R_i, with two exceptions: the antenna at the predecessor of x_i is oriented to cover b_i, and the antenna at z_i covers y_i. Finally, orient a_i's antenna to cover the entry point x_{i-1} in the sequence P_{i-1}, if $i > 1$, or the entry point x_i of the same sequence, if $i = 1$. Orient c_i's antenna to cover the entry point x_{i+1} in the sequence P_{i+1}.

If P_m contains $k+3$ points as well, then we use the same approach for antenna orientations, with the difference that c_m's antenna points to x_m. If P_m contains fewer than $k+3$ points, then we orient the antenna at each point to cover the next point in the sequence, with two exceptions: the antenna at the predecessor of x_m covers x_{m-1}, and the antenna at z_m covers y_m. A specific example is depicted in Figure 4. Now note that a path from c_i to x_{i+1} in $DG(1)$ can go through at most $k+1$ points in P_i (this accounts for all points in P_i, with the exception of y_i and a_i), and at most $\lceil (k+3)/2 \rceil$ points in P_{i+1} (because x_{i+1} is the midpoint of the sequence). So the number of hops from c_i to x_{i+1} in $DG(1)$ is bounded above by $\lceil (3k+3)/2 \rceil$. The link from c_i to x_{i+1} (and symmetrically, from a_i to x_{i-1}) is one of the longest links that needs to be realized by an antenna, therefore we set $r_{\text{dir}} = \lceil (3k+3)/2 \rceil$. Lemma 6 below (whose proof is omitted due to space restrictions) summarizes properties of the resulting communication graph, G_{dir}.

Lemma 6. *For each $i < m$, G_{dir} contains directed paths from a_i to x_i, and from c_i to x_i, each with at most $k+3$ hops. In addition, G_{dir} contains a directed cycle that passes through all points in $P_i \cup P_{i+1} \setminus \{a_i, c_{i+1}\}$, with at most $2k+4$ hops.*

Theorem 3. *Let P be a set of n points, each representing a transceiver equipped with a directional antenna of radius $\alpha \leq \pi/3$, and assume that $DG(1)$ is a path. Then, one can assign a direction to each of the n antennas, such that by fixing their transmission range to $\lceil (3k+3)/2 \rceil$, where $k = \lceil 2\pi/\alpha \rceil$, the resulting communication graph G_{dir} is a directed, strongly connected, $(2k+4)$-spanner of $DG(1)$, with respect to the number of hops.*

Proof. We prove that, for each edge pq of $DG(1)$, there exists a path from p to q in G_{dir} consisting of at most $(2k+4)$ hops. Note that p and q are either in the same sequence P_i, or in adjacent sequences P_i and P_{i+1}, for some $i \geq 1$. If P_m contains fewer than $k+3$ points, assume $i < m-1$. By Lemma 6, there is a cycle passing through all points in $P_i \cup P_{i+1} \setminus \{a_i, c_{i+1}\}$, with at most $2k+4$ hops. If both p and q are on this cycle, then at most $2k+3$ hops separate q from p. Otherwise, p and q must belong to a same sequence, because neither a_i nor c_{i+1} is a start or end point of a sequence (by our construction), and therefore it cannot be adjacent in $DG(1)$ to a point in a different sequence. If $p = a_i$, then one can go from p to x_i in at most $k+3$ hops (by Lemma 6), then from x_i to q in at most $k+1$ additional hops, for a total of at most $2k+4$ hops. The situations $p = c_{i+1}$, $q = a_i$, and $q = c_{i+1}$ are similar.

The case when P_m contains fewer than $k+3$ points and $i = m-1$ is similar: G_{dir} contains a directed cycle with at most $2k+4$ hops, that passes through all points in $P_m \cup P_{m-1} \setminus \{a_{m-1}\}$. From this point on, the analysis is similar to the one above. □

We note that the transmission radius for the path case discussed in this section is smaller than the transmission radius for the general case (Section 2) by a factor of 26.4. And, although the hop count established by Theorem 3 increases linearly with $1/\alpha$, (as opposed to logarithmically in the general case,) the hop count for the path case is smaller than the hop count for the general case for any $k \leq 16$, which corresponds to $\alpha \geq \pi/8$.

Acknowledgments. Many thanks to the Fields Institute for financial support. Work by M. Katz was partially supported by grant 1045/10 from the Israel Science Foundation, and by the Israel Ministry of Industry, Trade and Labor (consortium CORNET). Work by P. Carmi was partially supported by the Lynn and William Frankel Center for Computer Sciences and grant 2240-2100.6/2009 from the German Israeli Foundation for scientific research and development (GIF).

References

1. Ackerman, E., Gelander, T., Pinchasi, R.: On connected wedge-graphs. (manuscript) (2010)
2. Ben-Moshe, B., Carmi, P., Chaitman, L., Katz, M.J., Morgenstern, G., Stein, Y.: Direction Assignment in Wireless Networks. In: CCCG 2010, pp. 39–42 (2010)
3. Bhattacharya, B., Hu, Y., Shi, Q., Kranakis, E., Krizanc, D.: Sensor network connectivity with multiple directional antennae of a given angular sum. In: IPDPS 2009, pp. 1–11 (2009)
4. Bose, P., Guibas, L., Lubiw, A., Overmars, M., Souvaine, D., Urrutia, J.: The floodlight problem. J. Assoc. Comput. Mach. 9, 399–404 (1993)
5. Caragiannis, I., Kaklamanis, C., Kranakis, E., Krizanc, D., Wiese, A.: Communication in wireless networks with directional antennae. In: SPAA, pp. 344–351 (2008)
6. Carmi, P., Katz, M.J., Lotker, Z., Rosén, A.: Connectivity guarantees for wireless networks with directional antennas. (manuscript) (2009)
7. Damian, M., Flatland, R.: Spanning Properties of Graphs Induced by Directional Antennas. In: Electronic Proc. of the 20th Fall Workshop on Computational Geometry, Stony Brook, NY (2010)

8. Dobrev, S., Kranakis, E., Krizanc, E., Opatrny, J., Stacho, L.: Strong Connectivity in Sensor Networks with Given Number of Directional Antennae of Bounded Angle. In: Wu, W., Daescu, O. (eds.) COCOA 2010, Part II. LNCS, vol. 6509, pp. 72–86. Springer, Heidelberg (2010)
9. Kranakis, E., Krizanc, D., Morales, O.: Maintaining Connectivity in Sensor Networks Using Directional Antennae. In: Nikoletseas, S., Rolim, J. (eds.) Theoretical Aspects of Distributed Computing in Sensor Networks, Part 2, pp. 59–84. Springer, Heidelberg (2011), ISBN 978-3-642-14848-4
10. van Nijnatten, F.: Range Assignment with Directional Antennas. Master's Thesis. Technische Universiteit Eindhoven (2008)
11. Wu, W., Du, H., Jia, X., Li, Y., Huang, S.C.-H.: Minimum connected dominating sets and maximal independent sets in unit disk graphs. Theoretical Computer Science 352, 1–7 (2006)