A Cone-Based Distributed Topology-Control Algorithm for Wireless Multi-Hop Networks

Li (Erran) Li†
Department of Computer Science
Cornell University
lili@cs.cornell.edu

Joseph Y. Halpern†
Department of Computer Science
Cornell University
halpern@cs.cornell.edu

Paramvir Bahl
Microsoft Research
bahl@microsoft.com

Yi-Min Wang
Microsoft Research
ymwang@microsoft.com

Roger Wattenhofer*
ETH Zurich
wattenhofer@inf.ethz.ch

Abstract— **The topology of a wireless multi-hop network can be controlled by varying the transmission power at each node. In this paper, we give a detailed analysis of a cone-based distributed topology-control (CBTC) algorithm. This algorithm does not assume that nodes have GPS information available; rather it depends only on directional information. Roughly speaking, the basic idea of the algorithm is that a node u transmits with the minimum power $p_{u,\alpha}$ required to ensure that in every cone of degree α around u, there is some node that u can reach with power $p_{u,\alpha}$. We show that taking $\alpha = 5\pi/6$ is a necessary and sufficient condition to guarantee that network connectivity is preserved. More precisely, if there is a path from s to t when every node communicates at maximum power then, if $\alpha \leq 5\pi/6$, there is still a path in the smallest symmetric graph G_α containing all edges (u,v) such that u can communicate with v using power $p_{u,\alpha}$. On the other hand, if $\alpha > 5\pi/6$, connectivity is not necessarily preserved. We also propose a set of optimizations that further reduce power consumption and prove that they retain network connectivity. Dynamic reconfiguration in the presence of failures and mobility is also discussed. Simulation results are presented to demonstrate the effectiveness of the algorithm and the optimizations.**

I. Introduction

Multi-hop wireless networks, such as radio networks [11], adhoc networks [16], and sensor networks [4], [18], are networks where communication between two nodes may go through multiple consecutive wireless links. Unlike wired networks, which typically have a fixed network topology (except in case of failures), each node in a wireless network can potentially change the network topology by adjusting its transmission power to control its set of neighbors. The primary goal of topology control is to design power-efficient algorithms that maintain network connectivity and optimize performance metrics such as network lifetime and throughput. As pointed out by Chandrakasan et. al [2], network protocols that minimize energy consumption are key to the successful usage of wireless sensor networks. To simplify deployment and reconfiguration in the presence of failures and mobility, distributed topology-control algorithms that utilize only local information and allow asynchronous operations are particularly attractive.

The topology-control problem can be formalized as follows. We are given a set V of possibly mobile nodes located in the plane. Each node $u \in V$ is specified by its coordinates, $(x(u),y(u))$, at any given point in time. Each node u has a power function p where $p(d)$ gives the minimum power needed to establish a communication link to a node v at distance d away from u. Assume that the maximum transmission power p_{max} is the same for every node, and the maximum distance for any two nodes to communicate directly is R, i.e. $p(R) = p_{max}$. If every node transmits with power p_{max}, then we have an induced graph $G_R = (V,E)$ where $E = \{(u,v) |\, d(u,v) \leq R\}$ (where $d(u,v)$ is the Euclidean distance between u and v). Although this model is not always appropriate, Rodouplu and Meng [23] argue that it does capture various radio propagation environments.

It is undesirable to have nodes transmit with maximum power for two reasons. First, since the power required to transmit between nodes increases as the nth power of the distance between them, for some $n \geq 2$ [22], it may require less power for a node u to relay messages through a series of intermediate nodes to v than to transmit directly to v. Second, the greater the power with which a node transmits, the greater the likelihood of the transmission interfering with other transmissions.

Our goal in performing topology control is to find an undirected[1] subgraph G of G_R such that (1) G consists of all the nodes in G_R but has fewer edges, (2) if u and v are connected in G_R, they are still connected in G, and (3) a node u can transmit to all its neighbors in G using less power than is required to transmit to all its neighbors in G_R. Since minimizing power consumption is so important, it is desirable to find a graph G

This is a revised and extended version of "Analysis of a cone-based topology-control algorithm for wireless multi-hop networks", which appeared in *Proceedings of ACM Principles of Distributed Computing (PODC)*, 2001, and includes results from "Distributed topology control for power efficient operation in multihop wireless ad hoc networks", by R. Wattenhofer, L. Li, P. Bahl, and Y. M. Wang, which appeared in *Proceedings of IEEE INFOCOM*, 2001.

†The work of Halpern and Li was supported in part by NSF under grants grants IRI-96-25901, IIS-0090145, and NCR97-25251, and ONR under grants N00014-00-1-03-41, N00014-01-10-511, and N00014-01-1-0795.

*The work of Wattenhofer was supported in part by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

[1] Directed links complicate the design of routing and MAC protocols [19].

satisfying these three properties that minimizes the amount of power that a node needs to use to communicate with all its neighbors. Furthermore, for a topology control algorithm to be useful in practice, it must be possible for each node u in the network to construct its neighbor set $N(u) = \{v|(u,v) \in G\}$ in a distributed fashion. Finally, if G_R changes to G'_R due to node failures or mobility, it must be possible to reconstruct a connected G' without global coordination.

In this paper we consider a cone-based topology-control (CBTC) algorithm, and show that it satisfies all these desiderata. Most previous papers on topology control have utilized position information, which usually requires the availability of GPS at each node. There are a number of disadvantages with using GPS. In particular, the acquisition of GPS location information incurs a high delay, and GPS does not work in indoor environments or cities. By way of contrast, the cone-based algorithm requires only the availability of directional information. That is, it must be possible to estimate the direction from which another node is transmitting. Techniques for estimating direction without requiring position information are available, and discussed in the IEEE antenna and propagation community as the Angle-of-Arrival problem. The standard way of doing this is by using more than one directional antenna (see [12]). Specifically, the direction of incoming signals is determined from the difference in their arrival times at different elements of the antenna.[2]

The cone-based algorithm takes as a parameter an angle α. A node u then tries to find the minimum power $p_{u,\alpha}$ such that transmitting with $p_{u,\alpha}$ ensures that in every cone of degree α around u, there is some node that u can reach. We show that taking $\alpha = 5\pi/6$ is necessary and sufficient to preserve connectivity. That is, we show that if $\alpha \leq 5\pi/6$, then there is a path from u to v in G_R iff there is such a path in G_α (for all possible node locations) and that if $\alpha > 5\pi/6$, then there exists a graph G_R that is connected while G_α is not. Moreover, we propose several optimizations and show that they preserve connectivity. Finally, we discuss how the algorithm can be extended to deal with dynamic reconfiguration and asynchronous operations.

There were a number of papers on topology control prior to our work; as we said earlier, all assume that position information is available. Hu [9] describes an algorithm that does topology control using heuristics based on a Delauney triangulation of the graph. There seems to be no guarantee that the heuristics preserve connectivity. Ramanathan and Rosales-Hain [21] describe a centralized spanning tree algorithm for achieving connected and biconnected static networks, while minimizing the maximum transmission power. (They also describe distributed algorithms that are based on heuristics and are not guaranteed to preserve connectivity.) Rodoplu and Meng [23] propose a distributed position-based topology control algorithm that preserves connectivity; their algorithm is improved by Li and Halpern [13]. Other researchers working in the field of packet radio networks, wireless ad hoc networks, and sensor networks have also considered the issue of power efficiency and network lifetime, but have taken different approaches. For example, Hou and Li [8] analyze the effect of adjusting transmission power to reduce interference and hence achieve higher throughput as compared to schemes that use fixed transmission power [24]. Heinzelman et al. [7] describe an adaptive clustering-based routing protocol that maximizes network lifetime by randomly rotating the role of per-cluster local base stations (cluster-head) among nodes with higher energy reserves. Chen et al. [3] and Xu et al. [30] propose methods to conserve energy and increase network lifetime by turning off redundant nodes. Wu et al. [29] and Monks et al. [15] describe their power controlled MAC protocols to reduce energy consumptions and increase throughput. They do this through power control of unicast packets, but make no attempt at reducing the power consumption of broadcast packets.

After the initial publication of our results on CBTC [27], [14], there appeared a number of papers proposing different localized topology-control algorithms [28], [26], [10]. CBTC was the first algorithm that simultaneously achieved a variety of useful properties, such as symmetry, sparseness, and good routes; some of the recent topology also aim to simultaneously achieve a number of properties, most notably [26] and [10]. CBTC was also the first topology-control algorithm that did not require GPS information, but used only angle-of-arrival information. The only improvement towards this end that we are aware of is the XTC topology-control algorithm [28]. The XTC algorithm is somewhat similar in spirit to the SMECN algorithm [13], in that it removes an edge (u, v) if, according to some path-loss model, there is a two-hop path from u to v which nevertheless requires less energy than the direct path.

The rest of the paper is organized as follows. Section II presents the basic cone-based algorithm and shows that $\alpha = 5\pi/6$ is necessary and sufficient for connectivity. Section III describes several optimizations to the basic algorithm and proves their correctness. Section IV extends the basic algorithm so that it can handle the reconfiguration necessary to deal with failures and mobility. Section V describes network simulation results that show the effectiveness of the basic approach and the optimizations. Section VI summarizes this paper.

II. The Basic Cone-Based Topology Control Algorithm

We consider three communication primitives: broadcast, send, and receive, defined as follows:

- *bcast*(u, p, m) is invoked by node u to send message m with power p; it results in all nodes in the set $\{v|p(d(u, v)) \leq p\}$ receiving m.
- *send*(u, p, m, v) is invoked by node u to sent message m to v with power p. This primitive is used to send unicast messages, i.e. point-to-point messages.
- *recv*(u, m, v) is used by u to receive message m from v.

We assume that when v receives a message m from u, it knows the reception power p' of message m. This is, in general, less than the power p with which u sent the message, because of radio signal attenuation in space. Moreover, we assume that, given the transmission power p and the reception power p', u can estimate $p(d(u, v))$. This assumption is reasonable in practice.

For ease of presentation, we first assume a synchronous model; that is, we assume that communication proceeds in rounds, governed by a global clock, with each round taking one time unit. (We deal with asynchrony in Section IV.) In each round each

[2] Of course, if GPS information is available, a node can simply piggyback its location to its message and the required directional information can be calculated from that.

node u can examine the messages sent to it, compute, and send messages using the *bcast* and *send* communication primitives. The communication channel is reliable. We later relax this assumption, and show that the algorithm is correct even in an asynchronous setting.

The basic Cone-Based Topology-Control (*CBTC*) algorithm is easy to explain. The algorithm takes as a parameter an angle α. Each node u tries to find at least one neighbor in every cone of degree α centered at u. Node u starts running the algorithm by broadcasting a "Hello" message using low transmission power, and collecting Ack replies. It gradually increases the transmission power to discover more neighbors. It keeps a list of the nodes that it has discovered and the direction in which they are located. (As we said in the introduction, we assume that each node can estimate directional information.) It then checks whether each cone of degree α contains a node. This check is easily performed: the nodes are sorted according to their angles relative to some reference node (say, the first node from which u received a reply). It is immediate that there is a gap of more than α between the angles of two consecutive nodes iff there is a cone of degree α centered at u which contains no nodes. If there is such a gap, then u broadcasts with greater power. This continues until either u finds no α-gap or u broadcasts with maximum power.

Figure 1 gives the basic CBTC algorithm. In the algorithm, a "Hello" message is originally broadcasted using some minimal power p_0. In addition, the power used to broadcast the message is included in the message. The power is then increased at each step using some function *Increase*. As in [13] (where a similar function is used, in the context of a different algorithm), in this paper, we do not investigate how to choose the initial power p_0, nor do we investigate how to increase the power at each step. We simply assume some function *Increase* such that $Increase^k(p_0) = p_{max}$ for sufficiently large k. If transmission power can be set continuously in $[0, p_{max}]$, one can set $Increase(p) = 2p$ for fast convergence. If the initial choice of p_0 is less than the total power actually needed, then it is easy to see that this guarantees that u's estimate of the transmission power needed to reach a node v will be within a factor of 2 of the minimum transmission power actually needed to reach v. If transmission power can only be set to several discrete values, *Increase*(p) can be set to each value in increasing order. We adopt the latter approach in our simulation.

Upon receiving a "Hello" message from u, node v responds with an *Ack* message. Upon receiving the *Ack* from v, node u adds v to its set N_u of neighbors and adds v's direction $dir_u(v)$ (measured as an angle relative to some fixed angle) to its set D_u of directions. The test *gap*-$\alpha(D_u)$ tests if there is a gap greater than α in the angles in D_u. (We take *gap*-$\alpha(D_u) = 2\pi$ if $|D_u| = 1$.)

We use the following notation throughout the paper:
- $N_\alpha(u)$ is the final set of discovered neighbors computed by node u at the end of running CBTC(α).
- $p_{u,\alpha}$ is the corresponding final power.
- $N_\alpha = \{(u,v) \in V \times V : v \in N_\alpha(u)\}$.
- $G_\alpha = (V, E_\alpha)$, where V consists of all nodes in the network and E_α is the symmetric closure of N_α; that is, $(u,v) \in E_\alpha$ iff either $(u,v) \in N_\alpha$ or $(v,u) \in N_\alpha$.

CBTC(α)

$N_u \leftarrow \emptyset$; //the set of discovered neighbors of u
$D_u \leftarrow \emptyset$; //the directions from which the Acks have come
$p_u \leftarrow p_0$;

while $p_u < p_{max}$ and *gap*-$\alpha(D_u)$ **do**
    $p_u \leftarrow Increase(p_u)$;
    bcast$(u, p_u,$ ("Hello", p_u)) and gather Acks;
    $N_u \leftarrow N_u \cup \{v : v \text{ discovered}\}$;
    $D_u \leftarrow D_u \cup \{dir_u(v) : v \text{ discovered}\}$

Fig. 1. The basic cone-based algorithm running at each node u.

- *cone*(u', α, v') is the cone of degree α which is bisected by the line $\overline{u'v'}$, as in Figure 2.
- $N_c(u', \alpha, v')$ is the set of nodes inside *cone*(u', α, v').
- *circ*(u, r) is the circle centered at u with radius r.
- $rad^-_{u,\alpha}$ is the distance $d(u,v)$ of the neighbor v farthest from u in $N_\alpha(u)$; that is, $p(rad^-_{u,\alpha}) = p_{u,\alpha}$.
- $rad_{u,\alpha}$ is the distance $d(u,v)$ of the neighbor v farthest from u in E_α.

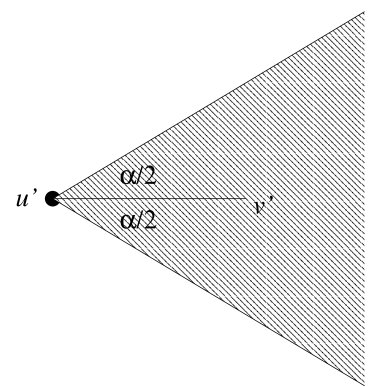

Fig. 2. *cone*(u', α, v')

Note that the N_α relation is not symmetric. As the following example shows, it is possible that $(v,u) \in N_\alpha$ but $(u,v) \notin N_\alpha$.

Example II.1: Suppose that $V = \{u_0, u_1, u_2, u_3, v\}$. (See Figure 3.) Further suppose that $d(u_0, v) = R$. Choose ϵ with $0 < \epsilon < \pi/12$ and place u_1, u_2, u_3 so that (1) $\angle vu_0u_1 = \angle vu_0u_2 = \pi/3 + \epsilon = \alpha/2$, (2) $\angle u_1vu_0 = \angle u_2vu_0 = \pi/3 - \epsilon$ (so that $\angle vu_1u_0 = \angle vu_2u_0 = \pi/3$), (3) $\angle vu_0u_3 = \pi$ (so that $\angle u_1u_0u_3 = \angle u_2u_0u_3 = 2\pi/3 - \epsilon$) and (4) $d(u_0, u_3) = R/2$. Note that, given ϵ and the positions of u_0 and v, the positions of u_1, u_2, and u_3 are determined. Since $\angle u_1u_0v > \angle u_0u_1v > \angle u_1vu_0$, it follows that $d(u_1, v) > d(u_0, v) = R > d(u_0, u_1)$; similarly $d(u_2, v) > R > d(u_0, u_2)$. (Here and elsewhere we use the fact that, in a triangle, larger sides are opposite larger angles.) Assume $2\pi/3 < \alpha \le 5\pi/6$. $N_\alpha(u_0) = \{u_1, u_2, u_3\}$, since there is no α-gap with this neighbor set. $N_\alpha(v) = \{u_0\}$, since v has to reach maximum power. Thus, $(v, u_0) \in N_\alpha$, but $(u_0, v) \notin N_\alpha$.
Example II.1 shows the need for taking the symmetric closure in computing G_α. Although $(u_0, v) \in G_R$, there would be no path from u_0 to v if we considered just the edges determined by N_α,

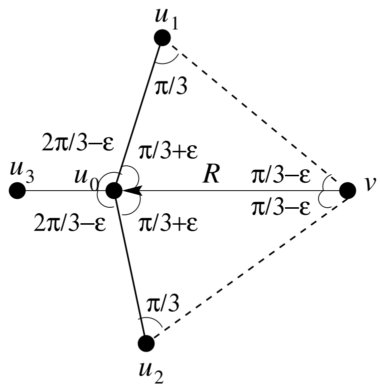

Fig. 3. N_α may not be symmetric.

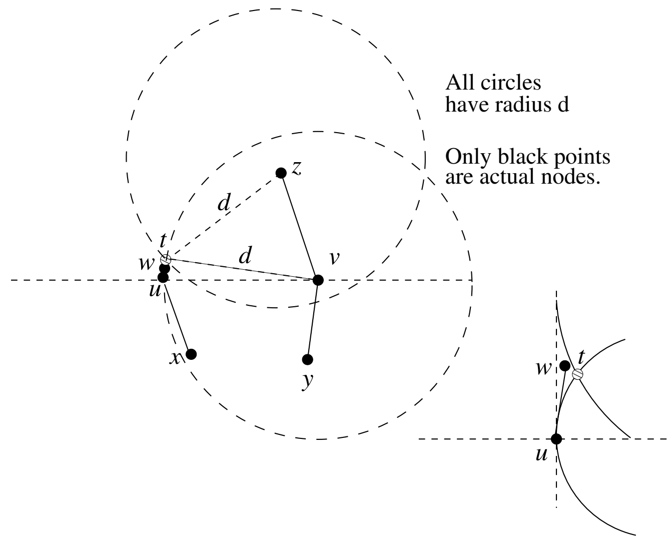

Fig. 4. Illustration for the proof of Lemma II.1.

without taking the symmetric closure. (The fact that $\alpha > 2\pi/3$ in this example is necessary. As we shall see in Section III-B, taking the symmetric closure is not necessary if $\alpha \leq 2\pi/3$.) As we have already observed, each node u knows the power required to reach all nodes v such that $(u, v) \in E_\alpha$: it is just the max of $p_{u,\alpha}$ and the power required by u to reach each of the nodes v from which it received a "Hello" message. (As we said earlier, if u receives a "Hello" from v, since it includes the power used to transmit it, u can determine the power required for u to reach v.)

We now prove the two main results of this paper: (1) if $\alpha \leq 5\pi/6$, then G_α preserves the connectivity of G_R and (2) if $\alpha > 5\pi/6$, then G_α may not preserve the connectivity of G_R. The following lemma will be used in the proof of (1).

Lemma II.1: If $\alpha \leq 5\pi/6$, and u and v are nodes in V such that $(u, v) \in E_R$ (that is, (u, v) is an edge in the graph G_R, so that $d(u, v) \leq R$), then either $(u, v) \in E_\alpha$ or there exist $u', v' \in V$ such that (a) $d(u', v') < d(u, v)$, (b) either $u' = u$ or $(u, u') \in E_\alpha$, and (c) either $v' = v$ or $(v, v') \in E_\alpha$.

Proof: If $(u, v) \in E_\alpha$, we are done. Otherwise, it must be the case that $d(u, v) > \max(rad^-_{u,\alpha}, rad^-_{v,\alpha})$. Thus, both u and v terminate CBTC(α) with no α-gap. It follows that $N_c(u, \alpha, v) \cap N_\alpha(u) \neq \emptyset$ and $N_c(v, \alpha, u) \cap N_\alpha(v) \neq \emptyset$. Choose $z \in N_c(v, \alpha, u) \cap N_\alpha(v)$ such that $\angle zvu$ is minimal. (See Figure 4.) Suppose without loss of generality that z is in the halfplane above $\overline{uv}$. If z is actually located in $cone(v, 2\pi/3, u)$, since $d(v, z) \leq rad^-_{v,\alpha} < d(u, v)$, it follows that $d(z, u) < d(u, v)$. For otherwise, the side zu would be at least as long as any other side in the triangle vzu, so that $\angle zvu$ would have to be at least as large as any other angle in the triangle. But since $\angle zvu \leq \pi/3$, this is impossible. Thus, taking $u' = u$ and $v' = z$, the lemma holds in this case. So we can assume without loss of generality that $z \notin N_c(v, 2\pi/3, u)$ (and, thus, that $N_c(v, 2\pi/3, u) \cap N_\alpha(v) = \emptyset$). Let y be the first node in $N_\alpha(v)$ that a ray that starts at vz would hit as it sweeps past vu going counterclockwise. By construction, y is in the half-plane below $\overline{uv}$ and $\angle zvy \leq \alpha$.

Similar considerations show that, without loss of generality, we can assume that $N_c(u, 2\pi/3, v) \cap N_\alpha(u) = \emptyset$, and that there exist two points $w, x \in N_\alpha(u)$ such that (a) w is in the halfplane above $\overline{uv}$, (b) x is in the halfplane below $\overline{uv}$, (c) at least one of w and x is inside $cone(u, \alpha, v)$, and (d) $\angle wux \leq \alpha$. See Figure 4.

If $d(w, v) < d(u, v)$, then the lemma holds with $u' = w$ and $v' = v$, so we can assume that $d(w, v) \geq d(u, v)$. Similarly, we can assume without loss of generality that $d(z, u) \geq d$. We now prove that $d(w, z)$ and $d(x, y)$ cannot both be greater than or equal to d. This will complete the proof since, for example, if $d(w, z) < d$, then we can take $u' = w$ and $v' = z$ in the lemma.

Suppose, by way of contradiction, that $d(w, z) \geq d$ and $d(x, y) \geq d$. Let t be the intersection point of $circ(z, d)$ and $circ(v, d)$ that is closest to u. Recall that at least one of w and x is inside $cone(u, \alpha, v)$. As we show in Appendix A, since node w must be outside (or on) both circles $circ(z, d)$ and $circ(v, d)$, we have $\angle wuv \geq \angle tuv$ (see the closeup on the far right side of Figure 4).

Since $d(t, z) = d(t, v) = d(u, v) = d$, and $d(z, v) < d$, it follows that $\angle zvt > \pi/3$. Thus,

$$\angle tvu = \angle zvu - \angle zvt < \angle zvu - \pi/3 \text{ and}$$
$$\angle tvu = \pi - 2 \times \angle tuv,$$

and so

$$\angle zvu - \pi/3 > \pi - 2 \times \angle tuv \text{ and,}$$
$$\angle tuv > 2\pi/3 - \angle zvu/2.$$

Since $\angle wuv \geq \angle tuv$, we have that

$$\angle wuv > 2\pi/3 - \angle zvu/2. \tag{1}$$

By definition of z, $\angle zvu \leq \alpha/2 \leq 5\pi/12$, so $\angle wuv > 2\pi/3 - 5\pi/24 = 11\pi/24 > \alpha/2$. Thus, it must be the case that $w \notin N_c(u, \alpha, v)$, so $x \in N_c(u, \alpha, v)$.

Arguments identical to those used to derive (1) (replacing the role of w and z by y and x, respectively) can be used to show that

$$\angle yvu > 2\pi/3 - \angle xuv/2. \tag{2}$$

From (1) and (2), we have

$$\begin{aligned} & \angle wuv + \angle xuv \\ > \ & (2\pi/3 - \angle zvu/2) + (4\pi/3 - 2 \times \angle yvu) \\ = \ & 2\pi - \angle zvu/2 - 2 \times \angle yvu. \end{aligned}$$

Since $\angle wuv + \angle xuv \leq \alpha \leq 5\pi/6$, we have that $5\pi/6 > 2\pi - \angle zvu/2 - 2 \times \angle yvu$. Thus,

$$\angle zvu/2 + 2\times\angle yvu = ((\angle zvu + \angle yvu) + 3\times\angle yvu)/2 > 7\pi/6.$$

Since $\angle zvu + \angle yvu \leq \alpha \leq 5\pi/6$, it easily follows that $\angle yvu > \pi/2$. As we showed earlier, $\angle zvu \geq \angle zvt > \pi/3$. Therefore, $\angle zvu + \angle yvu > 5\pi/6$. This is a contradiction. ∎

Theorem II.2: If $\alpha \leq 5\pi/6$, then G_α preserves the connectivity of G_R; u and v are connected in G_α iff they are connected in G_R.

Proof: Since G_α is a subgraph of G_R, it is clear that if u and v are connected in G_α, they must be connected in G_R.

We now prove the converse. Order the edges in E_R by length. We proceed by induction on the the rank of the edge in the ordering to show that if $(u,v) \in E_R$, then there is a path from u to v in G_α. For the base case, if (u,v) is the shortest edge in E_R, then it is immediate from Lemma II.1 that $(u,v) \in E_\alpha$. For note that, by construction, if $(u,v) \in E_R$ and $d(u',v') < d(u,v)$, then $(u',v') \in E_R$ and is a shorter edge than (u,v). For the inductive step, suppose that (u,v) is the kth shortest edge in E_R and, by way of contradiction, that (u,v) is not in E_α. By Lemma II.1, there exist $u',v' \in V$ such that (a) $d(u',v') < d(u,v)$, (b) either $u = u'$ or $(u,u') \in E_\alpha$, and (c) either $v = v'$ or $(v,v') \in E_\alpha$. As we observed, it follows that $(u',v') \in E_R$. Since $d(u',v') < d(u,v)$, by the inductive hypothesis, it follows that there is an path from u' to v' in G_α. Since E_α is symmetric, it is immediate that there is also a path from u to v in G_α. It immediately follows that if u and v are connected in G_R, then there is a path from u to v in G_α. ∎

The proof of Theorem II.2 gives some extra information, which we cull out as a separate corollary:

Corollary II.3: If $\alpha \leq 5\pi/6$, and u and v are nodes in V such that $(u,v) \in E_R$, then either $(u,v) \in E_\alpha$ or there exists a path $u_0 \ldots u_k$ such that $u_0 = u$, $u_k = v$, $(u_i, u_{i+1}) \in E_\alpha$, and $d(u_i, u_{i+1}) < d(u,v)$, for $i = 0, \ldots, k-1$.

Next we prove that degree $5\pi/6$ is a tight upper bound; if $\alpha > 5\pi/6$, then CBTC(α) does not necessarily preserve connectivity.

Theorem II.4: If $\alpha > 5\pi/6$, then CBTC(α) does not necessarily preserve connectivity.

Proof: Suppose $\alpha = 5\pi/6 + \epsilon$ for some $\epsilon > 0$. We construct a graph $G_R = (V, E_R)$ such that CBTC(α) does not preserve the connectivity of this graph. V has eight nodes: $u_0, u_1, u_2, u_3, v_0, v_1, v_2, v_3$. (See Figure 5.) We call u_0, u_1, u_2, u_3 the *u-cluster*, and v_0, v_1, v_2, v_3 the *v-cluster*. The construction has the property that $d(u_0, v_0) = R$ and for $i,j = 0,1,2,3$, we have $d(u_0, u_i) < R$, $d(v_0, v_i) < R$, and $d(u_i, v_j) > R$ if $i + j \geq 1$. That is, the only edge between the u-cluster and the v-cluster in G_R is (u_0, v_0). However, in G_α, the (u_0, v_0) edge disappears, so that the u-cluster and the v-cluster are disconnected.

In Figure 5, s and s' are the intersection points of the circles of radius R centered at u_0 and v_0, respectively. Node u_1 is chosen so that $\angle u_1 u_0 v_0 = \pi/2$. Similarly, v_1 is chosen so that $\angle v_1 v_0 u_0 = \pi/2$ and u_1 and v_1 are on opposite sides of the line $\overline{u_0 v_0}$. Because of the right angle, it is clear that, whatever $d(u_0, u_1)$ is, we must have $d(v_0, u_1) > d(v_0, u_0) = R$; similarly, $d(u_0, v_1) > R$ whatever $d(v_0, v_1)$ is. Next, choose u_2 so

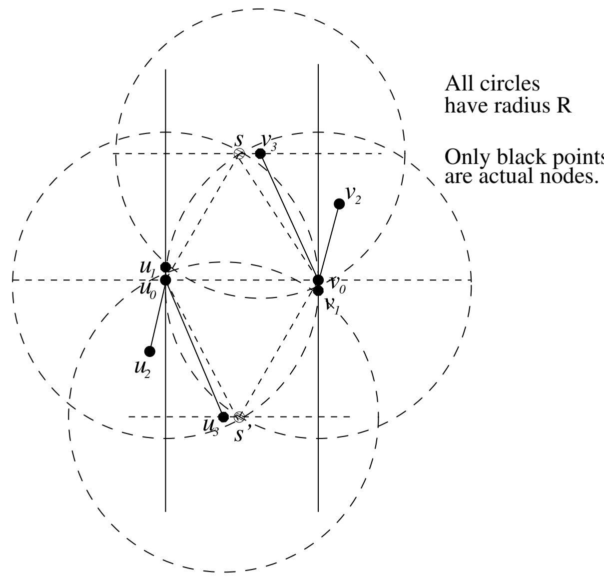

Fig. 5. A disconnected graph if $\alpha = 5\pi/6 + \epsilon$.

that $\angle u_1 u_0 u_2 = \min(\alpha, \pi)$ and $u_0 u_2$ comes after $u_0 u_1$ as a ray sweeps around counterclockwise from $u_0 v_0$. It is easy to see that $d(v_0, u_2) > R$, whatever $d(u_0, u_2)$ is, since $\angle v_0 u_0 u_2 \geq \pi/2$. For definiteness, choose u_2 so that $d(u_0, u_2) = R/2$. Node v_2 is chosen similarly. The key step in the construction is the choice of u_3 and v_3. Note that $\angle s' u_0 u_1 = 5\pi/6$. Let u_3 be a point on the line through s' parallel to $\overline{u_0 v_0}$ slightly to the left of s' such that $\angle u_3 u_0 u_1 < \alpha$. Since $\alpha = 5\pi/6 + \epsilon$, it is possible to find such a node u_3. Since $d(u_0, s') = d(v_0, s') = R$ by construction, it follows that $d(u_0, u_3) < R$ and $d(v_0, u_3) > R$. It is clearly possible to choose $d(v_0, v_1)$ sufficiently small so that $d(u_3, v_1) > R$. The choice of v_3 is similar.

It is now easy to check that when u_0 runs CBTC(α), it will terminate with $p_{u_0,\alpha} = \max(d(u_0, u_3), R/2) < R$; similarly for v_0. Thus, this construction has all the required properties. ∎

III. Optimizations

In this section, we describe three optimizations to the basic algorithm. We prove that these optimizations allow some of the edges to be removed while still preserving connectivity.

A. The shrink-back operation

In the basic CBTC(α) algorithm, u is said to be a *boundary node* if, at the end of the algorithm, u still has an α-gap. Note that this means that, at the end of the algorithm, a boundary node broadcasts with maximum power. An optimization would be to add a shrinking phase at the end of the growing phase to allow each boundary node to broadcast with less power, if it can do so without reducing its cone coverage. To make this precise, given a set *dir* of directions (angles) and an angle α, define $cover_\alpha(dir) = \{\theta$: for some $\theta' \in dir$, $|\theta - \theta'| \bmod 2\pi \leq \alpha/2\}$. We modify CBTC($\alpha$) so that, at each iteration, a node in N_u is tagged with the power used the first time it was discovered. Suppose that the power levels used by node u during the algorithm were $p_1, \ldots, p_k$. If u is a boundary node, p_k is the maximum power p_{max}. A boundary node successively removes nodes tagged with power p_k, then p_{k-1}, and so on, as long as their removal does not change the coverage. That is, let dir_i, $i = 1, \ldots, k$, be the set of directions found with all power levels p_i or less, then the minimum i such that $cover_\alpha(dir_i) = cover_\alpha(dir_k)$ is found. Let $N^s_\alpha(u)$ consist of all the nodes in $N_\alpha(u)$ tagged with power p_i or less. Let $N^s_\alpha = \{(u, v) : v \in N^s_\alpha(u)\}$, and let E^s_α be the symmetric closure of N^s_α. Finally, let $G^s_\alpha = (V, E^s_\alpha)$.

Theorem III.1: If $\alpha \leq 5\pi/6$, then G^s_α preserves the connectivity of G_R.

Proof: It is easy to check that the proof of Theorem II.2 depended only on the cone coverage of each node, so it goes through without change. In more detail, given any two nodes u and v in G^s_α, if $d(u, v) = d \leq R$ and and $(u, v) \notin E^s_\alpha$, then either both u and v did not use power sufficient to reach distance d in the basic CBTC algorithm or one or both of them used enough power to reach distance d but then shrank back. In either case, nodes u and v must still have neighbors in $N^s_\alpha(u)$ and $N^s_\alpha(v)$ fully covering the cones $cone(u, \alpha, v)$ and $cone(v, \alpha, u)$, respectively, since any shrink-back operation can only remove those neighbors that provide redundant cone coverage. Thus, the proof of Lemma II.1 goes through with no change. The remainder of the argument follows exactly the same lines as that of the proof of Theorem II.2. ∎

Note that this argument actually shows that we can remove any nodes from N_u that do not contribute to the cone coverage. However, our interest here lies in minimizing the power needed for broadcast, not in minimizing the number of nodes in N_u. There may be some applications where it helps to reduce the degree of a node; in this case, removing further nodes may be a useful optimization.

B. Asymmetric edge removal

As shown by Example II.1, in order to preserve connectivity, it is necessary to add an edge (u, v) to E_α if $(v, u) \in N_\alpha$, even if $(u, v) \notin N_\alpha$. In Example II.1, $\alpha > 2\pi/3$. This is not an accident. As we now show, if $\alpha \leq 2\pi/3$, not only don't we have to add an edge (u, v) if $(v, u) \in N_\alpha$, we can *remove* an edge (v, u) if $(v, u) \in N_\alpha$ but $(u, v) \notin N_\alpha$. Let $E^-_\alpha = \{(u, v) : (u, v) \in N_\alpha \text{ and } (v, u) \in N_\alpha\}$. Thus, while E_α is the smallest symmetric set containing N_α, E^-_α is the largest symmetric set contained in N_α. Let $G^-_\alpha = (V, E^-_\alpha)$.

Theorem III.2: If $\alpha \leq 2\pi/3$, then G^-_α preserves the connectivity of G_R.

Proof: We start by proving the following lemma, which strengthens Corollary II.3.

Lemma III.3: If $\alpha \leq 2\pi/3$, and u and v are nodes in V such that $(u, v) \in E_R$, then either $(u, v) \in N_\alpha$ or there exists a path $u_0 \ldots u_k$ such that $u_0 = u$, $u_k = v$, $(u_i, u_{i+1}) \in N_\alpha$, and $d(u_i, u_{i+1}) < d(u, v)$, for $i = 0, \ldots, k-1$.

Proof: Order the edges in E_R by length. We proceed by strong induction on the rank of an edge in the ordering. Given an edge $(u, v) \in E_R$ of rank k in the ordering, if $(u, v) \in N_\alpha$, we are done. If not, as argued in the proof of Lemma II.1, there must be a node $w \in N_c(u, \alpha, v) \cap N_\alpha(u)$. Since $\alpha \leq 2\pi/3$, the argument in the proof of Lemma II.1 also shows that $d(w, v) < d(u, v)$. Thus, $(w, v) \in E_R$ and has lower rank in the ordering of edges. Applying the induction hypothesis, the lemma holds for (u, v). This completes the proof. ∎

Lemma III.3 shows that if $(u, v) \in E_R$, then there is a path consisting of edges in N_α from u to v. This is not good enough for our purposes; we need a path consisting of edges in E^-_α. The next lemma shows that this is also possible.

Lemma III.4: If $\alpha \leq 2\pi/3$, and u and v are nodes in V such that $(u, v) \in N_\alpha$, then there exists a path $u_0 \ldots u_k$ such that $u_0 = u$, $u_k = v$, $(u_i, u_{i+1}) \in E^-_\alpha$, for $i = 0, \ldots, k-1$.

Proof: Order the edges in N_α by length. We proceed by strong induction on the rank of an edge in the ordering. Given an edge $(u, v) \in N_\alpha$ of rank k in the ordering, if $(u, v) \in E^-_\alpha$, we are done. If not, we must have $(v, u) \notin N_\alpha$. Since $(v, u) \in E_R$, by Lemma III.3, there is a path from v to u consisting of edges in N_α all of which have length smaller than $d(v, u)$. If any of these edges is asymmetric, i.e. in $N_\alpha - E^-_\alpha$, we can apply the inductive hypothesis to replace the edge by a path consisting only of edges in E^-_α. By the symmetry of E^-_α, such a path from v to u implies a path from u to v. This completes the inductive step. ∎

The proof of Theorem III.2 is now immediate from Lemmas III.3 and III.4. ∎

To implement asymmetric edge removal, the basic CBTC needs to be enhanced slightly. After finishing CBTC(α), a node u must send a message to each node v to which it sent an *Ack* message that is not in $N_\alpha(u)$, telling v to remove u from $N_\alpha(v)$ when constructing E^-_α. It is easy to see that the shrink-back optimization discussed in Section III-A can be applied together with the removal of these asymmetric edges.

There is a tradeoff between using CBTC($5\pi/6$) and using CBTC($2\pi/3$) with asymmetric edge removal. $p(rad^-_{u,5\pi/6})$) will be no greater than $p_{u,2\pi/3}$ if the $Increase$ function is the same, links are reliable, and Acks responding to one "Hello" message are received before the next one is sent. However, the power $p(rad_{u,5\pi/6})$ with which u needs to transmit may be greater than $p_{u,5\pi/6}$, since u may need to reach nodes v such that $u \in N_{5\pi/6}(v)$ but $v \notin N_{5\pi/6}(u)$. In contrast, if $\alpha = 2\pi/3$, then asymmetric edge removal allows u to still use $p_{u,2\pi/3}$ and may allow v to use power less than $p_{v,2\pi/3}$. Our experimental results confirm this. See Section V.

C. Pairwise edge removal

The final optimization aims at further reducing the transmission power of each node. In addition to the directional information, this optimization requires two other pieces of information. First, each node u is assigned a unique integer ID denoted ID_u, and that ID_u is included in all of u's messages. Second, although a node u does not need to know its exact distance from its neighbors, given any pair of neighbors v and w, node u needs to know which of them is closer. This can be achieved as follows. Recall that a node grows its radius in discrete steps. It includes its transmission power level in each of the "Hello" messages. Each discovered neighbor node also includes its transmission power level in the *Ack*. When u receives messages from nodes v_1 and v_2, it can deduce which of v_1 and v_2 is closer based on the transmission and reception powers of the messages.

Even after the shrink-back operation and possibly asymmetric edge removal, there are many edges that can be removed while still preserving connectivity. For example, if three edges form a triangle, we can clearly remove any one of them while still maintaining connectivity. In this section, we improve on this result by showing that if there is an edge from u to v_1 and from u to v_2, then we can remove the longer edge even if there is no edge from v_1 to v_2, as long as $d(v_1, v_2) < \max(d(u, v_1), d(u, v_2))$. Note that a condition sufficient to guarantee that $d(v_1, v_2) < \max(d(u, v_1), d(u, v_2))$ is that $\angle v_1 u v_2 < \pi/3$ (since the longest edge will be opposite the largest angle).

To make this precise, we use the notion of an edge ID. Each edge (u, v) is assigned an edge ID $eid(u, v) = (i_1, i_2, i_3)$, where $i_1 = d(u, v)$, $i_2 = max(\mathrm{ID}_u, \mathrm{ID}_v)$, and $i_3 = min(\mathrm{ID}_u, \mathrm{ID}_v)$. Edge IDs are compared lexicographically, so that $(i, j, k) < (i', j', k')$ iff either (a) $i < i'$, (b) $i = i'$ and $j < j'$, or (c) $i = i'$, $j = j'$, and $k < k'$.

Definition III.5: If v and w are neighbors of u, $\angle vuw < \pi/3$, and $eid(u, v) > eid(u, w)$, then (u, v) is a *redundant edge*.

As the name suggests, redundant edges are redundant, in that it is possible to remove them while still preserving connectivity. The following theorem proves this.

Theorem III.6: For $\alpha \leq 5\pi/6$, all redundant edges can be removed while still preserving connectivity.

Proof: Let E_α^{nr} consist of all the non-redundant edges in E_α. We show that if $(u, v) \in E_\alpha - E_\alpha^{nr}$, then there is a path from u to v consisting only of edges in E_α^{nr}. Clearly, this suffices to prove the theorem.

Let $e_1, e_2, \cdots, e_m$ be a listing of the redundant edges (i.e, those in $E_\alpha - E_\alpha^{nr}$) in increasing lexicographic order of edge ID. We prove, by induction on k, that for every redundant edge $e_k = (u_k, v_k)$ there is a path from u_k to v_k consisting of edges in E_α^{nr}. For the base case, consider $e_1 = (u_1, v_1)$. By definition, there must exist an edge (u_1, w_1) such that $\angle v_1 u_1 w_1 < \pi/3$ and $eid(u_1, v_1) > eid(u_1, w_1)$. Since e_1 is the redundant edge with the smallest edge ID, (u_1, w_1) cannot be a redundant edge. Since $\angle v_1 u_1 w_1 < \pi/3$, it follows that $d(w_1, v_1) < d(u_1, v_1)$. If $(w_1, v_1) \in E_\alpha$, then $(w_1, v_1) \in E_\alpha^{nr}$ (since (u_1, v_1) is the shortest redundant edge) and $(u_1, w_1), (w_1, v_1)$ is the desired path of non-redundant edges. On the other hand, if $(w_1, v_1) \notin E_\alpha$ then, since $d(w_1, v_1) < d(u_1, v_1) \leq R$ and $\alpha \leq 5\pi/6$, by Corollary II.3, there exists a path from w_1 to v_1 consisting of edges in E_α all shorter than $d(w_1, v_1)$. Since none of these edges can be redundant edge, this gives us the desired path.

For the inductive step, suppose that for every $e_j = (u_j, v_j)$, $1 \leq j \leq i - 1$, we have found a path H'_j between u_j and v_j, which contains no redundant edges. Now consider $e_i = (u_i, v_i)$. Again, by definition, there exists another edge (u_i, w_i) with $eid(u_i, v_i) > eid(u_i, w_i)$ and $\angle v_i u_i w_i < \pi/3$. If (u_i, w_i) is a redundant edge, it must be one of e_j's, where $j \leq i - 1$. Moreover, if the path H_i (from Corollary II.3) between v_i and w_i contains a redundant edge e_j, we must have $|e_j| < |e_i|$ and so $j \leq i - 1$. By connecting (u_i, w_i) with H_i and replacing every redundant edge e_j on the path with H'_j, we obtain a path H'_i between u_i and v_i that contains no redundant edges. This completes the proof. ∎

Although Theorem III.6 shows that all redundant edges can be removed, this doesn't mean that all of them should necessarily be removed. For example, if we remove some edges, the paths between nodes become longer, in general. Since some overhead is added for each link a message traverses, having fewer edges can affect network throughput. In addition, if routes are known and many messages are being sent using point-to-point communication between different senders and receivers, having fewer edges is more likely to cause congestion. Since we would like to reduce the transmission power of each node, we remove only redundant edges with length greater than the longest non-redundant edges. We call this optimization the *pairwise edge removal* optimization.

IV. Dealing with Reconfiguration, Asynchrony, and Failures

In a multi-hop wireless network, nodes can be mobile. Even if nodes do not move, nodes may die if they run out of energy. In addition, new nodes may be added to the network. We need a mechanism to detect such changes in the network. This is done by the Neighbor Discovery Protocol (NDP). A NDP is usually a simple beaconing protocol for each node to tell its neighbor that it is still alive. The beacon includes the sending node's ID and the transmission power of the beacon. A neighbor is considered failed if a pre-defined number of beacons are not received for a certain time interval τ. A node v is considered a new neighbor of u if a beacon is received from v and no beacon was received from v during the previous τ interval.

The question is what power a node should use for beaconing. Certainly a node u should broadcast with sufficient power to reach all of its neighbors in E_α (or E_α^-, if $\alpha \leq 2\pi/3$). As we will show, if u uses a beacon with power $p(rad_{u,\alpha})$ —the power needed for u to reach all its neighbors in E_α, then this is sufficient for reconfiguration to work with the basic cone-based algorithm (possibly combined with asymmetric edge removal if $\alpha \leq 2\pi/3$, in which case we can use power $p(rad^-_{u,\alpha})$).

We define three basic events:

- A $join_u(v)$ event happens when node u detects a beacon from node v for the first time;
- A $leave_u(v)$ event happens when node u misses some predetermined number of beacons from node v;
- An $aChange_u(v)$ event happens when u detects that v's angle with respect to u has changed. (Note this could be due to movement by either u or v.)

Our reconfiguration algorithm is very simple. It is convenient to assume that each node is tagged with the power used when it was first discovered, as in the shrink-back operation. (This is not necessary, but it minimizes the number of times that CBTC needs to be rerun.)

- If a $leave_u(v)$ event happens, and if there is an α-gap after dropping $dir_u(v)$ from D_u, node u reruns CBTC(α) (as in Figure 1), starting with power $p(rad^-_{u,\alpha})$ (i.e., taking $p_0 = p(rad^-_{u,\alpha})$).
- If a $join_u(v)$ event happens, u computes $dir_u(v)$ and the power needed to reach v. As in the shrink-back operation, u then removes nodes, starting with the farthest neighbor nodes and working back, as long as their removal does not change the coverage.

• If an $aChange_u(v)$ event happens, node u modifies the set D_u of directions appropriately. If an α-gap is then detected, then CBTC(α) is rerun, again starting with power $p(rad^-_{u,\alpha})$. Otherwise, nodes are removed, as in the shrink-back operation, to see if less power can be used.

In general, there may be more than one change event that is detected at a given time by a node u. (For example, if u moves, then there will be in general several *leave*, $join$ and *aChange* events detected by u.) If more than one change event is detected by u, we perform the changes suggested above as if the events are observed in some order, as long as there is no need to rerun CBTC. If CBTC needs to be rerun, it deals with all changes simultaneously.

Intuitively, this reconfiguration algorithm preserves connectivity. We need to be a little careful in making this precise, since if the topology changes frequently enough, the reconfiguration algorithm may not ever catch up with the changes, so there may be no point at which the connectivity of the network is actually preserved. Thus, what we want to show is that if the topology ever stabilizes, so that there are no further changes, then the reconfiguration algorithm eventually results in a graph that preserves the connectivity of the final network, as long as there are periodic beacons. It should be clear that the reconfiguration algorithm guarantees that each cone of degree α around a node u is covered (except for boundary nodes), just as the basic algorithm does. Thus, the proof that the reconfiguration algorithm preserves connectivity follows immediately from the proof of Theorem II.2.

While this reconfiguration algorithm works in combination with the basic algorithm CBTC(α) and in combination with the asymmetric edge removal optimization, we must be careful in combining it with the other optimizations discussed in Section III. In particular, we must be very careful about what power a node should use for its beacon. For example, if the shrink-back operation is performed, using the power to reach all the neighbors in G^s_α does not suffice. For suppose that the network is temporarily partitioned into two subnetworks G_1 and G_2; for every pair of nodes $u_1 \in G_1$ and $u_2 \in G_2$, the distance $d(u_1, u_2) > R$. Suppose that u_1 is a boundary node in G_1 and u_2 is a boundary node in G_2, and that, as a result of the shrink-back operation, both u_1 and u_2 use power $P' < p_{max}$. Further suppose that later nodes u_1 and u_2 move closer together so that $d(u_1, u_2) < R$. If P' is not sufficient power for u_1 to communicate with u_2, then they will never be aware of each other's presence, since their beacons will not reach each other, so they will not detect that the network has become reconnected. Thus, network connectivity is *not* preserved.

This problem can be solved by having the boundary nodes broadcast with the power computed by the basic CBTC(α) algorithm, namely p_{max} in this case. Similarly, with the pairwise edge removal optimization, it is necessary for u's beacon to broadcast with $p(rad_{u,\alpha})$, i.e., the power needed to reach all of u's neighbors in E_α, not just the power needed to reach all of u's neighbors in E^{nr}_α. It is easy to see that this choice of beacon power guarantees that the reconfiguration algorithm works.

It is worth noting that a reconfiguration protocol works perfectly well in an asynchronous setting. In particular, the synchronous model with reliable channels that has been assumed up to now can be relaxed to allow asynchrony and both communication and node failures. Now nodes are assumed to communicate asynchronously, messages may get lost or duplicated, and nodes may fail (although we consider only *crash* failures: either a node crashes and stops sending messages, or it follows its algorithm correctly). We assume that messages have unique identifiers and that mechanisms to discard duplicate messages are present. Node failures result in *leave* events, as do lost messages. If node u gets a message after many messages having been lost, there will be a *join* event corresponding to the earlier *leave* event.

V. Experimental Results

How effective is our algorithm and its optimizations as compared to other approaches? Before we answer this question, let us briefly review existing approaches. To our knowledge, among the topology-control algorithms in the literature [24], [8], [9], [21], [23], only Rodoplu and Meng's algorithm [23] attempts to optimize for energy efficiency while maintaining network connectivity. Following [13], we refer to Rodoplu and Meng's algorithm as the *MECN* algorithm (for *minimum-energy communication network*). The algorithms in [24], [8], [9] try to maximize network throughput; they do not guarantee network connectivity. Ramanathan and Rosales-Hain [21] have considered minimizing the maximum transmission power of all nodes by using centralized MST algorithms. However, their distributed heuristic algorithms do not guarantee network connectivity. Since we are only interested in algorithms that preserve connectivity and are energy efficient, it seems that the only relevant algorithm in the literature is the MECN algorithm. However, since the SMECN algorithm outperforms MECN [13], we will compare our algorithm with SMECN only.

We refer to the basic algorithm as CBTC, and to our complete algorithm with all applicable optimizations as OPT-CBTC.[3] Furthermore, we also make the comparison with the no-topology-control case, where each node always uses the maximum transmission power to send a packet (we refer to this approach as MaxPower). In the case of no-topology-control, the reason we choose maximum power is that it guarantees that there will be no network partitions due to insufficient transmission power.

A. Simulation Environment

The topology-control algorithms – CBTC, SMECN and MaxPower – are implemented in the ns-2 network simulator [20], using the wireless extension developed at Carnegie Mellon [6]. We generated 20 random networks, each with 200 nodes. Each node has a maximum transmission range of 500 meters and initial energy of 0.5 Joule. The nodes are placed uniformly at random in a rectangular region of 1500 by 1500 meters. Although there have been some papers on realistic topology generation [31], [1], most of them have focused on the Internet setting. Since large multihop wireless networks such as sensor networks are often deployed in a somewhat random fashion (for example, an airplane may drop sensors over some geographical region), we believe that assuming nodes are placed uniformly at random is not an unreasonable assumption.

[3] For brevity, we will omit the parameter α in our presentation when it is clear from the context.

We assume the two-ray propagation model for terrestrial communication [22]. A transmission from node u to node v takes power $p(u,v) = td(u,v)^n$ for some constant t at node u, where $n \geq 2$ is the path-loss exponent of outdoor radio propagation models, and $d(u,v)$ is the distance between u and v. The model has been shown to be close to reality in many environment settings [22]. Finally, we take the following parameter settings, which are chosen to simulate the 914MHz Lucent WaveLAN DSSS radio interface:

- the carrier frequency is 914MHz;
- the transmission raw bandwidth 2MHz;
- antennas are omni-directional with 0dB gain, and the antenna is placed 1.5 meters above a node;
- the receive threshold is -94dBW;
- the carrier sense threshold is -108dBW;
- the capture threshold is 10dB.

In order to simulate the effect of power control in the neighbor-discovery process, we made changes to the physical layer of the ns-2 simulation code to support eight discrete power levels. This seems to be more in keeping with current practice. For example, currently the Aironet PC4800 supports five transmission-power levels. Eight power levels seems sufficient to provide a realistic simulation of the kind of scenarios that arise in practice. In our simulation, power level 8 gives the maximum transmission range of 250 meters. The *Increase* function in Figure 1 moves from one power level to the next higher level. For the "Hello" packet in the CBTC algorithm, the transmission power level is controlled by the algorithm itself. Specifically, as we discussed in Section IV, node u broadcasts using the final power p_u (as determined by the *Increase* function in Figure 1). For point-to-point transmissions from a node u, the minimum power level needed to reach all of u's neighbors is used. We do not use different power levels for different neighbors because there is a delay associated with changing power levels in practice (in the order of 10 milliseconds [5] for certain wireless radio hardware), which some applications may not be able to tolerate.

To simulate interference and collision, we choose the WaveLAN-I [25] CSMA/CA MAC protocol. Since topology control by itself does not provide routing, we used the AODV [17] routing protocol in our simulation.

To simulate the network application traffic, we use the following application scenario: we choose 60 connections, i.e. 60 source-destination pairs. All the source and destination nodes are distinct. For each of these 60 connections in sequence, the source (if it is still alive) sends constant bit rate (CBR) traffic to its destination. The sending rate is 2 packets/sec and the packet size is 512 bytes. This traffic pattern seems to generate sufficient load in the network for our evaluation. We do not expect that the results would be qualitatively different if fewer or more connections were used. We use the same 60 connections in all our experiments. Since we conduct the experiments in 20 random networks, there is no need to randomize over the connections as well.

B. Network Topology Characteristics

Before comparing CBTC with SMECN and MaxPower through detailed network simulation, we first examine the topology graphs that result from using each of these approaches in the 20 random networks described previously.

Figure 6 illustrates how CBTC and the various optimizations improve network topology using the results from one of the random networks. Figure 6(a) shows a topology graph produced by MaxPower. Figures 6(b) and (c) show the corresponding graphs produced by CBTC($2\pi/3$) and CBTC($5\pi/6$), respectively. We can see that both CBTC($2\pi/3$) and CBTC($5\pi/6$) allow nodes in the dense areas to automatically reduce their transmission radius. Figures 6(d) and (e) illustrate the graphs after the shrink-back operation is performed. Figure 6(f) shows the graph for $\alpha = 2\pi/3$ as a result of the shrink-back operation and the asymmetric edge removal. Figures 6(g) and (h) show the topology graphs after all applicable optimizations. We can see that the optimizations are very effective in further reducing the transmission radius of nodes.

Table I compares the network graphs resulted from the cone-based algorithm parameterized by $\alpha = 2\pi/3$ and $\alpha = 5\pi/6$, in terms of average node degree and average radius. It also shows the corresponding results for SMECN and MaxPower. The results are averaged over the 20 random networks mentioned earlier. As expected, using a larger value of α results in a smaller node degree and radius. However, as we discussed in Section III-B, there is a tradeoff between using CBTC($2\pi/3$) and CBTC($5\pi/6$). Using the basic algorithm, we have $rad_{u,5\pi/6} = 205.4 < rad_{u,2\pi/3} = 220.6$. After applying asymmetric edge removal with $\alpha = 2\pi/3$, the resulting radius is 176.6. Hence, asymmetric edge removal can result in significant savings. After applying all applicable optimizations, both $\alpha = 2\pi/3$ and $\alpha = 5\pi/6$ end up with very similar results in terms of both average node degree and average radius. However, there are secondary advantages to setting $\alpha = 5\pi/6$. In general, CBTC($5\pi/6$) will terminate sooner than CBTC($2\pi/3$) and so expend less power during its execution (since $p_{u,5\pi/6} < p_{u,2\pi/3}$). Thus, if reconfiguration happens frequently, the advantage of using CBCT($5\pi/6$) over CBCT($2\pi/3$) in terms of reduction on power consumption can be significant.

The sixth row (MaxPower) gives the performance numbers for the case where each node uses the maximum transmission power of $p(250)$. We can see that using topology control cuts down the average degree by a factor of more than 3 (3.8 vs. 15.0) and the average radius by a factor of more than 2 (113.1 or 110.7 vs. 250). This clearly demonstrates the effectiveness of our topology-control algorithms.

The last row shows the results for SMECN. Recall that SMECN requires GPS position information, while the CBTC algorithms rely on only directional information. So our objective in the comparison is to study how well CBTC performs with the lack of distance information. The average radius numbers in Table I show that the performance of OPT-CBTC is in fact very close to (and slightly better than) that of SMECN (113.1 vs. 115.8). Note that SMECN does achieve a smaller average node degree (2.7 vs. 3.7). However, with SMECN, each node typically has more nodes within its radius that are not its neighbors. This is because for a node v to be considered a neighbor of u in SMECN, direct transmission has to take less energy than any two-hop path. Two-hop paths are less desirable than single-hop paths, they occupy the media for twice as long as one-hop transmissions. On the other hand, although OPT-CBTC reduces the

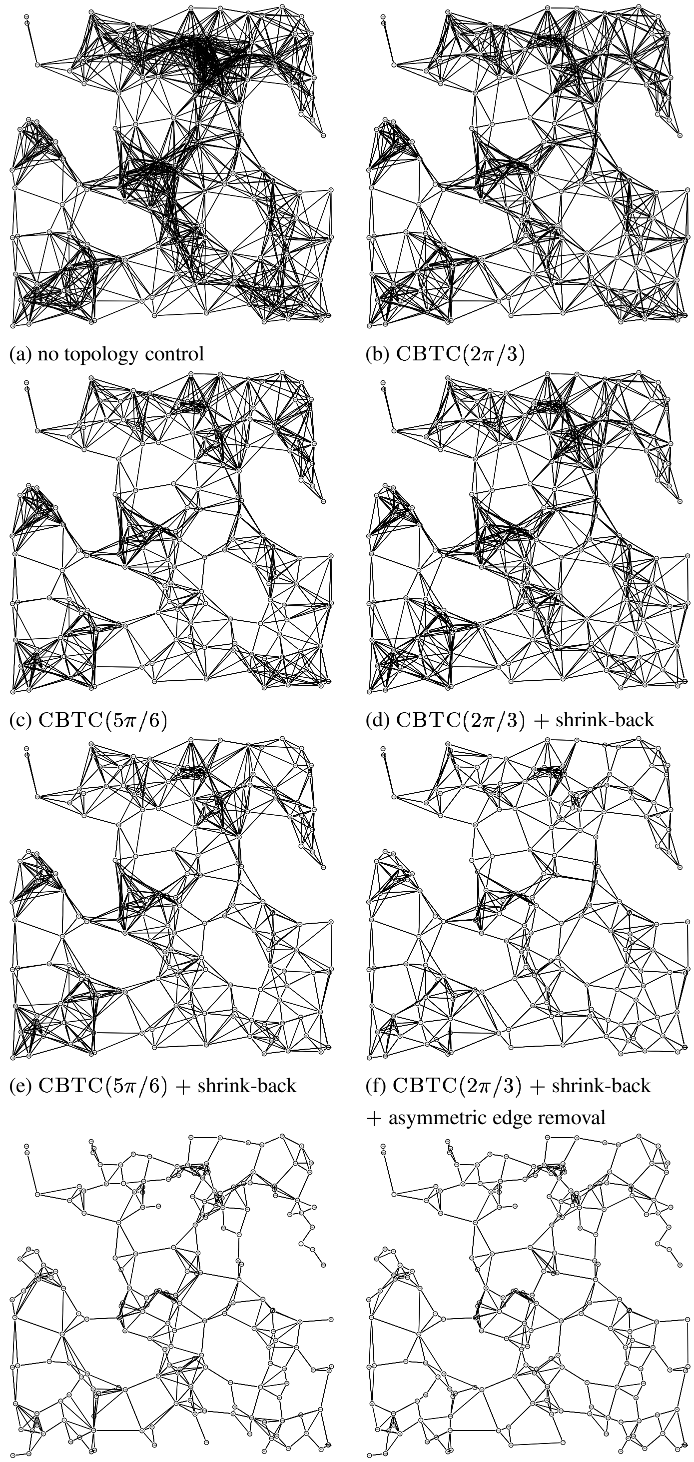

(a) no topology control

(b) $\mathrm{CBTC}(2\pi/3)$

(c) $\mathrm{CBTC}(5\pi/6)$

(d) $\mathrm{CBTC}(2\pi/3)$ + shrink-back

(e) $\mathrm{CBTC}(5\pi/6)$ + shrink-back

(f) $\mathrm{CBTC}(2\pi/3)$ + shrink-back + asymmetric edge removal

(g) $\mathrm{OPT} - \mathrm{CBTC}(5\pi/6)$

(h) $\mathrm{OPT} - \mathrm{CBTC}(2\pi/3)$

Fig. 6. The network graphs after different optimizations.

		Average Node Degree	Average radius
Basic	$\alpha = 5\pi/6$	8.8	205.4
	$\alpha = 2\pi/3$	10.9	220.6
with op_1	$\alpha = 5\pi/6$	8.3	194.3
	$\alpha = 2\pi/3$	10.1	209.4
with op_2	$\alpha = 2\pi/3$	6.9	176.6
with op_1 and op_2	$\alpha = 2\pi/3$	6.7	171.8
with all optimizations (OPT-CBTC)	$\alpha = 5\pi/6$	3.8	110.7
	$\alpha = 2\pi/3$	3.7	113.1
MaxPower	N/A	15.0	250
SMECN	N/A	2.7	115.8

TABLE I

AVERAGE DEGREE AND RADIUS OF THE CONE-BASED TOPOLOGY-CONTROL ALGORITHM WITH DIFFERENT α AND OPTIMIZATIONS (op_1–SHRINK-BACK, op_2–ASYMMETRIC EDGE REMOVAL).

power demand of nodes as much as SMECN does, SMECN has the additional property of preserving minimum-energy paths. If a different power level can be used for each neighbor, and the amount of unicast traffic is significantly greater than the amount of neighbor broadcast traffic, using SMECN can be beneficial.

C. Network Performance Analysis

We next use detailed network simulations to evaluate the algorithms in terms of energy consumption, number of delivered packets, and latency. Since the two CBTC settings $\alpha = 5\pi/6$ and $\alpha = 2\pi/3$ produced similar network graphs (as shown in Table I), we consider only $\alpha = 2\pi/3$ in the remaining experiments. [4] We simulate CBTC, MaxPower, and SMECN using the same traffic pattern and random networks for performance measurements. As the power available to a node is decreased after each packet reception or transmission, nodes in the simulation die over time. After a node dies, the network must be reconfigured. In our simulation, the NDP beacons trigger the reconfiguration protocol. The beacons are sent once per second for SMECN and CBTC, and each of them is jittered randomly before it is actually sent to avoid synchronization effects. For CBTC and OPT-CBTC, the beacons use power $p(rad^-_{u,2\pi/3})$. For SMECN, the beacons use the appropriate power level as computed by SMECN's neighbor discovery process. Note that no beacon is required in the MaxPower approach. For simplicity, we do not simulate node mobility, although some of the effects of mobility—that is, the triggering of the reconfiguration protocol—can already be observed when nodes run out of energy. In the rest of this section, we compare the performance of CBTC, OPT-CBTC, SMECN, and MaxPower. All results are averaged over the 20 random networks described in Section V-A.

C.1 Energy Consumption

We investigate the energy consumption of the three approaches in terms of the number of traffic sources alive and the average transmission power levels over time. As can be seen from Figure 7, OPT-CBTC has the best performance. CBTC performs worse than the SMECN algorithm, but uses only directional information. MaxPower has significantly worse performance than the other algorithms. Figure 7(a) shows the number of traffic sources that remain alive over time. We can see that when almost all the traffic sources in MaxPower are dead at time 600, about 45% and 30% of the traffic sources are still alive in SMECN and CBTC, respectively, and more than 79% of the traffic sources are still alive in OPT-CBTC. The basic CBTC algorithm does not perform as well as OPT-CBTC, but it still performs much better than MaxPower.

Next, we consider how the transmission power evolves over time as nodes die over time. Figure 7(b) shows the average power level averaged over all nodes. The "average power level" at time t is computed by considering, for each node u still alive at time t, the minimum power currently needed for u to reach all its neighbors (recall that this is the power that u uses in the simulation to send all messages except the NDP "Hello" beacons), and then averaging this number over all nodes still alive. For MaxPower, the average power is constant over time because the maximum power is always used. The curves show that, while the average power level of CBTC and SMECN increases rapidly over time as more nodes die, the power level of OPT-CBTC increases rather slowly and remains much lower.

C.2 Total Number of Packets Delivered and Latency

We collected packet delivery and latency statistics at the end of our simulation. SMECN, CBTC and OPT-CBTC were able to deliver 1.66, 1.44, and 2.94 times the amount of packets delivered by MaxPower, respectively, throughout the simulation. The statistics for packet delivery and the number of traffic sources still alive together show that it is undesirable to transmit with large radius because it increases energy consumption and causes unnecessary interference, and consequently decreases throughput. The average packet latencies in decreasing order are 271, 170, 126 and 79 msec for MaxPower, OPT-CBTC, CBTC and SMECN, respectively. MaxPower has the highest latency due to its low spatial reuse of the spectrum. That is, a successful transmission by MaxPower reserves a large physical area. Any node that hears the transmission within this area backs off and does not transmit itself. Therefore, the larger the area reserved,

[4] Since we use only a few discrete power levels, there is no significant benefit in using $\alpha = 5\pi/6$.

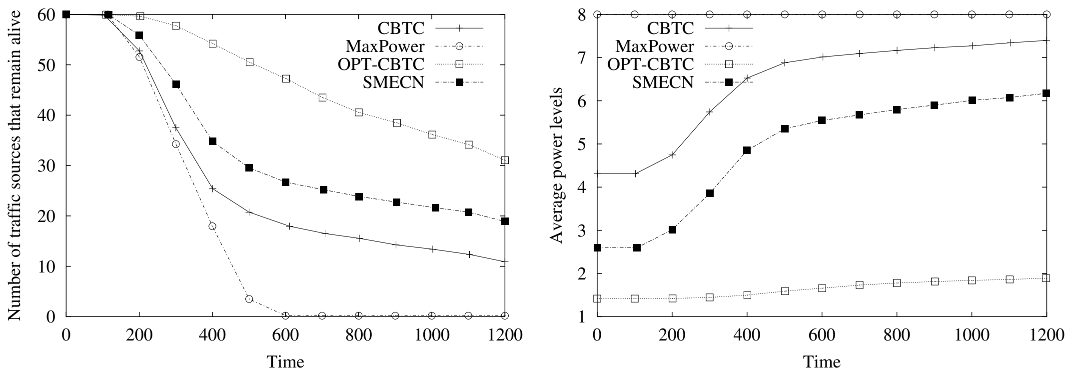

(a) Number of traffic sources that remain alive

(b) Average power levels

Fig. 7. Performance comparison through detailed network simulation.

the fewer nodes can transmit at any particular time. OPT-CBTC has higher latency than CBTC and SMECN because it typically takes longer routes due to the use of lower transmission power.

VI. CONCLUSION

We have analyzed the distributed cone-based algorithm and proved that $5\pi/6$ is a tight upper bound on the cone degree for the algorithm to preserve connectivity. We have also presented three optimizations to the basic algorithm—the shrink-back operation, asymmetric edge removal, and pairwise edge removal—and proved that they improve performance while still preserving connectivity. Finally, we showed that there is a tradeoff between using CBTC(α) with $\alpha = 5\pi/6$ and $\alpha = 2\pi/3$, since using $\alpha = 2\pi/3$ allows an additional optimization, which can have a significant impact on reducing the transmission radius. The algorithm extends easily to deal with reconfiguration and asynchrony. Most importantly, simulation results show that it is very effective in reducing power demands and increases the overall throughput.

Since the focus of this paper has been on reducing energy consumption, we conclude with some discussion of this goal. Reducing energy consumption has been viewed as perhaps the most important design metric for topology control. There are two standard approaches to doing this: (1) reducing the transmission power of each node as much as possible; (2) reducing the total energy consumption through the preservation of minimum-energy paths in the underlying network. These two approaches may conflict: reducing the transmission power required by each node may not result in minimum-energy paths or vice versa. Furthermore, there are other metrics to consider, such as network throughput and network lifetime. Reducing energy consumption tends to increase network lifetime. (This is particularly true if the main reason that nodes die is loss of battery power.) However, there is no guarantee that it will. For example, using minimum-energy paths for all communication may result in hot spots and congestion, which in turn may drain battery power and lead to network partition. Using approach (1) in this case may do better. If topology control is not done carefully, network throughput can be hurt. As we have already pointed out, eliminating edges may result in more congestion and hence worse throughput, even if it saves power in the short run. The right tradeoffs to make are very much application dependent. Therefore, an algorithm that adapts to the specific application setting is much needed. Reconfiguration in response to node mobility and failure consumes precious energy resources. Fast convergence of topology control is critical to keep the network functioning well. It would be interesting to investigate how much mobility CBTC can handle. We hope to explore these issues in more detail in future work.

REFERENCES

[1] K. Calvert, M. Doar, and E. W. Zegura. Modeling internet topology. *IEEE Communications Magazine*, 35(6):160–163, June 1997.

[2] A. Chandrakasan, R. Amirtharajah, S. H. Cho, J. Goodman, G. Konduri, J. Kulik, W. Rabiner, and A. Wang. Design considerations for distributed microsensor systems. In *Proc. IEEE Custom Integrated Circuits Conference (CICC)*, pages 279–286, May 1999.

[3] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proc. seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 85–96, 2001.

[4] L. P. Clare, G. J. Pottie, and J. R. Agre. Self-organizing distributed sensor networks. In *Proc. SPIE Conf. on Unattended Ground Sensor Technologies and Applications*, pages 229–237, April 1999.

[5] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian. Conserving transmission power in wireless ad hoc networks. In *Proc. International Conference on Network Protocols (ICNP)*, pages 24–34, November 2001.

[6] CMU Monarch Group. Wireless and mobility extensions to ns-2. http://www.monarch.cs.cmu.edu/cmu-ns.html, October 1999.

[7] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless micro-sensor networks. In *Proc. IEEE Hawaii Int. Conf. on System Sciences*, pages 4–7, January 2000.

[8] T.-C. Hou and V. O. K. Li. Transmission range control in multihop radio networks. *IEEE Trans. on Communications*, 34(1):38–44, January 1986.

[9] L. Hu. Topology control for multihop packet radio networks. *IEEE Trans. on Communications*, 41(10):1474 –1481, October 1993.

[10] L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. In *Proc. SPAA*, pages 220–229, 2003.

[11] R. E. Kahn. The organization of computer resources into a packet radio network. *IEEE Transactions on Communications*, COM-25(1):169–178, January 1977.

[12] K. Krizman, T. E. Biedka, and T.S. Rappaport. Wireless position location: fundamentals, implementation strategies, and source of error. In *IEEE 47th Vehicular Technology Conference*, pages 919–923, 1997.

[13] L. Li and J. Y. Halpern. Minimum energy mobile wireless networks revisited. In *Proc. IEEE International Conference on Communications (ICC)*, pages 278–283, June 2001.

[14] L. Li, J. Y. Halpern, P. Bahl, Y. M. Wang, and R. Wattenhofer. Analysis of distributed topology control algorithms for wireless multi-hop networks. In *Proc. ACM Symposium on Principle of Distributed Computing (PODC)*, pages 264–273, August 2001.

[15] J. Monks, V. Bharghavan, and W. Hwu. A power controlled multiple access protocol for wireless packet networks. In *Proc. IEEE Infocom*, pages 219–228, April 2001.
[16] C. E. Perkins. *Ad Hoc Networking*. Addison-Wesley, Reading, MA, 2001.
[17] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *Proc. 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, February 1999.
[18] G. J. Pottie and W. J. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.
[19] R. Prakash. Unidirectional links prove costly in wireless ad-hoc networks. In *Proc. International workshop on Discrete algorithms and methods for mobile computing and communications (dialM)*, pages 15–22, 1999.
[20] VINT Project. The UCB/LBNL/VINT network simulator-ns (Version 2). http://www.isi.edu/nsnam/ns.
[21] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proc. IEEE Infocom 2000*, pages 404–413, March 2000.
[22] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.
[23] V. Rodoplu and T. H. Meng. Minimum energy mobile wireless networks. *IEEE J. Selected Areas in Communications*, 17(8):1333–1344, August 1999.
[24] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Trans. on Communications*, 32(3):246–257, March 1984.
[25] Bruce Tuch. Development of WaveLAN, an ISM band wireless LAN. *AT&T Technical Journal*, 72(4):27–33, 1993.
[26] Y. Wang and X.Y. Li. Localized construction of bounded degree and planar spanner for wireless ad hoc networks. In *Proc. ACM DIALM-POMC Joint Workshop on Foundations of Mobile Computing*, pages 59–68, 2003.
[27] R. Wattenhofer, L. Li, P. Bahl, and Y. M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proc. IEEE Infocom*, pages 1388–1397, April 2001.
[28] R. Wattenhofer and A. Zollinger. XTC: A practical topology control algorithm for ad-hoc networks. Technical Report 407, Computer Science Department, ETH Zurich, 2003.
[29] S. L. Wu, Y. C. Tseng, and J. P. Sheu. Intelligent medium access control for mobile ad hoc networks with busy tones and power control. *IEEE Journal on Selected Areas in Communications*, 18(9):1647–1657, 2000.
[30] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad-hoc routing. In *Proc. seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 70–84, 2001.
[31] E. W. Zegura, K. Calvert, and S. Bhattacharjee. How to model an internetwork. In *Proc. IEEE Infocom*, volume 2, pages 594–602, 1996.

APPENDIX

I. PROOF FOR THEOREM II.2

Fact A.1: The distance d between any two points u, v in a $\beta < \pi/3$ sector of a circle is no greater than the circle radius r. If both u and v are not the center of the circle, then $d < r$.

Lemma A.2: In Figure 8, $circ(z, d)$ intersects $circ(v, d)$ on the arc from u clockwise to q at point t.

Proof: For any two points t', t'' on the arc from u clockwise to q, if $\angle t'vz > \angle t''vz$, then $d(t', z) > d(t'', z)$. This follows from a simple geometry argument. Consider triangles $\triangle t'vz$ and $\triangle t''vz$. Since $d(t', v) = d(t'', v) = d$ and the triangles have one side vz in common, $\angle t'vz > \angle t''vz$ implies $d(t', z) > d(t'', z)$. Since $d(u, z) \geq d$ (by assumption) and $d(q, z) < d$ (by Fact A.1), there must be a point t on the arc from u clockwise to q such that $d(t, z) = d$. ∎

Lemma A.3: Let line $\overline{ut}$ intersect $circ(u, d)$ at point f (if t is the same as u, then $\angle fuv = \pi/2$) in Figure 8. To cover $cone(u, \alpha, v)$, in the case of $d(w, v) \geq d$ of Lemma II.1, u must have at lease one neighbor in sector $\widehat{huq}$ of $circ(u, d)$ and outside $circ(v, d)$. Among these neighbors, let w be the one such that $\angle wuv$ is the smallest. w cannot lie within the $cone(u, \angle fuv, v)$.

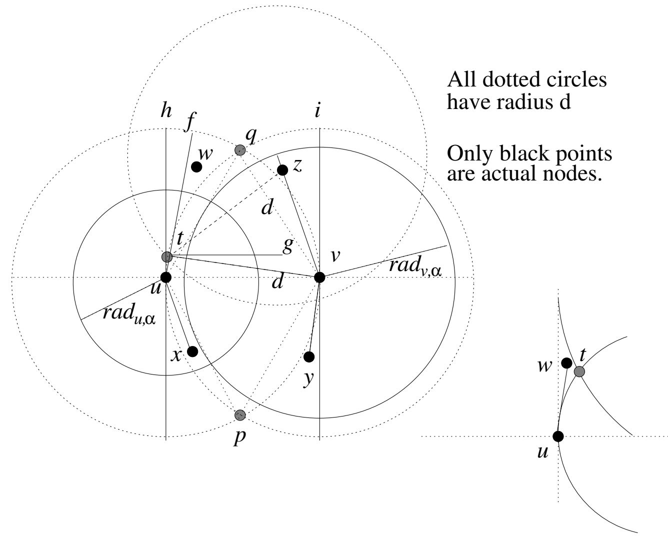

Fig. 8. Illustration for the proof of Lemma A.2 and Lemma A.3.

Proof: For the case of $d(w, v) \geq d$ of Lemma II.1, we only need to show that w cannot lie within the ftq region (the region inside sector $\widehat{fuq}$ of $circ(u, d)$ and outside of $circ(v, d)$). We prove by contradiction. Suppose w lies in that region. By the previous lemma, t lies in the arc from u to q. So both t and w are in the sector $\widehat{huq}$ of $circ(u, d)$. By Fact 1, $d(w, t) \leq d$. Our assumption is that $d(w, z) \geq d$. Thus, $d(w, z) \geq d(t, z) = d \geq d(w, t)$. Therefore,

$$\angle wtz \geq \pi/3 \tag{3}$$

Since $d(t, z) = d(t, v) = d$ (t is the intersection of $circ(z, d)$ and $circ(v, d)$),

$$\angle ztv = \pi - 2 * \angle zvt \tag{4}$$

Since z is inside $cone(v, \alpha, u)$,

$$\angle zvt \leq \alpha/2 - \angle tvu \leq 5\pi/12 - \angle tvu \tag{5}$$

Draw a line $\overline{tg}$ parallel to $\overline{uv}$. We have $\angle wtg = \angle wtz + \angle ztv - \angle vtg$. By Equation 3, $\angle wtg \geq \pi/3 + \angle ztv - \angle tvu$. By Equation 4 and 5, $\angle wtg \geq \pi/2 + \angle tvu$. Since $\angle tuv = (\pi - \angle tvu)/2 = \pi/2 - \angle tvu/2$, we have $\angle wtg \geq \angle tuv$. This contradicts our assumption of w's position. Thus, w must be outside $cone(u, \angle fuv, v)$. ∎

Li (Erran) Li (M '99; ACM M '99) received a B.E. in Automatic Control from Beijing Polytechnic University in 1993, a M.E. in Pattern Recognition from the Institute of Automation, Chinese Academy of Sciences, in 1996, and a Ph.D. in Computer Science from Cornell University in 2001 where Joseph Y. Halpern was his advisor. During his graduate study at Cornell University, he worked at Microsoft Research and Bell Labs Lucent as an intern, and at AT&T Research Center at ICSI Bekerley as a visiting student. He is presently a member of the Networking Research Center in Bell Labs. His research interests are in networking with a focus on wireless networking and mobile computing. His email address is erranlli@dnrc.bell-labs.com

Joseph Y. Halpern (SM '00; ACM F '02) received a B.Sc. in mathematics from the University of Toronto in 1975 and a Ph.D. in mathematics from Harvard in 1981. In between, he spent two years as the head of the Mathematics Department at Bawku Secondary School, in Ghana. He is currently a professor of computer science at Cornell University, where he has been since 1996. Together with his former student, Yoram Moses, he pioneered the approach of applying reasoning about knowledge to analyzing distributed protocols and multi-agent systems; he won a Gödel Prize for this work. He received the Publishers' Prize for Best Paper at at the International Joint Conference on Artificial Intelligence in 1985 (joint with Ronald Fagin) and in 1989. He has coauthored 6 patents, two books ("Reasoning About Knowledge" and "Reasoning About Uncertainty"), and over 100 journal publications and 100 conference publications. He is a former editor-in-chief of the Journal of the ACM. His email address is halpern@cs.cornell.edu

Victor Bahl (SM '97; ACM F '03) is a Senior Researcher and the Manager of the Networking Group in Microsoft Research. His research interests span a variety of problems in wireless networking including low-power RF communications; ubiquitous wireless Internet access and services; location determination techniques and services; self-organizing, self-managing multi-hop community mesh networks; and real-time audio-visual communications. He has authored over 65 scientific papers, 44 issued and pending patent applications, and book chapters in these areas. He is the founder and Chairman of the ACM Special Interest Group in Mobility (SIGMOBILE); the founder and past Editor-in-Chief of ACM Mobile Computing and Communications Review, and the founder and Steering Committee Chair of ACM/USENIX Mobile Systems Conference (MobiSys); He has served on the editorial board of the IEEE Journal on Selected Areas in Communications, and is currently serving on the editorial boards of Elsevier's Adhoc Networking Journal, Kluwer's Telecommunications Systems Journal, and ACM's Wireless Networking Journal. He has served as a guest editor for several IEEE and ACM journals and on networking panels and workshops organized by the National Science Foundation (NSF), the National Research Council (NRC) and European Union's COST. He has served as the General Chairman, Program Chair and Steering Committee member of several IEEE and ACM conferences and on the Technical Program Committee of over 40 international conferences and workshops. He is the recipient of Digital's Doctoral Engineering Award (1994) and the ACM SIGMOBILE's Distinguished Service Award (2001). Dr. Bahl received his Ph.D. in Computer Systems Engineering from the University of Massachusetts Amherst.

Yi-Min Wang received his B.S. degree from the Department of Electrical Engineering at National Taiwan University in 1986, and his Ph.D. degree from the Department of Electrical and Computer Engineering at University of Illinois at Urbana-Champaign in 1993, where he received the Robert T. Chien Memorial Award from the Graduate College for excellence in research. From 1993 to 1997, he was with AT&T Bell Labs and worked primarily in the area of checkpointing and rollback recovery, both in theory and practice. Since he joined Microsoft Research in 1998, Dr. Wang has expanded his research efforts into distributed systems and home networking. He is currently a Senior Researcher in the Systems and Networking group, leading an R&D effort in systems management and diagnostics.

Roger Wattenhofer (ACM M '99) received his doctorate in computer science in 1998 from ETH Zurich, Switzerland. From 1999 to 2001 he was in the USA, first at Brown University in Providence, RI, then at Microsoft Research in Redmond, WA. Currently he is an assistant professor at ETH Zurich. Roger Wattenhofer's research interests include a variety of algorithmic aspects in networking and distributed computing; in particular, peer-to-peer computing and ad-hoc networks. His email address is wattenhofer@inf.ethz.ch