Robust Topology Control Protocols*

Sukumar Ghosh Kevin Lillis Saurav Pandit Sriram Pemmaraju

August 5, 2004

Abstract

Topology control protocols attempt to reduce the energy consumption of nodes in an ad-hoc wireless network while maintaining sufficient network connectivity. Topology control protocols with various features have been proposed, but they all lack robustness and are extremely sensitive to faulty information from neighbors. For example, the XTC protocol (R. Wattenhofer and A. Zellinger, XTC: A practical topology control algorithm for ad-hoc networks, *WMAN 2004*) can be forced to construct a disconnected network even if two nodes in the network receive slightly faulty distance information from one neighbor each. A key step in most localized topology control protocols is one in which each node establishes a total ordering on its set of neighbors based on information received from them. In this paper, we propose a metric for *robustness* of localized topology control protocols and define an r-robust topology control protocol as one tha returns a correct output network even when its neighborhood orderings have been modified by upto $r-1$ adjacent swaps by a malicious adversary. We then modify XTC is a simple manner to derive a family of r-robust protocols for any $r > 1$. The price we pay for increased robustness is in terms of decreased network sparsity; however we can bound this decrease and we show that in transforming XTC from a 1-robust protocol (which it trivially is) into an r-robust protocol, the maximum vertex degree of the output network increases by a factor of $O(\sqrt{r})$. An extremely pleasant side-effect of our design is that the output network is both $\Omega(\sqrt{r})$-edge connected and $\Omega(\sqrt{r})$-vertex connected provided the input network is. Thus ensuring robustness of the protocol seems to give fault-tolerance of the output for free. Our r-robust version of XTC is almost as simple and practical as XTC and like XTC it only involves 2 rounds of communication between a node and its neighbors.

Keywords: Ad-hoc wireless networks, fault-tolerance, k-connectovity, robustness, topology control protocols.

1 Introduction

Ad-hoc wireless networks consist of autonomous devices or nodes communicating with each other by radio. Typically, each of these nodes has access to a tiny power source and this imposes stringent constraints on the amount of energy that a node can use for communicating with other nodes. *Topology control* protocols attempt to reduce the power consumption of nodes in order to increase the life of the network. Typically, the energy required by a node s to transmit a message to a node t increases at least quadratically with the distance between s and t. As a consequence, power consumption is significantly reduced if messages from s to t were routed through a sequence of intermediate nodes, such that the distance between consecutive nodes in the path is small. Topology control protocols choose a transmission power level for each node so that a node communicates with just a few nearby nodes. Reducing transmission power level also reduces collisions and therefore saves energy by reducing the number of retransmissions. However, the local choice of transmission power level for each node has to

*All authors are at the Department of Computer Science, The University of Iowa, Iowa City, IA 52242-1419. Contact e-mail addresses are [ghosh, lillis, spandit, sriram]@cs.uiowa.edu.

be such that the induced network topology satisfies certain global properties such as connectivity and the presence of multiple short paths between pairs of nodes. The two primary goals of topology control: (i) reducing transmission power level to save energy and (ii) maintaining connectivity and redundancy of short paths to increase routing efficiency, are clearly in conflict with each other. Any satisfactory solution to the topology control problem needs to address this key difficulty.

Let $G = (V, E)$ denote the ad-hoc network with vertex set V denoting the set of nodes and edge set E denoting the set of communication links. Let $c : E \to \mathbf{R}^+$ be a cost function that associates a non-negative real cost to each edge $e \in E$. For each vertex $u \in V$, let $N(u)$ denote the neighbors of u in G. During the course of a topology control protocol P, each vertex $u \in V$ chooses a subset $N_P(u) \subseteq N(u)$ of vertices to transmit to. Letting E_P denote the set of directed edges $\{(u, v) \mid u \in V, v \in N_P(u)\}$, we can view the output of P as the directed spanning subgraph $G_P = (V, E_P)$ of G. Typically, it is desired that G_P satisfy the following properties.

Symmetry If $v \in N(u)$ then $u \in N(v)$. As pointed out by [6, 8], without symmetry even the simple task of providing an ACK in response to a message received can become quite cumbersome. Symmetry implies that G_P can be viewed as an undirected graph. There is of course some cost to requiring symmetry, but this is not a property that is very difficult to impose. In describing the rest of the desired properties we assume that G_P is undirected.

Sparseness This property is typically quantified as $|E_P| = O(|V|)$. Often, a stronger property, that of bounded degree is desired. This property requires that for all vertices u, $|N(u)| \leq c$ for some constant c. Burkhart et. al. [2] point out that sparseness is often assumed to guarantee low interference, and while this may be true in an "average case" sense, it is not true in general. [2] also presents a reasonable definition of a metric for interference and one may, in addition to sparseness, require that G_P minimize this interference metric.

Connectivity G_P is required to be connected, provided G that is connected. Often, stronger versions of connectivity such as k-edge connectivity or k-vertex connectivity (for $k > 1$) are desired. These stronger versions of connectivity imply that G_P has multiple paths for routing between pairs of vertices and is more fault-tolerant to link or vertex failures.

Spanner Property For any pair of vertices u and v, let $C(u, v)$ (respectively, $C_P(u, v)$) denote the cost of the cheapest path between u and v in G (respectively, G_P). Then, the spanner property requires the existence of a constant t such that $C_P(u, v) \leq t \cdot C(u, v)$ for all pairs of vertices $u, v \in V$. If such a constant t exists, then G_P is called a *t-spanner* of G.

Less typically, certain other properties such as planarity of G_P are also desired. If G_P is planar, then geometric routing algorithms such as $\mathbf{GOAFR}^+$ provide efficient routing in the network [3].

Definitions and Notation. In addition to costs, the edges of the input graph G may have associated non-negative real *lengths*. The cost of an edge is usually distinct from, but related to its length. Often it is assumed that the vertices of the input graph G are embedded in some metric space. In this case, the *length* of an edge $\{u, v\}$, denoted $|uv|$ is equal to the distance between u and v in that space. If the vertices of G are embedded in a Euclidean space, then G is called a *Euclidean graph*. A special case of a Euclidean graph is a *unit disk graph*. G is a unit disk graph if its vertices are embedded in the Euclidean plane and for any pair of vertices u and v, $\{u, v\}$ is an edge of G iff $|uv| \leq 1$. As mentioned in [8], unit disk graphs are usually used to model an ad-hoc network where all the network nodes are placed in an unobstructed plane and have equal (normalized) transmission power and isotropic antennas, that is, antennas transmitting with identical power in every direction of the plane. The cost of an edge $\{u, v\}$, $c(u, v)$, is typically used to denote the amount of energy that one endpoint of the edge has to expend in order to communicate with the other endpoint. For a Euclidean graph, it is reasonable to assume that $c(u, v) = w(u, v)^\alpha$ for some $\alpha \geq 2$. If G_P, the output of a topology control protocol satisfies the spanner property with respect to edge costs, then it is called an *energy spanner*. It G_P satisfies the spanner property with respect to edge lengths, then G_P is called a *distance spanner*.

Related Work. Various topology control protocols have been proposed, each guaranteeing some subset of the above mentioned properties. Here we mention the two protocols that seem to provide strongest guarantees. Wang and Li [7] have proposed a local protocol for construction of symmetric, bounded degree, planar spanners for networks modeled by unit disk graphs. We will call this the WL protocol. Wattenhofer and Zollinger [8] have proposed a much simpler protocol called XTC that constructs symmetric, bounded degree, planar networks for networks modeled by unit disk graphs. In addition, XTC returns a symmetric, connected network even for input networks that have arbitrary edge lengths. In favour of the WL protocol is the fact that this protocol is guaranteed to return a spanner, whereas XTC provides no such guarantees. [8] does present experiments to suggest that the output of XTC may be a good spanner in the "average case." In favour of the XTC protocol is its extreme simplicity and the fact that the output graph is connected even when the input graph is not a unit disk graph. This implies that it may be appropriate to use XTC even when the terrain on which the nodes are distributed is not the 2-dimensional plane and even when there are obstacles in the terrain.

Our Results In this paper, we start by pointing out that existing protocols for the topology control problem, including the WL protocol and XTC, lack *robustness* and are extremely sensitive to faulty information from neighbors. For example, as we show in Section 2, the network constructed by XTC may end up becoming disconnected even when two nodes receive faulty distance information from one neighbor each. In a key step in the WL protocol, XTC, and other protocols such as the cone based protocol described in [4], each node u computes a total ordering $\prec_u$ on its neighborhood $N(u)$. In the WL protocol $\prec_u$ is based on degrees of vertices in $N(u)$, in XTC $\prec_u$ is based on the "quality" of the link between u and each vertex in $N(u)$, and in the cone based protocol $\prec_u$ is based on angles. In each case, correct information from neighbors is critical to the correctness of the neighborhood ordering and therefore critical to the correctness of the protocol itself.

Given that faulty information is almost par for the course in ad-hoc wireless networks, we seek robust topology control protocols whose output is guaranteed to satisfy the properties mentioned above, even in the presence of faulty information about neighbors. Faulty information from neighbors could be due to interference, due to feeble power supply at the sender node, due to the receiver having an incorrect estimate of the transmission range at the sender, etc. Even if information from neighbors is not faulty, it could simply be out-of-date because nodes may be mobile. In this paper we present a metric for the notion of robustness for topology control protocols that compute and use neighborhood orderings. We define (informally, for now) a k-*robust* protocol as one that can withstand a total of up to $k-1$ adjacent swaps performed on all the neighborhood ordering. We point out that XTC is not even 2-robust. We then present a simple modification to XTC that can turn it into a r-robust protocol for any integer $r > 0$. The price we pay for the increase in robustness is in terms of a decrease in the sparsity of the network. However, we bound this decrease. More specifically, in transforming XTC from a 1-robust protocol (which it is, trivially) to a r-robust protocol, for any integer $r > 1$, we increase the maximum vertex degree of the output graph by a factor of $O(\sqrt{r})$. Even with these modifications, XTC continues to be extremely simple and practical. An extremely pleasant side-effect of our design is that the output network is both $\Omega(\sqrt{r})$-edge connected and $\Omega(\sqrt{r})$-vertex connected. In other words, ensuring robustness of the protocol seems to provide fault-tolerance of the output for free.

2 XTC is not Robust

We start this section by reproducing the XTC protocol from [8].

1. Establish order $\prec_u$ over u's neighbors in G
2. Broadcast $\prec_u$ to each neighbor in G; receive orders from all neighbors
3. Select topology control neighbors:
4. $\quad N_u := \{\ \}$; $\widetilde{N}_u := \{\ \}$
5. $\quad$ **while** ($\prec_u$ contains unprocessed neighbors){
6. $\qquad v :=$ least unprocessed neighbor in $\prec_u$
7. $\qquad$ **if**($\exists w \in N_u \cup \widetilde{N}_u : w \prec_v u$)
8. $\qquad\quad \widetilde{N}_u := \widetilde{N}_u \cup \{v\}$
9. $\qquad$ **else**
10. $\qquad\quad N_u := N_u \cup \{v\}$
11. $\quad$ }

As mentioned in [8], the protocol consists of three main steps: (i) neighbor ordering (Line 1), (ii) neighbor order exchange (Line 2), and (iii) edge selection (Lines 3-11). In the edge selection step a vertex u decides to drop v from its set of neighbors if there is a vertex w that u and v both agree is mutually better. More precisely, u drops v from its neighborhood if there exists w such that $w \prec_u v$ and $w \prec_v u$. In the protocol, the variable N_u is the set of neighbors that u has chosen to retain and the variable $\widetilde{N}_u$ is the set of neighbors that u has chosen to drop. Let $E_{XTC} = \{(u, v) \mid v \in N_u\}$ and $G_{XTC} = (V, E_{XTC})$. Also, let $\prec = \{\prec_u \mid u \in V(G)\}$ denote the collection of neighborhood orderings. Note that the protocol leaves $\prec$ unspecified. Thus G_{XTC} is a function, not only of the input network G, but also of the neighborhood orderings $\prec$. This dependency will be important later and to emphasize this we use the notation $G_{XTC}(\prec)$ to denote the network constructed by the above protocol. In general, for a topology control protocol P, we use the notation $G_P(\prec)$ to denote the output of P. It is easily verified that $u \in N_v$ iff $v \in N_u$ and hence $G_{XTC}(\prec)$ can be thought of as undirected graph.

As mentioned in the introduction, XTC is extremely sensitive to small perturbations in the neighborhood orderings. In [8], it is shown that if G is a Euclidean graph and $\prec = \{\prec_u \mid u \in V(G)\}$, where $\prec_u$ is defined as

$$v \prec_u w \Leftrightarrow (|uv|, \min\{id_u, id_v\}, \max\{id_u, id_v\}) < (|uw|, \min\{id_u, id_w\}, \max\{id_u, id_w\}),$$

then $G_{XTC}(\prec)$ is symmetric and connected. We will call the above neighborhood ordering, a *distance-based* ordering. Note that in the distance-based ordering, ids are only used to break ties. We now present a simple example of a 4-vertex unit disk graph that illustrates the lack of robustness of XTC. We start with the neighborhood orderings $\prec$ as defined above, by Euclidean distance. We then make one swap each in the neighborhood orderings of two vertices to obtain new neighborhood orderings $\tilde{\prec}$. We point out that $G_{XTC}(\tilde{\prec})$ is not connected. Consider the unit disk graph shown in Figure 1. For the sake of being concrete, let the lengths of the edges be $|ab| = |dc| = \sqrt{3}/2$, $|ad| = |bc| = 1/2$, and $|ac| = |bd| = 1$. Then

$$\begin{array}{ccccc} d & \prec_a & b & \prec_a & c \\ c & \prec_b & a & \prec_b & d \\ b & \prec_c & d & \prec_c & a \\ a & \prec_d & c & \prec_d & b \end{array}$$

Now suppose that $\tilde{\prec}_a = \prec_a$, $\tilde{\prec}_d = \prec_d$, but

$$\begin{array}{ccccc} c & \tilde{\prec}_b & d & \tilde{\prec}_b & a \\ b & \tilde{\prec}_c & a & \tilde{\prec}_c & d \end{array}$$

Note that $\tilde{\prec}_b$ and $\tilde{\prec}_c$ are obtained by swapping one pair of elements each in $\prec_b$ and $\prec_c$. If XTC is run on the unit disk graph shown below with $\tilde{\prec} = \{\tilde{\prec}_a, \tilde{\prec}_b, \tilde{\prec}_c, \tilde{\prec}_d\}$, then $G_{XTC}(\tilde{\prec})$ contains just the two edges $\{a, d\}$ and $\{b, c\}$ and is therefore disconnected. Thus a total of two adjacent swaps were sufficient to break connectivity. Later in the paper we modify XTC in a simple manner into an r-robust protocol, one that can tolerate a total of up to $r - 1$ adjacent swaps on its neighborhood orderings.

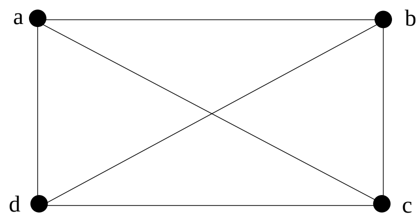

Figure 1: A unit disk graph for showing the sensitivity of XTC to small perturbations.

3 Characterizing Good Neighborhood Orderings

XTC's correctness and performance critically depends on $\prec$. Specifically, if $\prec$ is appropriately defined then the following two properties hold:

(i) For every triangle abc, $\prec_a$, $\prec_b$, and $\prec_c$ help vertices a, b, and c negotiate the dropping of one of the edges $\{a, b\}$, $\{b, c\}$, and $\{c, a\}$.

(ii) For every cut $(S, \overline{S})$ of G, $\prec$ prevents the dropping of some edge that crosses the cut $(S, \overline{S})$.

Property (i) implies that $G_{XTC}(\prec)$ is triangle-free, while (ii) implies that $G_{XTC}(\prec)$ is connected. Various properties of $G_{XTC}(\prec)$ proved separately in [8] immediately follow. Here we prove a general characterization of neighborhood orderings $\prec$ that guarantee properties (i) and (ii). It will be clear that the "distance-ordering" used in [8] satisfies this characterization. But more importantly, there are many other natural neighborhood orderings that also satisfy our characterization. For example, neighborhood orderings by increasing ids or by increasing angle also satisfy our characterization and therefore guarantee properties (i) and (ii).

The collection of neighborhood orderings $\prec$ induces a binary relation $\rightsquigarrow$ on the set of edges of G. For any two edges $e, e' \in E(G)$, $e \rightsquigarrow e'$ if e and e' share a common endpoint and if $e = \{u, v\}$ and $e' = \{u, w\}$, then $v \prec_u w$. Using this binary relation $\rightsquigarrow$ we can define a new graph $L(G, \prec)$ whose vertex set is the set of edges of G and whose set of (directed) edges is $\{(e, e') \mid e, e' \in E(G), e \rightsquigarrow e'\}$. We call $\prec$ *acyclic* if $L(G, \prec)$ is an acyclic graph. Note that is $L(G, \prec)$ is acyclic, then so is any subgraph of $L(G, \prec)$. Also note that any acyclic graph is guaranteed to contain at least one vertex with in-degree (respectively, out-degree) 0 and we call such a vertex, a *minimal* (respectively, *maximal*) vertex.

Theorem 1 *Let G be an arbitrary connected graph and $\prec$ be a collection of neighborhood orderings of G. $G_{XTC}(\prec)$ is triangle-free and connected if $\prec$ is acyclic.*

Proof: To show that $G_{XTC}(\prec)$ is triangle-free, we consider an arbitrary triangle abc in G. Since $L(G, \prec)$ is acyclic there is a triangle edge, say $\{a, b\}$, such that $\{b, c\} \rightsquigarrow \{a, b\}$ and $\{c, a\} \rightsquigarrow \{a, b\}$. This implies that $c \prec_b a$ and $c \prec_a b$. As a result XTC will drop edge $\{a, b\}$ and therefore the triangle abc is not part of $G_{XTC}(\prec)$. Since the choice of abc was arbitrary, $G_{XTC}(\prec)$ is triangle-free.

To show that $G_{XTC}(\prec)$ is connected, we consider a cut $(S, \overline{S})$ of G. Let $L_S(G)$ be the subgraph of $L(G, \prec)$ induced by the edges of G crossing the cut. Since $L(G, \prec)$ is acyclic, so is $L_S(G)$. Let e be a minimal vertex of $L_S(G)$. We now show that e is retained in $G_{XTC}(\prec)$. Let $e = \{u, v\}$ and suppose that e is not retained in $G_{XTC}(\prec)$. Then there is a vertex $w \in V(G)$ that is a common neighbor of u and v such that $w \prec_u v$ and $w \prec_v u$. Since $\{u, v\}$ crosses the cut $(S, \overline{S})$, at least one of $e_u = \{u, w\}$ or $e_v = \{v, w\}$ also crosses the cut. Without loss of generality suppose that e_u crosses $(S, \overline{S})$. Then, by the definition of $\rightsquigarrow$, $e_u \rightsquigarrow e$ and therefore e is not minimal in $L(G, \prec)$. This contradicts our choice of e as a minimal vertex.

Thus we have shown that for every cut $(S, \overline{S})$ of G, there an edge in $G_{XTC}(\prec)$ crossing the cut. This shows that $G_{XTC}(\prec)$ is connected. $\square$

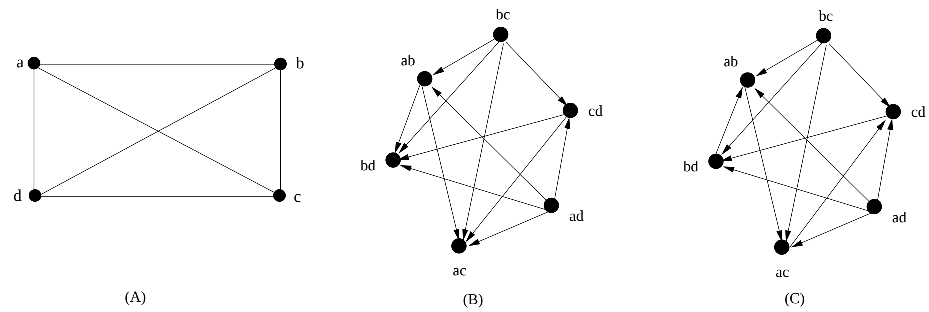

Figure 2: On the left is the unit disk graph from Figure 1. In the middle is $L(G, \prec)$, where $\prec$ is the distance-based ordering. It is easily verified that this is acyclic. Vertices ad and bc are minimal vertices in $L(G, \prec)$. On the right is $L(G, \widetilde{\prec})$, where $\widetilde{\prec}$ is obtained from $\prec$ by swapping a and d in $\prec_b$ and $\prec_c$. Note the cycle (ab, ac, cd, bd, ab) in $L(G, \widetilde{\prec})$.

It is easy to see that the distance-based ordering is acyclic. Let G be a Euclidean graph and let $e = \{u, v\}$ be the edge in G such that the triple $(|uv|, \min\{id_u, id_v\}, \max\{id_u, id_v\})$ is first in the increasing lexicographic ordering of all such triples. From the definition of the distance-based ordering, it follows that e is minimal in $L(G, \prec)$. If we assume that $L(G - e, \prec)$ is acyclic, then by induction it follows that so is $L(G, \prec)$. Similar arguments show that the following alternate orderings are also acyclic.

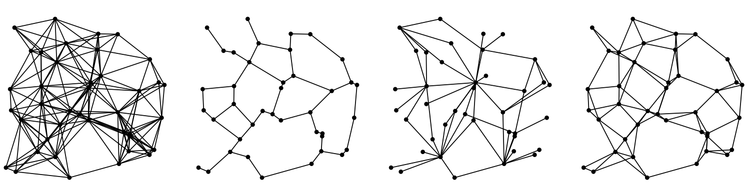

Figure 3: The graph on the left is a unit disk graph obtained by dropping 40 points uniformly at random on a 3×3 grid. It contains 197 edges. The second graph from the left is the output of XTC using a distance-based ordering and it contains 47 edges. The third graph from the left is the output of XTC using an id-based ordering and it contains 55 edges. The rightmost graph is the output of k-XTC using a distance-based ordering, for $k = 2$. It contains 88 edges.

1. The *id-based ordering* $\prec^{id}$. Let v and w be two neighbors of u. Then $v \prec_u^{id} w$ iff $id_v < id_w$. As before, $\prec^{id} = \{\prec_u^{id} \mid u \in V(G)\}$.
2. The *angle-based ordering* $\prec^a$. For any pair of vertices u and v in G, let $\alpha(u, v)$ denote the angle made by the line segment uv with the horizontal ray with origin u towards $+\infty$. For two neighbors v and w of u, $v \prec_u^a w$ iff

$$(\alpha(u, v), \min\{id_u, id_v\}, \max\{id_u, id_v\}) < (\alpha(u, w), \min\{id_u, id_w\}, \max\{id_u, id_w\}).$$

Of course, the id-based ordering is only well-defined when all vertices have (not necessarily distinct) ids and the angle-based ordering is only well defined when the vertices of G are embedded in Euclidean space and the vertices have ids because of the need to break ties when angle comparison does not distinguish neighbors.

The implication of the above characterization theorem is that XTC could have as well been run with the id-based ordering or the angle-based ordering instead of the distance-based ordering and the output graph would still have the properties: (i) symmetry, (ii) connectivity, and (iii) being triangle-free.

It should be noted that ignoring distances completely and using the id-based ordering or angle-based ordering is not, in general, a good idea. Though symmetry and connectivity are preserved, the output graph may have other undesirable features. Some of these are apparent in the graph in Figure 3 (third from left) that is constructed by XTC using the id-based ordering. For example, the degrees of certain vertices are quite high and moreover these high degree vertices have several pendant edges incident on them. These nodes are therefore prone to high congestion and the network is vulnerable to the failure of such nodes. While we are not suggesting the use of id-based ordering as an alternative to distance-based ordering, the result in Theorem 1 does suggest the possibility of using id-based ordering when distances to neighbors are similar (not necessarily the same). This may be another way to increase robustness of the protocol.

4 k-XTC: A Robust Version of XTC

In this section we propose a small modification to XTC that will result in a protocol that can tolerate arbitrary errors in a fixed number of pieces of information that are used by vertices to establish neighborhood orderings. The protocol, which we will call k-XTC is obtained from XTC by changing Line 7 to the following.

$$\text{if } (\exists W \subseteq N_u \cup \widetilde{N}_u \text{: } |W| = k \text{ and } \forall w \in W : w \prec_v u).$$

This modification simply means that the decision for u to drop v from its neighborhood needs the support of not one, but k other vertices that both u and v agree are mutually better. Let $G_{kXTC}(\prec)$ denote the output of k-XTC. Note that XTC is simply a special case of k-XTC with $k = 1$. A simple but important observation about the output of k-XTC is the following.

Proposition 2 *For any j, $1 \leq j < k$, $G_{jXTC}(\prec)$ is a subgraph of $G_{kXTC}(\prec)$.*

The rightmost graph in Figure 3 shows the output of k-XTC for $k = 2$. This graph has the same rough "shape" as the output of XTC (the graph that is second from left) but is denser and non-planar. As we will show later, this graph is k-edge connected as well as k-vertex connected. Therefore, every vertex in this graph has degree at least k.

We now quantify the notion of robustness as follows.

Definition: Let π and π' be two permutations of a finite, non-empty set S. We denote the fewest number of adjacent swaps needed to transform π to π' by $dist(\pi, \pi')$.

Definition: Let $\prec = \{\prec_u |\ u \in V(G)\}$ and $\widetilde{\prec} = \{\widetilde{\prec}_u |\ u \in V(G)\}$ be two collections of neighborhood orderings. Then we use $dist(\prec, \widetilde{\prec})$ to denote $\sum_u dist(\prec_u, \widetilde{\prec}_u)$.

Definition: A topology-control protocol P is said to be *r-robust for $\prec$* if $G_P(\widetilde{\prec})$ is connected for any collection of neighborhood orderings $\widetilde{\prec}$, where $dist(\prec, \widetilde{\prec}) < r$.

In other words, if P is r-robust for $\prec$, then P returns a connected subgraph when executed with $\prec$, but furthermore it returns a connected subgraph even when executed with a collection of neighborhood orderings that is obtained from $\prec$ using at most $r - 1$ adjacent swaps. Thus P can tolerate a total of $r - 1$ or fewer adjacent swaps made to its neighborhood orderings in the sense that it still returns a connected subgraph as output. Measuring the "distance" between orderings by the number of adjacent swaps provides a clean abstraction for quantifying a variety of errors that might cause vertices to believe a "false" ordering on neighbors. For example, if a vertex u underestimates the distance to a neighbor v then v might appear earlier than it should in $\prec_u$. If the (incorrectly) estimated distance to v is much smaller than the actual distance, then v's place in $\prec_u$ may be many adjacent swaps away from its correct place in $\prec_u$. We now prove the main result of this paper. Note that the result is proved for any collection of acyclic neighborhood orderings and not just for the distance-based ordering. Showing that k-XTC is k-robust is not hard, but showing a quadratic robustness needs the more intricate argument presented below.

Theorem 3 *k-XTC is $\frac{k(k+1)}{2}$-robust for any collection $\prec$ of acyclic neighborhood orderings.*

Proof: Let G be the input graph to k-XTC. Let $\prec$ be an arbitrary collection of *acyclic* neighborhood orderings and let $\widetilde{\prec}$ be an arbitrary collection of neighborhood orderings. From Theorem 1 and Proposition 2, we know that $G_{kXTC}(\prec)$ is connected. We will show that if $G_{kXTC}(\widetilde{\prec})$ is disconnected then $dist(\prec, \widetilde{\prec}) \geq k(k+1)/2$. This will imply that k-XTC is $\frac{k(k+1)}{2}$-robust.

We start by supposing that $G_{kXTC}(\widetilde{\prec})$ is disconnected and assuming for notational convenience, that $\widetilde{H} = G_{kXTC}(\widetilde{\prec})$. Since $\widetilde{H}$ is disconnected there is a cut $C = (S, \overline{S})$ such that there is no edge of $\widetilde{H}$ crossing $(S, \overline{S})$. On the other hand there is at least one edge in G crossing C. Let $E(C)$ be subset of edges in G crossing C. Since $L(G, \prec)$ is acyclic, the subgraph of $L(G, \prec)$ induced by edges in $E(C)$ is also acyclic. In the rest of the proof we use $L(C)$ to denote the subgraph of $L(G, \prec)$ induced by $E(C)$.

Our proof is constructive and what we now describe is the first iteration of the construction procedure. Let $e = \{u, v\}$ be a minimal edge in $L(C)$. Without loss of generality, suppose that $u \in S$ and $v \in \overline{S}$. The edge e does not appear in $\widetilde{H}$ and this can only happen because there is a set W of k vertices such that for all $w \in W$, w is a common neighbor of u and v, $w \mathrel{\widetilde{\prec}_u} v$, and $w \mathrel{\widetilde{\prec}_v} u$. Let (W_u, W_v) be a partition of W such that $W_u \subseteq S$ and $W_v \subseteq \overline{S}$. Let $k_u = |W_u|$ and $k_v = |W_v|$. Note that $k_u + k_v = k$. Also note that for each $w \in W_v$, edge $\{u, w\}$ crosses C and similarly for each $w \in W_u$, edge $\{v, w\}$ crosses C. Also

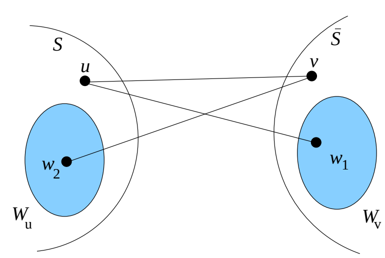

Figure 4: The edges $\{u, w_1\}$ and $\{v, w_2\}$ cross the cut $(S, \overline{S})$. Furthermore, $v \prec_u w_1$ and $w_1 \mathrel{\widetilde{\prec}_u} v$. Also, $u \prec_v w_2$ and $w_2 \mathrel{\widetilde{\prec}} u$.

note that since $\{u, v\}$ is a minimal edge in $L(C)$, $v \prec_u w$ for all $w \in W_v$ and $u \prec_v w$ for all $w \in W_u$. Thus, we have:

(i) for all $w \in W_v$, $w \mathrel{\widetilde{\prec}_u} v$ and $v \prec_u w$ and

(ii) for all $w \in W_u$, $w \mathrel{\widetilde{\prec}_v} u$ and $u \prec_v w$.

See Figure 4 for an example. Item (i) implies that $dist(\prec_u, \widetilde{\prec}_u) \geq k_v$ and item (ii) implies that $dist(\prec_v, \widetilde{\prec}_v) \geq k_u$. These inequalities together imply that $dist(\prec, \widetilde{\prec}) \geq k$.

Remark: Actually, something stronger can be claimed. Even if we wanted to transform $\prec_u$ into an ordering $\prec'_u$ such that $w \prec'_u v$ for all $w \in W_v$, but $\prec_u$ and $\prec'_u$ match in the pairwise ordering of all others pairs of elements, it would take at least k_v adjacent swaps. In other words, $\prec'_u$ is along the way between $\prec_u$ and $\widetilde{\prec}_u$ and just getting to $\prec'_u$ from $\prec_u$ takes at least k_v adjacent swaps. Getting to $\widetilde{\prec}_u$ from $\prec'_u$ may take additional adjacent swaps and we account for these separately in future iterations of the construction procedure. Similar remarks can be made about the "distance" between $\prec_v$ and $\widetilde{\prec}_v$.

The choice of edge $e = \{u, v\}$ described above, ends the first iteration of our construction procedure. Let $B_1 = \{e\}$ and let $V_1 = \{u, v\}$. The set V_1 represents the endpoints of the edge in B_1. To state our induction hypothesis we need additional notation. For any set X of vertices, let $dist_X(\prec_u, \tilde{\prec}_u)$ be the minimum number of adjacent swaps we need to make on $\prec_u$ so that every element $v \in X \cap N(u)$ is in the same relative position in $\prec_u$ as in $\tilde{\prec}_u$. More precisely, $dist_X(\prec_u, \tilde{\prec}_u) = \min_{\prec'_u} dist(\prec_u, \prec'_u)$, where the min operation is over all $\prec'_u$ such that for any $v \in X \cap N(u)$ and for any $w \in N(u)$, $v \prec_u w \Leftrightarrow v \prec'_u w$. Here is a small example to illustrate this definition.

Example. Define the permutations $\pi = (54321)$ and $\pi' = (12345)$. It is easy to see that $dist(\pi, \pi') = 10$. Now let $X = \{1, 4\}$. What is $dist_X(\pi, \pi')$? It is again easy to verify that $dist_X(\pi, \pi') = dist(\pi, \pi'') = 9$, where $\pi'' = (13245)$. This is because $dist_X(\pi, \pi')$ is the number of adjacent swaps needed to transorm π into a permutation in which 1 appears before all other elements and 4 appears after all other elements except 5. Thus the positions of elements 1, 4, and 5 are fixed.

For any collection $\prec$ of neighborhood orderings, let $dist_X(\prec, \tilde{\prec}) = \sum_{u \in X} dist_X(\prec_u, \tilde{\prec}_u)$. We also need the following two elementary facts about transforming one permutation into another via adjacent swaps.

Fact 1. For any $X \subseteq Y \subseteq N(u)$,

$$dist_X(\prec_u, \tilde{\prec}_u) \leq dist_Y(\prec_u, \tilde{\prec}_u).$$

Fact 2. Let $X \subseteq Y \subseteq N(u)$ and $x \in Y - X$. Suppose there is a set $W \subseteq N(u)$ such that for all $w \in W$, $x \prec_u w$ and $w \tilde{\prec}_u x$ then

$$dist_X(\prec_u, \tilde{\prec}_u) + |W| \leq dist_Y(\prec_u, \tilde{\prec}_u).$$

Our induction hypothesis is the following.

Induction hypothesis: For any $i \geq 1$, after i iterations of this procedure, we have a set B_i of i edges from $E(C)$ such that there are no edges from $E(C) - B_i$ into B_i, though there may be edges from B_i into $E(C) - B_i$. Let V_i be the set of endpoints of edges in B_i. Then $dist_{V_i}(\prec, \tilde{\prec}) \geq k + (k-1) + \cdots + (k-i+1)$.

We have shown that at the end of the first iteration of the construction procedure, $|B_1| = 1$, there are no edges from $E(C) - B_1$ into B_1, and $dist_{V_1}(\prec, \tilde{\prec}) \geq k$. This is the base case of our proof.

We now make the following claim about the $(i+1)$st iteration of our construction procedure. We will prove this claim later; for now we will assume that it holds and complete the proof of the induction step.

Claim: In the $(i+1)$st iteration it is possible to pick an edge $e' \in E(C) - B_i$ such that (i) e' has at least one endpoint not in V_i, and (ii) in-degree of e' in $L(C)$ is at most i.

Assuming this claim, we proceed in a manner that is similar to the argument for the first iteration. Let $e' = \{u', v'\}$, $u' \in S$, $v' \in \overline{S}$, and without loss of generality, $v' \notin V_i$. The fact that e' is not in $\tilde{H}$ implies that there is a set W of k vertices such that for all $w \in W$, w is a common neighbor of u and v, $w \tilde{\prec}_{u'} v'$ and $w \tilde{\prec}_{v'} u'$. Using the fact (derived from the above claim) that the in-degree of e' in $L(C)$ is at most i, we conclude, using an argument similar to the one for the first iteration, that there exist subsets $W_{u'} \subseteq W \cap S$ and $W_{v'} \subseteq W \cap \overline{S}$, such that $|W_{u'}| + |W_{v'}| = (k - i)$ and

(i) for all $w \in W_{v'}$, $w \tilde{\prec}_{u'} v'$ and $v' \prec_{u'} w$ and

(ii) for all $w \in W_{u'}$, $w \tilde{\prec}_{v'} u'$ and $u' \prec_{v'} w'$.

Let $k_{u'} = |W_{u'}|$ and $k_{v'} = |W_{v'}|$, $B_{i+1} = B_i \cup \{e'\}$, and V_{i+1} be the endpoints of vertices in B_{i+1}. Item (i) along with Fact 2 implies that $dist_{V_{i+1}}(\prec_{u'}, \tilde{\prec}_{u'}) \geq dist_{V_i}(\prec_{u'}, \tilde{\prec}_{u'}) + k_{v'}$. Item (ii) implies that $dist_{V_{i+1}}(\prec_{v'}, \tilde{\prec}_{v'}) \geq k_{u'}$. These inequalities together along with Fact 1 imply that $dist_{V_{i+1}}(\prec, \tilde{\prec}) \geq dist_{V_i}(\prec, \tilde{\prec}) + (k - i)$. This completes the induction step. If we repeat the induction step until $i = k$, then we have a set V_k of vertices such that $dist_{V_k}(\prec, \tilde{\prec}) \geq k(k+1)/1$. Since $V_k \subseteq V$, by Fact 1 we have that $dist(\prec, \tilde{\prec}) = dist_V(\prec, \tilde{\prec}) \geq dist_{V_k}(\prec, \tilde{\prec}) \geq k(k+1)/1$.

We now prove the above claim that guarantees the existence of e'.

Proof of Claim: Let T_i be the set of edges not in B_i, that have both endpoints in V_i. Consider the subgraph of $L(C)$ obtained by deleting $B_i \cup T_i$. Call this L_i. Since $L(C)$ is acyclic, L_i is also acyclic and let e' be a minimal vertex in L_i. If e' is not incident on any vertex in V_i, then e' is also minimal in $L(C)$ and we are done. So we assume that e' is incident on at least on vertex in V_i. Since, e' was picked from $L(C) - B_i - T_i$, e' cannot be incident on two vertices in V_i, because otherwise e' will be in T_i. Therefore, we are left with the case in which e' is incident on one vertex in V_i. Now let $e' = \{b, g\}$ and suppose that $b \in V_i$ and $g \notin V_i$. Figure 5 illustrates the situation for $i = 3$. Suppose that there are x edges in B_i

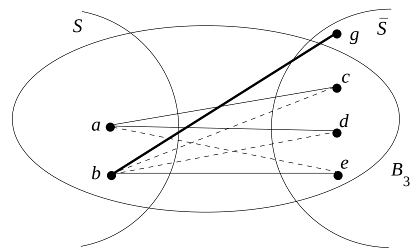

Figure 5: This figure illustrates the proof of the Claim in the proof of Theorem 3. Here $B_3 = \{\{a, c\}, \{a, d\}, \{b, e\}\}$, $T_3 = \{\{b, d\}, \{b, c\}, \{a, e\}\}$, and $e' = \{b, g\}$. The set $V_3 = \{a, b, c, d, e\}$.

incident on b and y edges in T_i incident on b. In Figure 5, $x = 1$ and $y = 2$. The in-degree of e' in $L(C)$ is therefore bounded above by $x + y$ and since $x \leq |B_i| = i$, we get the upper bound $i + y$. Now note that for every edge $\{b, b'\}$ in T_i, there is an edge in B_i incident on b' that does not share any endpoints with edge $\{a, b\}$. In other words, for every edge e'' of T_i such that $e'' \rightsquigarrow e'$ there is a unique edge f in B_i such that $f \not\rightsquigarrow e'$. This gives the upper bound of i on the in-degree of e'. $\square$

The above result was stated in terms of tolerance to spurious swaps performed by a possibly malicious adversary. This result can also be recast (though, less cleanly) in terms of "reading errors." The $k(k+1)/2 - 1$ swaps can be tolerated by the system, could be the result of a vertex u reading information from $k - 1$ neighbors erroneously such that their relative distances are completely reversed as compared to their original distances. Alternately, k different vertices could each read information about a single neighbor incorrectly such that the new (erroneous) neighborhood ordering at each of these k vertices is $(k - 1)/2$ swaps away from the correct neighborhood ordering.

An extremely pleasant side-effect of our design of k-XTC is the fault-tolerance of G_{kXTC}. We prove in the following two theorems that if G is k-edge connected (respectively, k-vertex connected) then G_{kXTC} is also k-edge connected (respectively, k-vertex connected). Localized protocols for constructing such fault-tolerant spanning subgraphs appear in [1, 5], but k-XTC is far simpler than these. Furthermore, k-XTC provides robustness, bounded degree in case the input is a unit disk graph, and also preserves k-connectivity for arbitrarily input graphs with arbitrary edge lengths. Also note that the following two theorems are proved for any acyclic collection of neighborhood orderings, not just for distance-based orderings.

Theorem 4 *For any collection of acyclic neighborhood orderings $\prec$, $G_{kXTC}(\prec)$ is k-edge connected provided G is k-edge connected.*

Proof: Suppose G is k-edge connected, but G_{kXTC} is not. For any cut $C = (S, \overline{S})$ of V, let $E(C)$ denote the edges in G crossing the cut and similarly, let $E_{kXTC}(C)$ denote the edges of G_{kXTC} crossing the cut C. Since G is k-edge connected, but G_{kXTC} is not, there there is a cut $C = (S, \overline{S})$ of V such that $|E(C)| \geq k$ and $|E_{kXTC}(C)| < k$. Let L be the subgraph of $L(G, \prec)$ induced by $E(C) - E_{kXTC}(C)$. Note that L is non-empty and since $L(G, \prec)$ is acyclic, so is L. Let $e = \{u, v\}$ be a minimal vertex in L. Since $\{u, v\} \in E - E_{kXTC}$, there exists a vertex set W, $|W| = k$, such that for all $w \in W$, w is a common neighbor of u and v and

$$w \prec_u v \text{ and } w \prec_v u. \tag{1}$$

Let $W_1 = W \cap S$ and $W_2 = W \cap \overline{S}$. Then,

$$\chi = \{\{u, x\} \mid x \in W_2\} \cup \{\{y, v\} \mid y \in W_1\}$$

is a subset of $E(G)$ of edges that cross the cut $(S, \overline{S})$. Note that $|\chi| = k$ and therefore not all edges in χ can belong to $E_{kXTC}(C)$. Let $\{a, b\} \in \chi - E_{kXTC}(C)$. Thus $\{a, b\} \in E(C) - E_{kXTC}(C)$ and is therefore a vertex in L. Note that $\{a, b\}$ is either incident on u or incident on v. Without loss of generality, assume that $a = u$. Then, from (1) it follows that $b \prec_u v$. This means that $\{a, b\} \rightsquigarrow \{u, v\}$, contradicting that fact that $e = \{u, v\}$ is minimal in L. $\square$

Theorem 5 *For any collection of acyclic neighborhood orderings $\prec$, $G_{kXTC}(\prec)$ is k-vertex connected provided G is k-vertex connected.*

Proof: Suppose that G is k-vertex connected and G_{kXTC} is not. Since G_{kXTC} is not k-vertex connected, there exists $V' \subseteq V$ such that $|V'| = k-1$ and $G'_{kXTC} = G_{kXTC} - V'$ is disconnected. Since G is k-vertex connected, $G' = G - V'$ is connected. Since G'_{kXTC} is disconnected, there exists cut $C = (S, \overline{S})$ of $V - V'$ such that no edges in G'_{kXTC} cross cut C. However, since G' is connected, there exists a non-empty set of edges E_C in G' that cross cut C. Let L be the subgraph of $L(G, \prec)$ induced by E_C. Let $e = \{u, v\}$ be a minimal vertex in E_C. Without loss of generality suppose that $u \in S$ and $v \in \overline{S}$. Since there are no edges in G'_{kXTC} that cross the cut C, e is not in G_{kXTC}. Hence, there exists $W \subseteq V$, $|W| = k$, such that for all $w \in W$, w is a common neighbor of both u and v, and $w \prec_u v$ and $w \prec_v u$. Since $|W| = k$ and $|V'| = k - 1$, there exists a vertex $w \in W - V'$. Therefore, w is a vertex in G' and in G'_{kXTC}. Without loss of generality assume that $w \in \overline{S}$. Therefore, edge $\{u, w\}$ crosses the cut C and belongs to E_C. Furthermore, since $w \prec_u v$, $\{u, w\} \rightsquigarrow \{u, v\}$ contradicting the fact that $\{u, v\}$ is minimal in L. $\square$

We use $\Delta(G)$ to denote the maximum degree of a vertex in G. We now show that the argument for the upper bound 6 [8] on $\Delta(G_{XTC})$ if G is a unit disk graph carries over cleanly to the k-XTC, giving an upper bound of $6k$ on $\Delta(G_{kXTC})$. Note that the argument if specific to distance-based orderings, not just any acyclic ordering. In fact, as mentioned before for the case of the id-based ordering, not all acyclic orderings will satisfy this upper bound result.

Theorem 6 *If G is unit disk graph and $\prec$ is the collection of distance-based neighborhood orderings, then $\Delta(G_{kXTC}(\prec)) \leq 6k$.*

Proof: To prove this theorem, we show that $k+1$ adjacent edges in G_{kXTC} cannot enclose an angle less than $\frac{\pi}{3}$. More precisely, assume that a vertex u has $k+1$ neighbors $v_0, v_1, \ldots, v_k$ in G_{kXTC}, listed in counterclockwise order starting at some arbitrary neighbor v_0. Further assume that $\angle v_0 u v_k < \frac{\pi}{3}$. Figure 6 illustrates the situation.

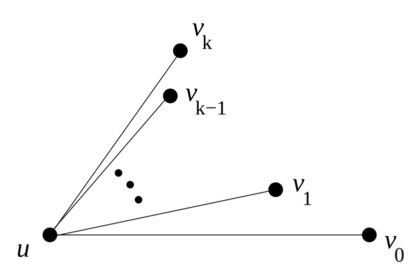

Figure 6: The neighbors $v_0, v_1, \ldots, v_k$ of u. For the proof we suppose that $\angle v_0 u v_k < \pi/3$.

Suppose that among the neighbors $v_0, v_1, \ldots, v_k$, the neighbor v_i for some i, $0 \le i \le k$, is considered last by k-XTC. Since v_i is considered last we have $v_j \prec_u v_i$ for all $j \ne i$, $0 \le j \le k$. Since $\prec$ is the distance-based ordering, this implies that $|uv_j| \le |uv_i|$, for all $j \ne i$, $0 \le j \le k$.

Now consider a triangle uv_iv_j, $j \ne i$, $0 \le j \le k$. Since $|uv_j| \le |uv_i|$, uv_j is not the longest edge of the triangle. Also since $\angle v_j u v_i < \frac{\pi}{3}$, the line segment v_iv_j is strictly shorter than at least one of the other two line segments in the triangle, namely uv_i and uv_j. Combining this with the fact that $|uv_j| \le |uv_i|$, we have $|v_iv_j| < |uv_i|$, implying that $v_j \prec_{v_i} u$. Thus, we have $v_j \prec_u v_i$ and $v_j \prec_{v_i} u$ for all $j \ne i$, $0 \le j \le k$. This means that edge $\{u, v_i\}$ will not be included in G_{kXTC}, contradicting the fact that v_i is a neighbor of u in G_{kXTC}. □

5 Future Directions

The spanner properties of G_{kXTC} remain unexplored and there are several interesting questions one could ask. For example, as k increases G_{kXTC} becomes more dense and we expect it to become a better spanner for G. One could ask if given any $t \ge 1$ and a unit disk graph G, whether there is a $k = k(t)$ such that G_{kXTC} is a t-spanner for G. It is possible that the answer to this question is "no" and one could then ask the following "smaller" question. Are there a pair of constants (t, k) such that for any unit disk graph G, G_{kXTC} is a t-spanner. Again, it is possible that for any k, one could construct a unit disk graph G that would force G_{kXTC} to be an arbitrarily bad spanner. In this case, one could focus on random unit disk graphs (those obtained by distributing points uniformly at random in a bounded planar region) and investigate spanner properties of G_{kXTC} in this setting. This would also be an attempt at analytically proving the conclusion, experimentally derived in [8], that G_{XTC} is a good spanner for random unit disk graphs.

Acknowledgement. We thank Ted Herman for illuminating discussions on real world sensor and ad-hoc wireless networks.

References

[1] M. Bahramgiri, M. Hajiaghayi, and V. S. Mirrokni. Fault-tolerant and 3-dimensional distributed topology control algorithms in wireless multi-hop networks. In *Proceedings of the 11th IEEE International Conference on Computer Communications and Networks (IC3N)*, pages 392–398, 2002.

[2] M. Burkhart, P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Does topology control reduce interference? In *Proceedings of the 4th ACM International Symposium on Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, 2003.

[3] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric ad-hoc routing: Of theory and practice. In *Proceedings of the 22nd ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.

[4] L. Li, J. Halpern, P. Bahl, Y. Wang, and R. Wattenhofer. Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 264–273, 2001.

[5] N. Li and J.C. Hou. FLSS: A fault-tolerant topology control algorithm for wireless networks. In *Proceedings of MOBICOM*, 2004.

[6] R. Prakash. Unidirectional links prove costly in wireless ad-hoc networks. In *Proceedings of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communication (DIAL-M)*, 1999.

[7] Y. Wang and X. Y. Li. Localized construction of bounded degree planar spanner for wireless ad hoc networks. In *Proceedings of the 2003 Joint Workshop on Foundations of Mobile Computing*, pages 59–68, 2003.

[8] R. Wattenhofer and A. Zollinger. XTC: A practical topology control algorithm for ad-hoc networks. In *Proceedings of the 4th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN 04)*, 2004.