Global connectivity from local geometric constraints for sensor networks with various wireless footprints

Raissa D'Souza
Comp. Sci. and Eng.
Univ. of California
Davis, CA 95616
rmdsouza@ucdavis.edu

David Galvin
Dept. of Mathematics
Univ. of Pennsylvania
Philadelphia, PA 19104
dgalvin@math.upenn.edu

Cristopher Moore
Dept. of Computer Science
Univ. of New Mexico
Albuquerque, NM 87131
moore@cs.unm.edu

Dana Randall
College of Computing
Georgia Inst. of Tech.
Atlanta, GA 30332
randall@cc.gatech.edu

***Abstract*— Adaptive power topology control (*APTC*) is a local algorithm for constructing a one-parameter family of θ-graphs, where each node increases power until it has a neighbor in every θ sector around it. We show it is possible to use such a local geometric θ-constraint to determine whether full network connectivity is achievable, and consider tradeoffs between assumptions of the wireless footprint and constraints on the boundary nodes. In particular, we show that if the boundary nodes can communicate with neighboring boundary nodes and all interior nodes satisfy a $\theta_I < \pi$ constraint, we can guarantee connectivity for any arbitrary wireless footprint. If we relax the boundary assumption and instead impose a $\theta_B < 3\pi/2$ constraint on the boundary nodes, together with the $\theta_I < \pi$ constraint on interior nodes, we can guarantee full network connectivity using only a "weak-monotonicity" footprint assumption. The weak-monotonicity model, introduced herein, is much less restrictive than the disk model of coverage and captures aspects of the spatial correlations inherent in signal propagation and noise. Finally, assuming the idealized disk model of coverage, we show that when $\theta < \pi$, *APTC* constructs graphs that are sparse, and when $\theta < 2\pi/3$, the graphs support greedy geometric routing.**

Key words: ad hoc networks, topology control, adaptive power, connectivity, graph theory.

I. Introduction

We consider global properties of communications networks that can be guaranteed solely from local rules, particularly in the context of ad hoc networks which are typically both dynamic and temporary. A fundamental challenge is determining how to ensure global network connectivity using minimal overhead even when locations of nodes, and their linkages, can change over time. For ad hoc networks made of mobile nodes, the connectivity must evolve as the nodes move. Even for networks made of stationary nodes (such as some sensor networks), local connectivity can change over time due to the dynamic and noisy nature of wireless channels. We study a distributed and local construction (called Adaptive Power Topology Control, *APTC*) for building up communication edges between initially isolated nodes located on a two-dimensional plane, similar to the cone based topology control algorithm introduced by Wattenhofer et al. [1], and analyzed by Li et al. [2]. The only underlying information necessary for the construction is the local value of the angles formed between adjacent edges (i.e., links) E_v incident to each vertex (node) $v \in V$. These angles must all be less than a specified value θ, for all v. We call the graph describing the node positions and resulting edges at any time a θ-graph, denoted G_θ. Consider the graph G_R that is formed if we include all achievable linkages when each node broadcasts at maximal power. Assuming the broadcast region around each node is a uniform disk, Li et al. [2] provide an elegant geometric proof showing that if $\theta < 5\pi/6$, then the constructed graph G_θ preserves the connectivity of G_R, but is more sparse and therefore more power efficient. Though very useful if G_R is fully connected, meaning there exists at least one path connecting each pair of nodes, this does not give any method for testing the connectivity of G_R. Furthermore, it relies intrinsically on the uniform disk coverage model which, while a useful idealization for analysis, is not a realistic model for wireless footprints (see Sec. II, and Fig. 1(a)).

A. Our results

We show it is possible to use local geometric constraints to determine whether full network connectivity is achievable for any arbitrary wireless footprint, provided certain conditions are met. We define several tradeoffs between requirements of the boundary nodes and assumptions about the wireless footprint. Most previous algorithms impose constraints only on interior nodes and make strong assumptions about the wireless footprint. We show that with modest boundary requirements, the constraints on interior nodes and footprints can be greatly relaxed. This is an important consideration because when the network covers a large area, the boundary nodes will typically comprise only $O(\sqrt{n})$ of the n nodes. We might, for instance, carefully deploy a boundary region of sensors, then scatter sensors haphazardly in the interior. Further, in cases where deployment is inexpensive (consider a sensor network deployed by a robotic arm), internal nodes can be moved from dense regions to regions where the θ-constraint is not yet satisfied. If sensors are not moveable, existing sensor network protocols such as sleep cycling schemes could be easily employed by unnecessary nodes.

More precisely, we show that a modification of the *APTC* algorithm provably achieves global connectivity in a variety of scenarios. The more restrictive the boundary constraints, the weaker the assumptions required for the wireless footprint. (1) If the boundary nodes are known to be able to communicate with each other, then we can guarantee the entire network is connected provided all interior nodes satisfy a local $\theta_I < \pi$ requirement, for any arbitrary wireless footprint. (2) If we relax the communication requirement on the boundary nodes, but instead impose a local $\theta_B < 3\pi/2$-constraint on the boundary, and require all internal nodes to satisfy a $\theta_I < \pi$ constraint, then we can guarantee the entire network is connected for footprints that obey at least a "weak-monotonicity" constraint. Weak-monotonicity (introduced in Sec. III) is much less restrictive than the standard monotone footprint assumption and takes into account angular correlations between connections. Weak-monotonicity is also sufficient to ensure connectivity when all nodes satisfy the $\theta < \pi$ constraint on the sphere and the infinite plane, where there are no boundary nodes. (3) If the individual footprints are not uniform disks but the average over all footprints is approximately so, we show connectedness with high probability if $\theta_I < \pi$. Boundary nodes would need only to be connected to the interior, with no θ-constraint. These proofs all hold regardless of how a network is constructed, requiring only that local geometric constraints on θ are satisfied together with the appropriate boundary conditions.

This provides a general test for network connectivity that could easily be executed on a deployed system where nodes have access to local geometric information. Of course any individual node on its own would not be able to know if the network is fully connected, however the local information can be aggregated. If it is known that there are N nodes deployed, and all N send and receive messages that they satisfy their θ-constraint, we can locally learn of the global connectivity. If N is unknown, we can guarantee that our algorithm constructs the largest possible component that exists in G_R.

Finally, we prove additional properties of the *APTC* network. If the wireless footprints conform to the idealized disk coverage model, for $\theta < \pi$, the resulting graphs are sparse, and for $\theta < 2\pi/3$, the graphs provably support greedy geometric routing. If the footprints are not circular, but contain some smaller region which is circular, we show it may still be possible to support greedy geometric routing.

B. Related work

We study the *APTC* algorithm introduced by D'Souza et al. [3] which is similar to the construction by Wattenhofer et al. [1]. Although we deal with connectivity issues and not explicitly network performance, we note that in [3] the algorithm was shown to have extremely favorable performance characteristics, especially with regard to reducing power consumption and the timescale associated with discovery of the full network topology. Such optimizations could be particularly useful when coupled with routing algorithms relying on on-demand topology discovery, as studied by Perkins and Royer [4]. We show that when $\theta < 2\pi/3$ that greedy routing always works, assuming the disk model of coverage.

Most of the related previous work (e.g., [1] and [2]) relies on a priori knowledge of global network properties, such as the connectivity of the maximum power graph G_R. Poduri et al. [5] recently proved connectivity using only local geometric properties. However, this construction relies fundamentally on the uniform disk coverage model to achieve a supergraph of the Random Neighbor Graph. Wattenhofer and Zollinger [6] provide one of the first papers addressing local conditions for connectivity without assuming a unit disk model of coverage. In fact, their algorithm applies to three-dimensional systems, as well as nodes on a two-dimensional plane. The flexibility comes from requiring only an ordering on the quality of links, with no reference to geometry. Yet geometric constructions have some advantages. They can be simple to test and deploy, and enable geometric routing. Furthermore, many studies have already analyzed the performance characteristics of geometric ad hoc networks, showing them favorable.

II. Background and Terminology

A. Basic network operation assumptions

Ad hoc or sensor networks are composed of nodes equipped with wireless radios, allowing them to broadcast to, and receive messages from, other nodes over a shared wireless channel. Messages are exchanged directly between nodes within each other's broadcast range. However, exchanges with more distant devices can require relaying messages along a path of intermediary nodes. Thus data exchange relies fundamentally on devices cooperating in relaying one another's data.

The broadcast nature of a wireless network means a transmission interferes with all other simultaneous transmissions, with the greatest impact on transmissions sent by devices within close spatial range. We prefer devices to broadcast at low power to reduce interference, and moreover to conserve battery life (which can be the more important of the two criteria in the sensor network setting). The broadcast power, however, cannot be too low. It must be high enough to ensure neighboring devices can communicate and, at a larger scale, form a fully connected network (*i.e.*, a network where all devices have some, potentially multihop, path to all other devices). Understanding at what level to set each node's broadcast range has been the subject of numerous investigations.

B. Geometric graphs

The networks we consider can be modeled by geometric graphs. A geometric graph $G = (V, E)$ has vertices V (i.e., the wireless devices which are the nodes of the communication network) and a metric defining a distance between vertices. The edges of the graph E connect specific pairs of vertices. If a communication link exists between two nodes in the communication network, an edge between those two nodes exists in G. We consider the special case where the vertices inhabit a two-dimensional Euclidean plane, where a given vertex i has coordinates $(x_i, y_i) \in \mathbf{R}^2$, and we refer to the distance between nodes i and j as $d(i, j)$. Geometric graphs are convenient to describe the structure of many ad hoc networks, including some sensor networks, where nodes are constrained to lie in two-dimensions. In constrast, many other classes of networks exist in a space with no geometry, for instance the World Wide Web. For a recent comprehensive treatment of random geometric graphs see, e.g., [7].

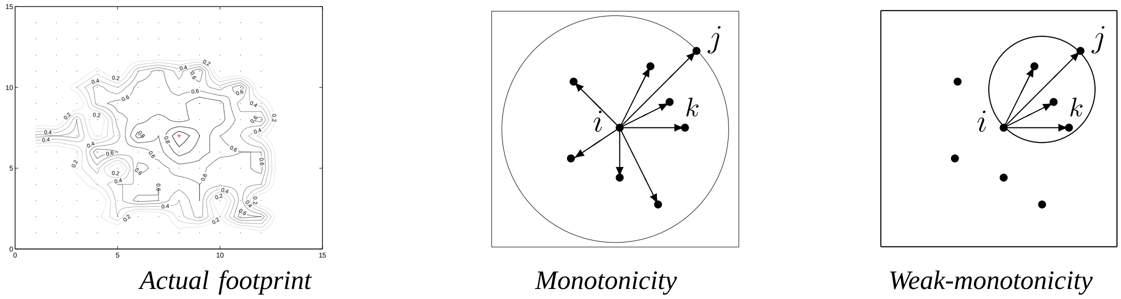

Fig. 1. (a) Example of an actual wireless footprint, reprinted from [8]. A central node broadcasts packets. The contours of probability for receiving the transmission are outlined. (b) Connectivity for node i assumed by the disk coverage model. (c) Connectivity for node i assumed by weak-monotonicity.

C. Wireless footprints

In principle, signals that are broadcast from a wireless device decay in an isotropic manner polynomially with distance from the source as $1/d^{\alpha}$, where $\alpha > 2$. Thus most models of connectivity conceptualize the broadcast region (or "footprint") as the *disk model of coverage* with a circular disk of radius r centered on each device i. For all points interior to the disk, all transmissions are considered successful, and the points connected to i. For all points with $d > r$, the signal is considered too small to distinguish from background noise so no transmission is ever received, and these points are considered not connected to i. The second part of Fig. 1 depicts this *monotonicity* assumption, where a successful connection between vertices i and j at the current level of transmission implies that i is also connected to all other closer vertices. Empirical studies of wireless sensor networks, however, show footprints are much less regular and can have large random deviations from a uniform disk. See, e.g., [8], and in particular Fig. 5 therein (reprinted here in Fig. 1). When a central node broadcasts, there is a complicated landscape of contours of probability of packet reception surrounding it with hills, voids and islands. As in [8], one can define a "good link" as one where the probability of packet reception is greater that Γ, where they take $\Gamma = 0.65$. The assumption is that with error correction techniques, etc., one can boost such a raw packet signal to adequate reception levels. Regardless, large deviations from a unit disk remain.

D. Distributed topology control algorithm for building G_θ

Consider V vertices distributed in $\mathbf{R}^2$. Details of the distribution are not pertinent for now. We begin from the isolated nodes and consider an algorithm for establishing the edges, E, and building up a graph $\vec{G}_\theta$ very similar to the one described in [3]. A fundamental requirement for the algorithm is access to directional information obtainable, for instance, from directional antennae, GPS, triangulation, or various other methods (see for instance [10]).

Each initially isolated node begins by transmitting at low power, incrementally ramping up until satisfying a geometric constraint on connectivity, as described below and illustrated in Fig. 2. As the node ramps up power incrementally, it broadcasts connection requests and processes acknowledgements of such request, thus establishing communication links with other nearby nodes. The node will first establish a link with the closest accessible node within its communication footprint, then with the next closest, etc. (Notice that we need not make any assumptions about isotropy or monotonic decay of the footprint; there could be nodes located at a closer spacial distance which do not get linked to since they are not in the accessible footprint). With each new connection made, the geometric information is assessed. In general, at each step, we consider the vectors drawn originating from a node and ending at its say m neighbors. These vectors divide the area around the central node into m disjoint sectors. If the angle of each sector is less than θ, the constraint is satisfied and the node sets its operating power at the current value. If any angle is greater than or equal to θ, the construction continues. If the node reaches its maximum allowed broadcast power level before satisfying the constraint, the node halts execution and lowers its broadcast power back down to the level where the last new connection to a neighbor was first made (or to zero if it has no neighbors in its broadcast range). The construction for $\theta = \pi$ was introduced by D'Souza et al. [3] and we refer to this algorithm as adaptive power topology control (*APTC*).

Each node i sets its operating range r_i independently of all other nodes, hence the resulting links may be unidirectional (i.e., the edges of $\vec{G}_\theta$ are directed). For both theoretical and implementation reasons we want all links to be bidirectional (resulting in an undirected graph G_θ). This can be achieved in many ways. We choose to do so at graph construction time.

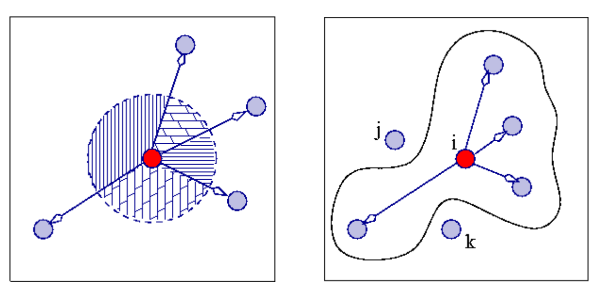

Fig. 2. (a) The vectors from a node to its m connected neighbors divide a unit circle around that node into m disjoint sectors. If the angle of each sector is less than or equal to θ, the geometric constraint is satisfied. (b) An example wireless footprint for node i. It does not connect to nodes j or k, even though they are closer in distance than other connected neighbors, yet i still satisfies its geometric constraint.

When node i broadcasts an acknowledgement to an in-link request from node k it must create a link to k, even if the length of that link $d_{ik} > r_i$. Node i would transmit with range r_i at all times, except when it needs to send a transmission directly to node k.[1] We refer to the underlying undirected graph as G_θ.

The algorithm used to generate G_θ can be integrated with standard wireless protocols such as the IEEE 802.11 wireless network MAC[11], and more specialized sensor network protocols such as sleep cycling schemes (see for instance [12]). In addition since the construction is local and distributed, it could be reiterated whenever a node notices its neighbors have changed.

III. Proofs of connectivity

We now show that we can ensure network connectivity using only local geometric constraints. The results hold for finite size systems, not just the asymptotic limit, however special consideration must be paid to nodes on the boundary. We assume *boundary nodes* on the convex hull of the pointset are identified in advance. We call these nodes B, and we say two nodes are *adjacent in* B if they are neighbors in the description of the convex hull, regardless of the distance between them. All other nodes are called *interior nodes*. We consider a family of boundary constraints on B. In general, the more restrictive the boundary constraints the less restrictions that need be imposed on the wireless footprint to guarantee connectivity.

A. Connected boundaries and arbitrary footprints

If the boundary nodes are identified as such and we know ahead of time that they are all connected, then the θ-constraint on all the internal nodes is sufficient to ensure global connectivity. This straightforward observation is formalized in the following theorem that will be used again in the following sections where we make less restrictive assumptions about the boundary nodes.

Theorem 1: If $G(V, E)$ satisfies the θ-constraint at every internal node with $\theta < \pi$ and all of the boundary nodes are known to be connected, then $G(V, E)$ is fully connected.

Proof: We need only show every internal node v has a path in $G(V, E)$ to some node on the boundary. Consider any line ℓ through the vertex v. Since v satisfies the θ_I-constraint, it must have some neighbor in each half-plane defined by ℓ. Consider one of these neighbors v_1, and for simplicity say v_1 lies to the "right" of ℓ. If v_1 is a boundary vertex we are done. Otherwise, let ℓ_{v_1} be the line parallel to ℓ through v_1. Vertex v_1 must have a neighbor v_2 to the right of ℓ_{v_1}. Continuing in this fashion, we must eventually find a vertex v_k on the boundary in the same connected component of v. ■

[1]For instance each node could keep an internal table of connected neighbors (already required by various routing protocols such as [9]), and in addition corresponding broadcast ranges.

B. Weak-monotonicity

We now relax the requirement that boundary nodes be connected to one another. In what follows we consider a variant on the *APTC* algorithm to produce (θ_I, θ_B) graphs where internal nodes satisfy the θ_I-constraint and boundary nodes satisfy the θ_B-constraint. We call the output of the algorithm a G_{θ_I,θ_B} graph. Notice that the θ_B-constraint allows the boundary nodes to stop increasing power once the constraint is satisfied. Previously under APTC boundary nodes were required to connect to all links reachable when using the maximum power. The geometrical interpretation of $\theta_B < 3\pi/2$ is that the links incident to any boundary node cannot be confined to a single quadrant around the node. Similarly, the $\theta_I < \pi$ constraint can be interperented as saying that links incident to an interior node cannot be confined to a single half-plane defined by a line through the node. To analyze this algorithm, we introduce *weak-monotonicity*, a less restrictive footprint model than the uniform disk model that captures spatial correlations inherent in signal propagation and noise. Under weak-monotonicity we will first show connectivity for G_{θ_I,θ_B} graphs, then generalize the result to sensors on a sphere, and then to the infinite plane.

Definition 1: Weak-monotonicity (see Fig. 3) implies that if $\vec{ij}$ is an edge and k is a node where $\angle jik = \alpha$ and $d(i,k) \leq \sin(\alpha) \cdot d(i,j)$, then $\vec{ik}$ is also an edge.

Weak-monotonicity is equivalent to saying that if $\vec{ij}$ is an edge, then i has a link to all other vertices in the circle of *diameter* $d(i,j)$ centered at the midpoint of the edge $\vec{ij}$. Note in contrast, the uniform disk model assumes i has a link to all other vertices in the circle of *radius* $d(i,j)$ centered at i. The first two parts of Fig. 3 depict the links that are inferred from an edge (i,j) under the monotone (disk model of coverage) and weak-monotone footprint assumtions. Notice that weak-monotonicity no longer assumes that signal propagation is monotone and isotropic, just that there are strong spatial correlations along directions of good and bad signal reception. Though this does not capture an arbitrary wireless footprint, it allows us to broaden the class of acceptable footprints far beyond the uniform disk model.

Connectivity for any G_{θ_I,θ_B} graph: Let G_{θ_I,θ_B} be the graph formed by *APTC* with the weak-monotonicity footprint model. We now show that if $\theta_I < \pi$, $\theta_B < 3\pi/2$, and these local θ constraints are satisfied at every internal and boundary node, then G_{θ_I,θ_B}, is connected. Since G_{θ_I,θ_B} is a subgraph of G_R (the graph formed when all pairs of nodes within distance R are connected), G_R is thus connected. We start by presenting a crucial lemma that says that two distinct components cannot have crossing edges, one from each component. This lemma uses the weak-monotonicity condition but does not require any knowledge of how the graph is connected.

Lemma 2: Let $\vec{G} = (V, \vec{E})$ be any directed graph satisfying the weak-monotonicity condition (i.e., for all $i, j, k \in V$ with $\angle ikj = \alpha$, if $\vec{ij} \in E$ and $d(i,k) \leq \sin(\alpha) \cdot d(i,j)$

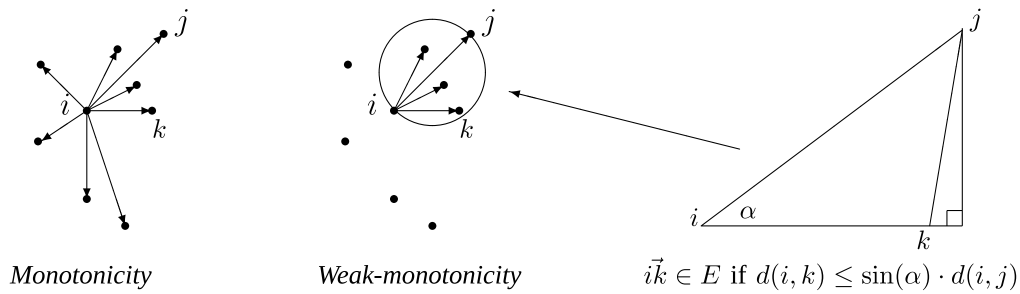

Fig. 3. Monotonicity and weak-monotonicity implications of the edge $\vec{ij}$

then $\vec{ik} \in E$). Let $G = (V, E)$ be the undirected version of $\vec{G}$ formed by making all edges bidirectional. Then any two crossing edges in G must belong to the same component.

Proof: Suppose G has two components C_1 and C_2 such that $\vec{ij} \in C_1$ and $\vec{kl} \in C_2$ cross. The quadrilateral (i, k, j, l) is depicted in Fig. 4. At least one angle of the quadrilateral must be greater than or equal to $\pi/2$, and we assume without loss of generality that it is $\angle ikj$. Then $d_{ik} \leq \sin(\alpha) \cdot d_{ij}$. Since $\vec{ij}$ is an edge in $\vec{G}$, by the assumption of weak-monotonicity, $\vec{ik}$ is also an edge in $\vec{G}$. This edge $\vec{ik}$ connects $\mathcal{C}_1$ and $\mathcal{C}_2$ in G, so they lie in the same component. ■

We now show that Lemma 2 is enough to ensure connectivity of G_{θ_I,θ_B} under the weak-monotonicity footprint model.

Theorem 3: Let $\theta_I < \pi$ and $\theta_B = 3\pi/2$. If G_{θ_I,θ_B} satisfies the θ_I-constraint at every internal node and the θ_B-constraint at every boundary node, then G_{θ_I,θ_B} is connected.

Proof: First, we observe that the proof of Theorem 1 shows that there is a path from each internal node to some vertex on the boundary. It remains only to show that all boundary vertices lie in the same connected component.

Suppose this is not true, and let x and y be the closest consecutive boundary vertices that lie in different components. Let ℓ' be the line through x and y, let ℓ_x be the line perpendicular to ℓ' through x and ℓ_y be the line perpendicular to ℓ' through y. (See Fig. 5.) Since the external angle around any point on the convex hull is at least π, $\theta_B > 3\pi/2$ implies that both x and y must have neighbors in the interior of the infinite rectangle delineated on three sides by ℓ_x, ℓ' and ℓ_y. We call the neighbor of x in this rectangle x_1 and, for the sake of terminology, we say that it lies to "the right" of ℓ_x. We call the neighbor of y in this rectangle y_1 and say it lies to "the left" of ℓ_y. As before, we continue building a path p_x from x that heads to the right at each step and a path p_y from y that heads to the left. These paths must end at boundary vertices x' and y'. If the paths intersect or cross, then by Lemma 2 they must lie in the same component and we have reached a contradiction. If they do not intersect or cross, then x' is a boundary vertex to the left of y' on the opposite side of the convex hull. If they are not nearest neighbors on the convex hull, find any two nearest neighbors on the hull lying between them that lie in different components in G_{θ_I,θ_B} and call these x' and y' instead. Notice that since we assumed that x and y were the closest boundary nodes lying in different components, we have $d(x', y') > d(x, y)$; therefore the edge (x', y') cannot be parallel to the edge (x, y) since x' and y' lie between ℓ_x and ℓ_y. Suppose without loss of generality that the lines through (x, y) and (x', y') intersect to the right of ℓ_y. As before, let $\ell_{x'}$ be the line perpendicular to the edge (x', y') through x', and similarly $\ell_{y'}$. There must be paths from x' and y' that cross or stay within the infinite rectangle delineated by $(x', y'), \ell_{x'}$ and $\ell_{y'}$. Since the path $p_{y'}$ originating at y' must reach a point on the convex hull to the left of y, it must intersect the path p_x. From Lemma 2 this proves that x and y lie in the same component in G_{θ_I,θ_B}. ■

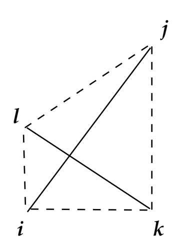

Fig. 4. The quadrilateral formed by crossing edges $\vec{ij}$ and $\vec{kl}$.

Connectivity on a sphere: These proofs can be generalized to a finite set of sensors on a sphere, where it is now possible to avoid the boundary constraints altogether. We assume that if two vertices are connected, then they take

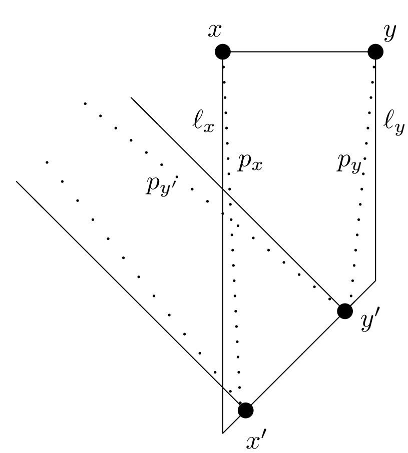

Fig. 5. Proof of Theorem 3

the shortest path around the sphere. In other words, even operating at full power, we can assume that there is no link that has length greater than half the circumference of a great circle. We show that if the spherical angles at each node satisfy the $\theta < \pi$ constraint, then G_θ is fully connected. The proof is similar in spirit to the finite planar setting. We first generalize Lemma 2 to the sphere, and then show that any two components must have crossing edges. Together this is sufficient to demonstrate that any spherical network satisfying the θ-constraint everywhere must be connected.

Lemma 4: Let $\vec{G} = (V, \vec{E})$ be a graph embedded on the unit sphere that satisfies the weak-monotonicity condition (i.e., for all $i, j, k \in V$ with $\angle ikj = \alpha$, if $\vec{ij} \in E$ and $d(i,k) \leq \sin(\alpha) \cdot d(i,j)$, then $\vec{ik} \in E$) and let G be the undirected version of $\vec{G}$. Then any two crossing edges must belong to the same component.

Proof: Let $\vec{ij}$ and $\vec{kl}$ be two crossing edges. The vertices $\{i, k, j, l\}$ form a quadrilateral. This quadrilateral divides the sphere into two pieces, and we refer to the piece containing the edges $\vec{ij}$ and $\vec{kl}$ as the *interior* of the quadrilateral. Since the length of $\vec{ij}$ and $\vec{kl}$ are less than half the circumference of any great circle, there must be an interior angle of the quadrilateral that exceeds $\pi/2$. We can use the proof of Lemma 2 to show that one of the edges of the quadrilateral must also be a link, assuming the weak-monotone model of coverage. ■

One additional lemma will be useful before stating and proving the main theorem for points on a sphere.

Lemma 5: Let P be a polygon on the sphere, where all the edges of P have length at most half the circumference of any great circle. If all the angles exceed π traveling around the polygon in one direction (viewed from one side of the polygon), then P lies in one half-sphere.

Proof: Let $e_1 = (p_1, p_2), e_2 = (p_2, p_3), e_3 = (p_3, p_4)$ be three consecutive edges on the polygon P, and let c be the great circle containing e_2. If all the angles exceed π traveling in one direction around P, then both e_1 and e_3 lie on the same half-sphere defined by c. The circles containing e_1 and e_3 intersect at antipodal points; let q be the one that lies in the same half-sphere (defined by c) as e_1 and e_3. We show P must lie inside the triangle defined by p_2, p_3 and q (where the interior of the triangle is the side bounded by angles that are less than π). It then follows that P lies on a half-sphere. If P is not contained in triangle (p_2, p_3, q), then there are at least two edges that starts outside this triangle and end at a vertex in or on the triangle. Following the polygon P around starting with e_2 in the direction of e_3, let e_i be the first edge that ends outside the triangle. If e_i crosses the circle containing e_3, then there must be an angle that exceeded π among the first i edges. If instead it crosses the circle containing e_1, all edges crosssing the boundary of triangle (p_2, p_3, q) must cross the circle containing e_1. Repeating the argument starting at e_2 and proceeding around the polygon in the other direction (first through e_2), we can similarly conclude that all edges crossing the boundary of triangle (p_2, p_3, q) must cross the circle containing e_3. This is a contradiction, so all of P must lie within the triangle and hence a half-sphere. ■

Theorem 6: If G_θ lies on the sphere and satisfies the θ-constraint everywhere, with $\theta < \pi$, then it is connected.

Proof: Suppose that there is more than one connected component in G_θ, and call two of these components C_1 and C_2. Notice that if every vertex $i \in V$ satisfies the θ-constraint, then every vertex has degree at least 3 and each component can be decomposed into a collection of minimal cells containing no other points from that component. If there are no crossing edges, then all of C_1 must lie within a single cell of C_2 (and, because all the points are lying on a sphere, this is equivalent to saying that all of C_2 lies in within a single cell of C_1). If we consider the vertices comprising these two cells, c_1 in C_1 and c_2 in C_2, it is not difficult to see that they cannot all be satisfying the θ-constraint if $\theta < \pi$. In particular, if the θ-constraint is satisfied by the vertices in c_1, then from lemma 5 c_1, and hence all of C_2, lies in one half-sphere. But then the constraint cannot be satisfied by its boundary cell c_2. ■

Connectivity in the infinite setting: For completeness we include the proof that in the infinite setting we can establish network connectivity over $\mathbf{R}^2$ using just the θ-constraint on the interior nodes, where $\theta = \theta_I < \pi$, under the weak-monotonicity assumption. This mathematical result inspired our definition of (θ_I, θ_B) graphs, but the proofs are somewhat technical. This section can be skipped for readers solely interested in finite realizations.

For $x \neq y \in \mathbf{R}^2$ we write $[x, y]$ for the (straight) line segment joining x and y. For $V \subset \mathbf{R}^2$ consider a graph $G = (V, E)$ on vertex set V. We refer to the set $\cup_{\{x,y\} \in V}[x, y]$ as the *realization* of G in $\mathbf{R}^2$ and say that G is a *θ-graph* if for each $x \in V$ every sector at x determined by the realization of G has angle less than θ.) Throughout this section we make the usual abuse of notation, identifying a subset of the vertices of a graph with the subgraph it induces.

Our conditions will be that $\theta < \pi$, that there is a uniform upper bound on the lengths of edges in E, that every finite disk in $\mathbf{R}^2$ contains only finitely many points of V, and that the neighborhood of each vertex obeys weak-monotonicity.

We show the following theorem.

Theorem 7: Let $V \subset \mathbf{R}^2$ satisfy the condition that its intersection with every disk of finite radius is finite. Let $G = (V, E)$ be a θ-graph on V with $\theta < \pi$. Then $G = (V, E)$ is connected and spans $\mathbf{R}^2$.

The proofs follow the general outline of the proofs from the finite setting, although they are much more sensitive. We defer these proofs for an Appendix in the final version.

C. Connectivity for footprints that are monotonic on average

Up until now we began with considering constraints on the boundary nodes, and from there determined requirements for the wireless footprints. Instead here we begin with constraints on the footprints. Though any individual footprint may have random deviations from a uniform disk (as shown in Fig. 1), here we assume that the average over all footprints is monotonic and isotropic. Given this, we can relax all constraints on boundary nodes and still show connectivity, with high probability, provided $\theta_I < \pi$. In such cases boundary nodes would just follow the APTC protocol and set their operating power accordingly. Recall the discussion in Sec. II of empirical wireless footprints and the definition of a "good link" (the probability of packet reception is greater than Γ). This leads us to the following definition.

Definition 2: For an arbitrary footprint, let $P(d)$ be the probability of packet reception at distance d from the source. We say the footprint is *isotropic and monotonic on average* if $P(d)$ has no dependence on angle (isotropic), and decays monotonically with d.

Note, for the disk graph assumption of strict monotonicity for a disk of radius r, $P(d) = 1$ if $d \leq r$, and $P(d) = 0$ if $d > r$. In addition, considering the definition of a "good link", if i and j are vertices in $G(V, E)$ and $\vec{ij}$ is an edge in $G(V, E)$, then $P(d_{ij}) > \Gamma$.

Theorem 8: If broadcast footprints decay monotonically on average, G_θ has one component with high probability.

Consider crossing edges of components C_1 and C_2 satisfying the θ-constraint. We let M be the number of crossing edges. Each pair of crossing edges (see Fig. 4) forms a quadrilateral where some $d(i,k) \leq d(i,j)$ while edge $\vec{ij} \in G(V, E)$. We know that $P[d(i,k)] \geq P[d(i,j)] > \Gamma$, and this holds for each set of crossing edges independently. The probably the components are not merged by a particular crossing edge is less than $(1 - \Gamma)$. The probability they are not merged by M independent crossing edges is less than $(1 - \Gamma)^M$. Setting $\Gamma = 0.65$ as in [8], if $M = 5$ the probability that a crossing edge will have merged C_1 and C_2 exceeds 99.5%.

IV. Behavior on Random Distributions

There are many advantages to assuming the idealized uniform disk coverage model. From an implementation perspective, it simplifies protocols and ensures reciprocity of signal reception. From an analytic perspective, it simplifies analysis, and allows us to prove additional features of the algorithm. We prove that under the disc model of coverage, the graph is sparse, the radii of the disks are tightly distributed, and moreover, when $\theta \leq 2\pi/3$, greedy routing works.

A. Sparseness of G_θ

Consider a Poisson distribution of points on a two-dimensional plane. Starting with an isolated node we consider the process of that node building up connectivity via the *APTC* algorithm. But, we no longer impose an upper cutoff to the maximum allowed broadcast power (hence we consider a supergraph of G_R considered up until now). We show the supergraph is sparse (hence so is G_θ).

Theorem 9: If the vertices are distributed uniformly at random, G_θ is sparse (i.e., $E = O(V)$).

Proof: Consider an individual node ramping up power according to the *APTC* algorithm. The node accumulates connected neighbors which divide the area around it into conic sectors. The node stabilizes its operating power when the angle of the largest conic section is less than θ, where $\theta = 2\pi A$ for some fixed $A \in (0, 1)$. For instance, if $A = 1/2$ then $\theta = \pi$, this is equivalent to stopping once the point is inside the convex hull.

If $Q(t)$ is the probability that this is true after t points, then the out-degree distribution $P(t)$ of the adaptive power model is the probability that it *first* happens after t points, i.e.

$$P(t) = Q(t) - Q(t-1) \ .$$

Now, recall that, for $t \geq 2$, choosing numbers $a_1, a_2, \ldots, a_t$ uniformly conditioned on $\sum_{i=1}^{t} a_i = 1$ is equivalent to choosing a uniform point $\vec{a}$ inside a $(t-1)$-dimensional equilateral simplex S of height 1, where the a_i are the lengths of the perpendiculars from $\vec{a}$ to the t, $(t-2)$-dimensional faces. Then the event that the largest angular gap is less than θ is equivalent to the event that $\vec{a}$ is within a distance A of every face (giving us an excluded area).

If $A = 1/2$ (i.e. $\theta = \pi$) the excluded areas are t simplices of height $1/2$. Each of these contains a fraction $1/2^{t-1}$ of the volume of S, so we have

$$Q(t) = 1 - \frac{t}{2^{t-1}}$$

for $t \geq 1$, and

$$P(t) = \frac{t-2}{2^{t-1}}$$

for $t \geq 2$. Amusingly, the average out-degree is then an integer:

$$\bar{t} = \sum_{t=2}^{\infty} P(t)\, t = 1 + \sum_{t=0}^{\infty} \frac{t(t-2)}{2^{t-1}} = 5$$

and the variance, σ^2, is 4.

Considering the stronger constraint $\theta = 2\pi/3$, the expected out degree is higher, yet the graph still sparse. In this case $A = 2/3$. Now each pair of excluded simplices has an intersection consisting of a simplex of height $1/3$ lying on the center of one edge. By inclusion-exclusion, we have

$$Q(t) = 1 - t\left(\frac{2}{3}\right)^{t-1} + \binom{t}{2}\left(\frac{1}{3}\right)^{t-1}$$

for $t \geq 1$, and so

$$P(t) = (t-3)\,\frac{2^{t-2} - t + 1}{3^{t-1}}$$

for $t \geq 2$. The average out-degree is then $71/8 = 8.875$ and the variance is $783/64 = 12.2344$. ∎

It is easy to show that the radius and link length distributions are tightly concentrated in the following sense: there is a constant C such that, in a network of n nodes uniformly distributed in the unit square, with high probability no radius is more than $C\sqrt{(\log n)/n}$.

B. Greedy routing works

One intuitive approach to routing on a wireless network is to pass the packet from its current location s to whichever neighbor is closest to the destination t. This greedy approach seems to have been first considered by Finn [13], who noted that it can get stuck at a local optimum where every neighbor

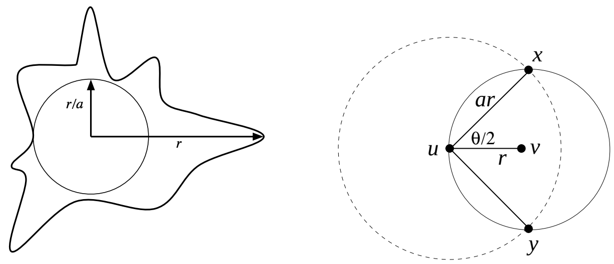

Fig. 6. Left, a footprint with eccentricity a; right, the proof of Theorem 10.

of s is farther from t than s is. Karp and Kung [9] called the space between s and t a "void", and proposed a protocol called Greedy Perimeter Stateless Routing (GPSR) that moves counterclockwise around the face of the graph containing the void until we reach the destination or greedy routing can resume. In order to ensure that this approach works, they first "planarize" the graph by reducing it to the Relative Neighborhood Graph (RNG)[14] or the Gabriel Graph (GG)[15].

In Ref. [16] the authors remark on the fact that greedy routing always works, assuming the uniform disk footprint model, and that the angular gap between neighbors is at most $2\pi/3$. Here we prove a more general result about when greedy routing works even if the footprint is not a uniform disk. Instead we require that the footprint contain some smaller region which is a uniform disk, as shown in Fig. 6(a). More precisely we require that each vertex contains a disk whose radius is some constant fraction of the distance to their farthest neighbor. Let us say that a network has *eccentricity* $a \geq 1$ where a is the smallest constant with the following property: for every u and v, if u and v are connected, then u is connected to every w such that $d(u,w) \leq d(u,v)/a$.

The next theorem states that as long as $a < 2$, there is some $\theta = \theta(a)$ such that if the angular gap between neighbors is at most θ, then greedy routing succeeds. For simplicity we ignore edge effects and assume that the network is spread throughout the plane.

Theorem 10: Suppose a network has eccentricity a where $a < 2$. Let $\theta = 2\cos^{-1}(a/2)$ and let $\epsilon > 0$, and suppose that every vertex u has at least one neighbor in every sector of angle $\theta - \epsilon$. Then for every pair of vertices u and v, u has at least neighbor w such that $d(w,v) < d(u,v)$.

Proof: Consider the right-hand part of Figure 6. By hypothesis, u has a neighbor w somewhere in the sector between x and y. If this neighbor is inside the circle centered on v, then $d(w,v) < d(u,v)$ and we are done; but if it is outside the dashed circle centered on u, then u and v are neighbors by the definition of eccentricity. By inspection we have $\cos(\theta/2) = a/2$. ■

When $a = 1$, we have the uniform disk model of coverage, and find that $\theta = 2\pi/3$, in agreement with the remark in [16]. Unfortunately, if $a > 2$ then there are arrangements of vertices in the plane such that greedy routing fails: for example, if the destination v is surrounded by a ring of vertices which are connected to each other but not to v.

V. Discussion

We have shown it is possible to guarantee global connectivity using only local geometric constraints. We explore tradeoffs between constraints on interior and boundary nodes and show that with modest boundary requirements, the constraints on interior nodes and footprints can be relaxed while connectivity still guaranteed. Many such tradeoffs exist in cooperative networked environments.

We introduce a "weak-monotonicity" model of wireless footprints which is much less restrictive than the uniform disk model, the latter being the most common model currently used for analysis. Typically strength of signal reception from a wireless source is not isotropic but is correlated with spatial directions. Weak-monotonicity captures this spatial correlation without needing to assume isotropy.

Our proofs are constrained to nodes on $\mathbf{R}^2$ or a sphere. Determining a corresponding geometric constraint for general three-dimensional systems would be extremely interesting.

References

[1] R. Wattenhofer, L. Li, P. Bahl, and Y. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of IEEE INFOCOM*, 2001.

[2] L. Li, J. Y. Halpern, P. Bahl, Yi-Min Wang, and R. Wattenhofer. Analysis of a cone-based topology control algorithm for wireless multi-hop networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, 2001.

[3] R. M. D'Souza, S. Ramanathan, and D. Temple Lang. Measuring performance of ad hoc networks using timescales for information flow. In *Proceedings of IEEE INFOCOM*, 2003.

[4] C. Perkins and E. Royer. Ad-hoc on-demand distance vector routing. In *MILCOM '97 panel on Ad Hoc Networks*, 1997.

[5] S. Poduri, S. Pattem, B. Krishnamachari, and G. Sukhatme. NET graphs: using local geometry for topology control in wireless sensor networks. *In Press*, 2004.

[6] R. Wattenhofer and A. Zollinger. XTC: A practical topology control algorithm for ad-hoc networks. In *Proceedings of the 18th International Parallel and Distributed Processing Symposium*, 2004.

[7] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, 2003.

[8] D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. TR UCLA/CSD-TR 02-0013, 2002.

[9] B. Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proc. 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 2000.

[10] L. Girod and D. Estrin. Robust range estimation using acoutic and multimodal sensing. In *Proceedings of the IEEE/RSJ Intl Conf on Intelligent Robots and Systems (IROS)*, 2001.

[11] IEEE Community LAN MAN Standards community Wireless LAN medium access control (MAC) and physical layer (PHY) specif., 1997.

[12] A. Cerpa and D. Estrin. Ascent: adaptive self-configuring sensor network topologies. *IEEE Transations Journal on Mobile Computing, Special Issue on Mission-Oriented Sensor Networks*, 3(3):272–285, 2004.

[13] G.G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical report, 1987.

[14] G. Toussaint. The relative neighborhood graph of a finite planer set. *Pattern Recognition*, 12(4):261–268, 1980.

[15] K. Gabriel and R.Sokal. A new statistical approach to geographic variation analysis. *Systematic Zoology*, 18:259–278, 1969.

[16] Q. Fang, J. Gao, and L. J. Guibas. Locating and bypassing routing holes in sensor networks. In *Proceedings of IEEE INFOCOM*, 2004.