

XTC: A Practical Topology Control Algorithm for Ad-Hoc Networks

Roger Wattenhofer and Aaron Zollinger
{wattenhofer, zollinger}@inf.ethz.ch
Department of Computer Science, ETH Zurich
8092 Zurich, Switzerland

Abstract

The XTC ad-hoc network topology control algorithm introduced in this paper shows three main advantages over previously proposed algorithms. First, it is extremely simple and strictly local. Second, it does not assume the network graph to be a Unit Disk Graph; XTC proves correct also on general weighted network graphs. Third, the algorithm does not require availability of node position information. Instead, XTC operates with a general notion of order over the neighbors' link qualities. In the special case of the network graph being a Unit Disk Graph, the resulting topology proves to have bounded degree, to be a planar graph, and—on average-case graphs—to be a good spanner.

1 Introduction

An ad-hoc or sensor network consists of mobile nodes featuring, among other components, a processor, some memory, a wireless radio, and a power source; physical constraints often demand the power source to be feeble—a weak battery or a small solar cell.

Since energy is the limiting factor for lifetime and operability of an ad-hoc network, researchers have developed a variety of mechanisms and algorithms to conserve energy. These mechanisms and algorithms are often dubbed "topology control." The currently best proposals feature an impressive list of properties. However, these algorithms also often require many unrealistic assumptions: First, most algorithms assume that all the nodes know their exact positions, by means of a global positioning system (GPS) for example. Second, the algorithms assume that the world is flat and without buildings. In this paper we present the XTC¹ topology control algorithm that works i) without GPS and ii) even in a mountainous and obstructed environment. Surprisingly the XTC algorithm features all the relevant properties of topology control while being faster than any previous proposals. We believe that XTC is the currently most realistic topology control algorithm available.

The paper is organized as follows: In the remainder of the Introduction we discuss topology control in more detail and compare our contribution with related work. After providing preliminary definitions in Section II, we describe the XTC algorithm in Section III. For illustration, Section IV proves XTC's properties when

¹By the date of submission the authors have not yet been able to agree on the meaning of the letter "X" in "XTC". The candidate list comprises terms such as "exotic", "extreme", "exceptional", or "exemplary", but also "extravagant" or even "extraterrestrial". Consensus has however been achieved concerning the pronunciation of the algorithm name.

employed on Euclidean and Unit Disk Graphs. The algorithm's behavior on general weighted graphs is the subject of Section V. The subsequent section provides an evaluation of XTC on average-case random graphs. Section VII finally concludes the paper.

1.1 Topology Control

For two communicating ad-hoc nodes u and v , the energy consumption of their communication grows at least quadratically with their distance. Having one or more relay nodes between u and v therefore helps to save energy.

The primary target of a topology control algorithm is to abandon long-distance communication links and instead route a message over several small (energy-efficient) hops. For this purpose each node in the ad-hoc network chooses a "handful" of "close-by" neighbors "in all points of the compass" (we are going to fill in the details later). Clearly nodes cannot abandon links to "too many" far-away neighbors in order to prevent the ad-hoc network from being partitioned or the routing paths from becoming non-competitively long. In general there is a trade-off between network connectivity and sparseness.

Let the graph $G = (V, E)$ denote the ad-hoc network before running the topology control algorithm, with V being the set of ad-hoc nodes, and E representing the set of communication links. There is a link (u, v) in E if and only if the two nodes u and v can communicate directly. Running the topology control algorithm will yield a sparse subgraph $G_{tc} = (V, E_{tc})$ of G , where E_{tc} is the set of remaining links. The resulting topology G_{tc} should have the following properties:

Property 1 (Symmetry) The resulting topology G_{tc} should be symmetric, that is, node u is a neighbor of node v if and only if node v is a neighbor of node u .

Asymmetric communication graphs are unpractical, because many communication primitives become unacceptably complicated. A simple ACK message confirming the receipt of a message, for example, is already a nightmare in an asymmetric graph [16].

Property 2 (Connectivity) Two nodes u and v are connected if there is a path from u to v , potentially through multiple hops. If two nodes are connected in G , then they should still be connected in G_{tc} .

Although a minimum spanning tree (MST) is a sparse connected subgraph, it is often not considered a good topology, since close-by nodes in the original graph G might end up being far away in G_{tc} (G being a ring, for instance). Therefore Property 2 is usually strengthened:

Property 2+ (Spanner) For any two nodes u and v , if the optimal path between u and v in G has cost c , then the optimal path between u and v in G_{tc} has cost $f(c)$. If $f(c)$ is bounded from above by a linear function in c , the graph G_{tc} is called a spanner.

Researchers have studied a selection of cost metrics, the most popular being i) Euclidean distance and ii) various energy metrics. The cost of a link in model i) is the Euclidean distance of the link, in model ii) the distance is raised to a predefined power. In both models the cost of a path is commonly defined to be the sum of the costs of all links in the path.

As mentioned, the primary target of a topology control algorithm is to abandon long-distance neighbors, or more formally:

Property 3 (Sparseness) The remaining graph G_{tc} should be sparse, that is, the number of links should be in the order of the number of nodes, i.e. $|E_{tc}| = O(|V|)$.

This reflects that not too many close-by nodes must be chosen, which reduces interference and thus saves energy. Since there still might be some nodes with many neighbors (e.g. a star graph), also Property 3 features an improved version.

Property 3+ (Low Degree) Each node in the remaining graph G_{tc} has a small number of neighbors. In particular the maximum degree in the graph G_{tc} should be bounded from above by a constant.

Since connectivity and sparseness run against each other, topology control has been a thriving research area. In addition to the properties 1, 2, and 3, one can often find secondary targets. For instance, it is popular (and often for free) to ask the remaining graph to be planar in order to run a geometric (a.k.a. geographic, location-based, position-based) routing algorithm, such as GOAFR/GOAFR⁺ [11, 13], or GFG/GPSR [2, 9].

1.2 Contribution

This paper features three major contributions.

First, it is agreed upon that the subgraph G_{tc} should not be computed by a heavyweight *global* algorithm, but instead with a *distributed* algorithm. To account for mobility, it is often argued that an algorithm should not only be *distributed* but even *local*: Each node is allowed to exchange messages with its neighbors a few times and then must decide which links it wants to keep. Many naive topologies such as the MST can provably not be computed locally and are therefore not realistic. To the best of our knowledge we present the fastest algorithm so far, where each node only communicates with its neighbors twice.

Second, it is often assumed that the ad-hoc nodes are represented by points in a Euclidean plane. Two nodes are connected in the original graph G if and only if their Euclidean distance is at most 1 (a normalized transmission radius). The graph G is known as the Unit Disk Graph (UDG). We believe that a UDG is a good first step towards understanding ad-hoc algorithms. However, a UDG is not practical, since it is based on several assumptions:

- i) All nodes are homogeneous.
- ii) Antennas are perfect isotropic radiators, such that all transmission radii are equal.
- iii) Attenuation is uniform, that is, the Euclidean plane is flat and free of blocking objects such as walls. Radio propagation is as in vacuum.

Especially Assumption iii) is questionable in any realistic environment. We believe that ad-hoc network algorithms should work, or in other words be correct, also in a more hostile environment,

that goes beyond the UDG. There was an attempt to generalize the UDG [12]. In this paper we however even allow the original communication graph G not to comply with any of the above assumptions. Instead G is a general graph without any geometric assumptions. For instance—as opposed to many other algorithms—XTC also works correctly if nodes are located in three-dimensional space, as in a building. When studying efficiency (not correctness), analytically and by simulation, we will make geometric assumptions to prove stronger results that compare better with related work. To the best of our knowledge we present the first topology control algorithm that works for general graphs.

A third assumption that is commonly made is that the nodes have detailed information about their neighbors. It is often assumed that all the nodes know their exact coordinates in the plane, for instance by means of a global positioning system (GPS). The most notable exception is the cone-based topology control (CBTC) algorithm [23], where nodes conclude information about their neighbors merely based on their relative signal strength and the signal arrival angle. It is in this third respect that this paper’s main contribution lies. For the correctness of the algorithm it is sufficient that the network nodes order their neighbors according to a general concept of link quality.

We feel that the present paper represents a paradigm shift in topology control. Where recent research tried to improve existing algorithms by enhancing them with various new features (and thus rendering the algorithms more complicated), we actually present an algorithm that is simpler, faster, and works without unrealistic assumptions.

1.3 Related Work

Early work in topology control focused on the special case of randomly distributed nodes. Hou and Li [5] can be considered originators of topology control following up a paper by Takagi and Kleinrock [19].

A next wave of topology control algorithms advocated the employment of classic computational geometry structures, such as the Delaunay Triangulation [6], the Minimum Spanning Tree [17], or a generalized version of the Gabriel Graph [18]. Unfortunately neither the Delaunay Triangulation nor the Minimum Spanning Tree can be computed locally, and therefore these solutions lack practicality. The Gabriel Graph can be computed locally and is symmetric (Property 1) as well as an energy-spanner (Property 2+). It also features sparseness (Property 3, but not low degree as in Property 3+), an additional property that was not paid much attention to in [18]. Since the Gabriel Graph is similar to our algorithm not only with respect to all of these properties but also regarding its simplicity of construction—provided that node position information is available—, we compare our results against this graph in Section 6.

The cone-based topology control algorithm (CBTC) [23] was the first algorithm to achieve and prove several properties at the same time: Apart from being symmetric (Property 1), an energy spanner (Property 2+) and a sparse graph (Property 3), there is an optional distributed (but recursive and therefore not local) second phase that reduces the degree of the graph (Property 3+).² There have been a series of improvements collecting properties, some

²CBTC is often misinterpreted as being an implementation of the Yao Graph [24]. This is not quite accurate. The Yao Graph sectorizes the plane at each node, and tries to find a neighbor for the node in each sector. CBTC only needs neighbors as long as not all cones contain a neighbor. This passive definition gives more flexibility and as a consequence fewer neighbors on average.

based on CBTC [7, 14], some based on other geometric structures, such as a local version of the Delaunay triangulation [15]. The most recent proposal [22] features a local algorithm that computes a planar distance (and energy) spanner with constant node degree.

Sometimes cluster-based solutions are considered to be topology control. The most popular and promising approaches are based on (connected) dominating sets [1, 8, 10]. In particular [10] presents a solution to build a dominating set in a local fashion and constant time. This is currently the only constant-time algorithm with worst-case guarantees; surprisingly enough the algorithm works on general graphs, not only on Unit Disk Graphs. Although dominating sets also achieve symmetry (Property 1), connectivity (Property 2), and sparseness (Property 3), the resulting graphs often do not try to be distance or energy, but instead hop spanners. Whereas distance and energy spanners are related, hop spanners will produce orthogonal results; results so different in fact that we would consider cluster-based solutions to be topology control in a very broad sense only.

2 Preliminaries

This section provides formal definitions of basic concepts essential for the understanding of the paper.

In a *weighted graph* $G = (V, E)$ every edge $(u, v) \in E$ is attributed a weight ω_{uv} . When referring to a weighted graph we assume that the weights are symmetric: $\omega_{uv} = \omega_{vu}$.

The nodes of a *Euclidean graph* are assumed to be located in a Euclidean plane. Furthermore the edge weight of an edge (u, v) is defined to be $\omega_{uv} = |uv|$, where $|uv|$ is the Euclidean distance between the nodes u and v . Note that the definition of Euclidean graphs does not contain a statement on the existence of certain edges.

A *Unit Disk Graph* is a Euclidean graph containing an edge (u, v) if and only if $|uv| \leq 1$. Unit Disk Graphs are often employed to model an ad-hoc network where all network nodes are placed in an unobstructed plane and have equal (normalized) transmission power and isotropic antennas, that is antennas sending with identical power in every direction of the plane.

Strongly related to edge weights is the *cost of an edge*. The cost of an edge $c(u, v)$ can be considered to represent the effort an algorithm is required to expend in order to send a message over (u, v) . Common definitions of edge cost metrics include the hop or link metric $c(u, v) \equiv 1$, the Euclidean metric $c(u, v) = |uv|$, and the energy metric $c(u, v) = |uv|^e$ for an attenuation exponent $e \geq 2$.

A *path* $p(u, v)$ from a node u to a node v being a sequence of consecutively contingent edges starting at u and ending at v , the *cost of a path* $|p(u, v)|$ is accordingly defined to be the sum of the costs of all edges contained in the path.

3 XTC Algorithm

In this section we describe our topology control algorithm XTC. The algorithm consists of three main steps:

- I) Neighbor ordering,
- II) neighbor order exchange, and
- III) edge selection.

In the first step each network node u computes a total order \prec_u over all its neighbors in the network graph G . From an abstract point of view, this order is intended to reflect the quality of the

XTC Algorithm

- 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: $N_u := \{\}; \tilde{N}_u := \{\}$
- 5: while (\prec_u contains unprocessed neighbors) {
- 6: $v :=$ least unprocessed neighbor in \prec_u
- 7: if ($\exists w \in N_u \cup \tilde{N}_u : w \prec_v u$)
- 8: $\tilde{N}_u := \tilde{N}_u \cup \{v\}$
- 9: else
- 10: $N_u := N_u \cup \{v\}$
- 11: }

links to the neighbors. A node u will consider its neighbors in G (in the third step of the algorithm) according to \prec_u ordered with respect to decreasing link quality: The link to a neighbor appearing early in the order \prec_u is regarded as being of higher quality than the link to a neighbor placed later in \prec_u . A neighbor w appearing before v in order \prec_u is denoted as $w \prec_u v$. For illustration we assume in Section 4 that \prec_u corresponds to the order of the neighbors' Euclidean distances from u . It is however conceivable that the neighbor order reflects a much more general notion of link quality, such as signal attenuation or packet arrival rate.

In the second step the neighbor order information is exchanged among all neighbors. Typically a node u broadcasts its own neighbor order while receiving the orders established by all of its neighbors.

During the third step, which does not require any further communication, each node locally selects those neighboring nodes which will form its neighborhood in the resulting topology control graph, based on the previously exchanged neighbor order information. For this purpose a node u traverses \prec_u with decreasing link quality: "Good" neighbors are considered first, "worse" ones later. Informally speaking, a node u only builds a direct communication link to a neighbor v if u has no "better" neighbor w that can be reached more easily from v than u itself.

Although the XTC algorithm is executed at all nodes, the detailed description as shown in the above box assumes the point of view of a node u . Lines 1 and 2 correspond to Steps I) and II). Lines 3-11 define Step III) in more detail: First the two sets N_u and \tilde{N}_u are initialized to be empty. Now the neighbor ordering \prec_u established in Line 1, is traversed in increasing order. In Line 7 the neighbor order \prec_v of the currently considered neighbor v is examined: If any of u 's neighbors w already processed appears in v 's order before u ($w \prec_v u$) node v is included in \tilde{N}_u (Line 8); otherwise v is added to N_u (Line 10).

After completion of the algorithm, the set N_u contains u 's neighbors in the topology control graph G_{XTC} . More formally, the edge set E_{XTC} of the graph $G_{XTC} = (V, E_{XTC})$ is $E_{XTC} = \{(u, v) \mid \exists u: v \in N_u\}$.

In the algorithm as described above, each node constructs in Step I) a total order over all its neighbors in G . In a variant of the algorithm a node u could apply a growing radius technique—starting with the "best" neighbor—to decide on a neighbor v 's inclusion in N_u or \tilde{N}_u —based on \prec_v —immediately when iden-

tifying v as the next “worse” neighbor found so far. Applying such interleaving of steps I), II) and III), u could terminate earlier, that is, as soon as having found “enough” neighbors (where Theorem 4.2 would provide a termination criterion in the case of G being a Unit Disk Graph).

Property 1 as described in the Introduction, that is symmetry of the resulting graph, often has to be enforced by topology control algorithms (for instance by a propose-accept cycle) [21, 23]. The following theorem shows that in contrast XTC is guaranteed to “automatically” compute a graph with Property 1, without any assumptions whatsoever on the neighbor orders:

Theorem 3.1 (Symmetry). *The edges in G_{XTC} are symmetric: A node u includes a neighbor v in N_u if and only if v includes u in N_v .*

Proof. Assume for the sake of contradiction that u includes v in N_u (Assumption 1), whereas v does not include u in N_v (Assumption 2). According to Assumption 2, there exists a node $w \in N_v \cup \tilde{N}_v$ when v decides to include u in \tilde{N}_v in Line 8 ($w \prec_v u$), such that $w \prec_u v$ (Line 7). Since $w \prec_u v$ holds, $w \in N_u \cup \tilde{N}_u$ at the point of time when u decides about v ’s inclusion in N_u ; together with $w \prec_v u$, it follows that v is included in \tilde{N}_u , which is a contradiction to Assumption 1. \square

4 XTC on Euclidean Graphs

The main purpose of this section is to provide an illustration of the graph resulting from the topology control algorithm. We make three assumptions:

- i) Every node u has a unique identifier id_u . The identifiers are comparable, that is, there exists a total order “ $<$ ” defined over the set of all identifiers.
- ii) The nodes are placed in a Euclidean plane.
- iii) Every edge (u, v) is attributed a weight defined to be the triple $(|uv|, \min(id_u, id_v), \max(id_u, id_v))$, where $|uv|$ is the Euclidean distance between nodes u and v . The neighbor orders computed in Step I) of the XTC algorithm are based on the lexicographic order³ of these edge weights, that is
$$w \prec_u v \iff (|uw|, \min(id_u, id_w), \max(id_u, id_w)) < (|uv|, \min(id_u, id_v), \max(id_u, id_v))$$

Assumption i) is common in the context of distributed algorithms and viable for practical networks. Assumption ii)—although often made in order to model ad-hoc networks—is less realistic; nevertheless we adopt this assumption in this section for the sake of illustration.

Assumption iii) can for instance be realized by having each node initially transmit a control signal together with a message containing information on the control signal transmission power. With the additional assumption that the employed antennas are isotropic and that the signal can propagate without obstruction, the control signal receivers can compute an order over the Euclidean distances to the senders from the receive and transmission power levels. If all nodes send with equal transmission power, the order \prec_u is even equivalent to the relative order of only the receive power levels sensed at a node u .

³The lexicographic order of two triples is defined according to the order of the first components, or—if the first components are equal—according to the second components, or—if both the first and the second components, respectively, are equal—according to the third components. Formally: $(a_1, b_1, c_1) < (a_2, b_2, c_2) \iff (a_1 < a_2) \vee ((a_1 = a_2) \wedge (b_1 < b_2)) \vee ((a_1 = a_2) \wedge (b_1 = b_2) \wedge (c_1 < c_2))$.




Figure 1. As described in the proof for Theorem 4.2, if $\alpha < \pi/3$ and $|uv| < |uw|$, it follows that $|vw| < |uw|$.

The following theorem proves Property 2 as defined in the Introduction, that is connectivity of the topology control graph G_{XTC} . Note that this theorem does not require G to be a Unit Disk Graph; G being a Euclidean Graph is sufficient.

Theorem 4.1 (Connectivity). *Given a Euclidean Graph G , two nodes u and v are connected in G_{XTC} if and only if they are connected in G . Consequently, the graph G_{XTC} is connected if and only if G is connected.*

Proof. Since XTC exclusively considers edges in G , u and v can only be connected if they are connected in G . In order to prove the opposite direction of the above equivalence, we assume for contradiction that G_{XTC} contains at least one pair of non-connected nodes that are connected in G . Consider the pair u, v with minimum value $(|p(u, v)|, \min(id_u, id_v), \max(id_u, id_v))$, where $|p(u, v)|$ is the Euclidean cost of the shortest path connecting u and v on G , among all pairs of nodes u and v that are not connected in G_{XTC} but connected in G . The nodes u and v must be connected directly by the edge (u, v) in G ; otherwise a different pair of nodes w, x lying on the path connecting u and v would have a value $|p(w, x)|$ less than $|p(u, v)|$, and u, v would not be the pair with minimum $|p(u, v)|$. Since the edge (u, v) is in G , the cost of the shortest path connecting u and v is their Euclidean distance: $|p(u, v)| = |uv|$. According to the assumption, u includes v in \tilde{N}_u , that is, at the moment u decides so, there is a node $w \in N_u \cup \tilde{N}_u$ such that $w \prec_v u$. Since $w \in N_u \cup \tilde{N}_u$, we also have $w \prec_u v$, or $(|uw|, \min(id_u, id_w), \max(id_u, id_w)) < (|uv|, \min(id_u, id_v), \max(id_u, id_v))$. Since $(|uv|, \min(id_u, id_v), \max(id_u, id_v))$ is the least such value for any pair of non-connected nodes in G_{XTC} , and as u and w are connected in G (w is contained in \prec_u), u and w must also be connected in G_{XTC} . For the same reason and since $w \prec_v u$, also v and w must be connected in G_{XTC} , which contradicts the assumption that u and v are not connected in G_{XTC} . \square

For the remainder of this section we assume that G is a Unit Disk Graph. The following theorem proves that in this case G_{XTC} does not only feature sparseness (Property 3, as described in the Introduction), but even bounded degree (Property 3+):

Theorem 4.2 (Bounded Degree). *Given a Unit Disk Graph G , G_{XTC} has degree at most 6.*

Proof. We prove that no two adjacent edges in G_{XTC} enclose an angle less than $\pi/3$, from which the theorem follows. Assume for contradiction that the two edges (u, v) and (u, w) enclose an angle $\alpha < \pi/3$ at node u . Furthermore let v be u ’s neighbor that was included in N_u before w , that is $v \prec_u w$ or $(|uv|, \min(id_u, id_v), \max(id_u, id_v)) <$




Figure 2. The quadrangle $uvwx$ from the proof of Theorem 4.3 contains at least one angle $\beta \geq \pi/2$. It follows that if G is a Unit Disk Graph, the edge (u, w) is not contained in G_{XTC} .

$(|uw|, \min(id_u, id_w), \max(id_u, id_w))$. Since $|uv| \leq |uw|$ and $\alpha \leq \pi/3$, it follows that $|vw| < |uw|$ (cf. Figure 1). G being a Unit Disk Graph, also the edge (v, w) is in G , as $|vw| < |uw| \leq 1$. Consequently $v \prec_w u$, implying that u included w in \tilde{N}_u , which is however a contradiction to the assumption that the edge (u, w) is in G_{XTC} . \square

As an additional property, G_{XTC} contains no two intersecting edges, which allows its employment for geometric routing.

Theorem 4.3 (Planarity). *Given a Unit Disk Graph G , G_{XTC} is planar, that is, it contains no two intersecting edges.*

Proof. Suppose for the sake of contradiction that the two edges (u, w) and (v, x) intersect in G_{XTC} (cf. Figure 2). Of the quadrangle $uvwx$, at least one angle has size not less than $\pi/2$. Let this angle be w.l.o.g. β , adjacent to node v . Since $\beta \geq \pi/2$, $|uv| < |uw|$ and $|vw| < |uw|$. When node u considered the inclusion of w in N_u , v was consequently already in $N_u \cup \tilde{N}_u$ and $v \prec_w u$, causing u to include w in \tilde{N}_u , which is however a contradiction to the assumption that the edge (u, w) is in G_{XTC} . \square

In the following theorem we describe the relationship between G_{XTC} and the Relative Neighborhood Graph of G [20]. The Relative Neighborhood Graph is defined to contain all edges $(u, v) \in G$, such that there exists no node w with $|uw| < |uv| \wedge |vw| < |uv|$ (cf. Figure 4(a)).

Theorem 4.4. *Given a Unit Disk Graph G , G_{XTC} is a subgraph of the Relative Neighborhood Graph computed on G . If G contains no node having two or more neighbors at exactly the same distance, G_{XTC} is identical to the Relative Neighborhood Graph.*

Proof. The subgraph relationship follows from the fact that if G_{XTC} contains an edge (u, v) , there exists no node w with $|uw| < |uv| \wedge |vw| < |uv|$, which implies that (u, v) is also contained in the Relative Neighborhood Graph RNG. Furthermore XTC excludes an edge (u, v) that is preserved in RNG—there is no w with $|uw| < |uv| \wedge |vw| < |uv|$ —only if there exists a node w with $w \prec_u v \wedge w \prec_v u$, which is in total possible only if $|uw| = |uv|$ and the enclosing angle $\angle vuw \leq \pi/3$ or $|vw| = |uv|$ and the enclosing angle $\angle uvw \leq \pi/3$. \square

5 XTC on General Weighted Graphs

In realistic ad-hoc networks nodes are not located in a plane and received transmission power does not only depend on the distance to the sender, but above all on physical obstacles between




Figure 3. The edge weights ω_{uv} , ω_{uw} , and ω_{vw} reflect that signal propagation between u and w is impaired by a physical obstacle (wall, building, hill): $\omega_{uv} < \omega_{uw}$ and $\omega_{vw} < \omega_{uw}$. In contrast to typical topology control algorithms based on node positions, XTC does not include the edge (u, w) in its result graph, but exploits that a better connection exists via node v .

sender and receiver. As one of the main properties of such real ad-hoc networks, however, symmetry of physics is preserved: The attenuation factor of a link between two network nodes is identical to signal propagation in either direction. Accordingly an ad-hoc network can be modeled by a weighted graph, where each edge is attributed a weight representing the corresponding signal attenuation factor (cf. Figure 3). More abstractly the edge weights can be considered qualities of links between node pairs. Assuming isotropic antennas, a node can obtain its neighbor order with a technique similar to the one described in Section 4 by sending a control signal. If the edge weights are considered link quality indicators in a more general sense, these weights and consequently the neighbor ordering can be established by exchange of probe messages.

In this section we show that the XTC algorithm computes a connected topology even in general weighted graphs modeling realistic ad-hoc networks. We assume that the neighbor order of a




Figure 4. Definitions of the Relative Neighborhood Graph (a) and the Gabriel Graph (b). In the Relative Neighborhood Graph the edge (u, v) exists if and only if the hatched lune (excluding its boundary) does not contain a third node. In the Gabriel Graph the edge (u, v) exists if and only if the hatched circle (including its boundary) does not contain a third node.




Figure 5. The Unit Disk Graph G (left), the Gabriel Graph of G (center), and G_{XTC} of 1400 nodes placed randomly and uniformly on a square field of 20 units side length.

node u —as employed in the algorithm—corresponds to the order over the weights of the edges adjacent to u .

Theorem 5.1 (Connectivity). *Given a general weighted graph G , two nodes u and v are connected in G_{XTC} if and only if they are connected in G . Consequently, the graph G_{XTC} is connected if and only if G is connected.*

Proof. This theorem can be proved in analogy to the proof of Theorem 4.1, substituting edge weights for Euclidean distances. \square

What furthermore can be stated for G being a general weighted graph with respect to G_{XTC} 's sparseness (Property 3), is that G_{XTC} cannot contain cycles of length three:

Theorem 5.2. *Given a general weighted graph G , G_{XTC} has girth 4, that is, the shortest cycle in G_{XTC} is of length 4.*

Proof. It is sufficient to show that G_{XTC} does not contain any cycles of length 3. We suppose for contradiction that there exists such a cycle through the nodes u, v , and w , that is, all three edges (u, v) , (v, w) , and (u, w) are contained in G_{XTC} . Let us further assume w.l.o.g. that $w \prec_u v$, or $\omega_{uw} < \omega_{uv}$, where ω_{uv} is the weight of the edge (u, v) . At the point of time when u considered the inclusion of v in N_u , w had already been processed, which means that since (u, v) is in G_{XTC} , $u \prec_v w$ must hold, and consequently also $\omega_{uv} < \omega_{vw}$. Applying the same argument from v 's point of view yields that also $v \prec_w u$ or $\omega_{vw} < \omega_{uw}$ must hold, which provokes a contradiction. \square

Although G_{XTC} has girth 4, it does not feature sparseness: Constructed on a general weighted graph G , G_{XTC} can have degree $\Theta(n)$ and contain $\Theta(n^2)$ edges. An example for such a graph is $K_{n/2, n/2}$, the complete bipartite graph with $n/2$ nodes in each partition set, in which each node has degree $n/2$ and which consequently contains $n^2/4$ edges in total. Network graphs resulting from real ad-hoc networks can however be expected to be considerably sparser.

6 Average-Case Evaluation

In this section we present the properties of the topology control graph G_{XTC} on average-case Euclidean graphs, that is on

graphs generated by randomly and uniformly placing nodes in a given field. In particular we study the spanner and bounded degree properties—Properties 2+ and 3+, as stated in the Introduction—in the context of average graphs. The bounded degree property having been shown to hold for G_{XTC} in Section 4, we demonstrate in this section that also the spanner property holds on average graphs. Since, G_{XTC} being a planar graph, it lends itself to geometric routing, we furthermore examine the influence of G_{XTC} on such routing algorithms.

In order to model the physical network in our average-case evaluation we adopt the Unit Disk Graph definition, in which an edge exists if and only if its Euclidean length is less than one unit. To assess the average-case properties of G_{XTC} , we compare it with the Gabriel Graph [4]. Being one of the most prominent topology control structures on the one hand, the Gabriel Graph is on the other hand particularly well suited for comparison with respect to the spanner property, since it is not only a spanner with respect to the energy metric, but even contains the energy-minimal path between any pair of points.

The Gabriel Graph is—similarly to the Relative Neighborhood Graph—defined such that the presence of an edge (u, v) depends on whether a certain geometric area contains a third node w or not. In the case of the Gabriel Graph this geometric area is the circle (including the boundary) having the line segment \overline{uv} as a diameter (cf. Figure 4(b)).

Figure 5 illustrates the Gabriel Graph and G_{XTC} constructed from a sample Unit Disk Graph. The figure shows well that, informally speaking, areas with high edge density in the Unit Disk Graph are thinned out by topology control, while preserving connectivity of the graph. This tendency can be observed even more clearly for G_{XTC} than for the Gabriel Graph.

As studied in percolation theory [3], *network density* is an important parameter influencing the properties of average-case networks. The transition between the two extremes with respect to network density—very low densities, where hardly any pair of nodes is connected, and very high densities, where disconnection of the network is extremely improbable—takes place in a relatively narrow *critical density range* roughly around 5 nodes per unit disk. In order to account for this effect we acquired our measurements and simulation results over a spectrum of network densities. For each considered network density the number of nodes




Figure 6. Stretch factors of G_{XTC} with respect to energy (solid) and Euclidean (dotted) metric, stretch factor of Gabriel Graph with respect to Euclidean metric (dashed) over network density spectrum. Mean values are plotted in black, maximum values in gray. The Gabriel Graph contains the energy optimal path; its stretch factor curve ($\equiv 1$, light gray) is plotted for reference.

corresponding to the density was randomly and uniformly placed on a square field with side length 20 units.

6.1 Spanner Property

In order to study the spanner property of G_{XTC} on randomly generated networks we calculated the stretch factor of a pair of nodes u, v

$$s(u, v) := \frac{|p_{tc}(u, v)|}{|p(u, v)|},$$

that is the ratio between cost of the shortest path between u and v on the topology control graph and the shortest path on the Unit Disk Graph G . For each considered network density we generated 2000 networks and also randomly selected a pair of nodes u, v to calculate $s(u, v)$. As topology control graphs we employed G_{XTC} and the Gabriel Graph. Edge and correspondingly path costs were considered with respect to Euclidean edge length $|uv|$ and to energy with an attenuation exponent 2, that is $|uv|^2$.

Figure 6 depicts our results over the considered network density range. Since the Gabriel Graph contains the energy-minimal path connecting any pair of nodes, the line $s(u, v) \equiv 1$ is plotted for reference. Although G_{XTC} is not an energy spanner in the strict sense, the results show that this graph has a good energy spanning property on average graphs: The corresponding mean value of the stretch factor does not exceed 1.06, whereas even the maximum value stays below 1.9. With respect to the Euclidean metric, the mean value curves for both the Gabriel Graph and G_{XTC} remain—throughout the density range—almost constant at low values of below 1.1 and 1.25, respectively; the corresponding maximum curve for G_{XTC} is less stable than for the Gabriel Graph, but only rarely reaches values above 3.

In summary, the results show that G_{XTC} is a good average-case spanner with respect to the Euclidean metric, but especially for the energy metric.

6.2 Bounded Degree Property

G_{XTC} being shown to have degree at most 6 in Theorem 4.2, we study here its average-case behavior with respect to node degree by comparison with the corresponding behavior of the Gabriel Graph. The results therefor were obtained similarly as for the spanner property. From each of the 2000 random networks—generated for each considered network density—a randomly chosen node was examined regarding its degree in the Unit Disk Graph, the Gabriel Graph G , and G_{XTC} .

Figure 7 shows the acquired results. Mean and maximum degree values for the Unit Disk Graph rise in accordance with network density. The mean degree curves for G_{XTC} and the Gabriel Graph increase slowly for very low densities, remain however—once beyond the critical density range—constant at approximately 2.5 and 3.9, respectively. Although inherently less stable, the maximum degree numbers for G_{XTC} and the Gabriel Graph stay within a narrow range from 4 to 5 and between 7 and 9. The low degree values of G_{XTC} suggest its suitability to reduce interference in ad-hoc networks.

6.3 Performance of Geometric Routing

G_{XTC} 's planarity property enables it to be employed for geometric routing. We study its influence on geometric routing again by comparison with the Gabriel Graph.

For this purpose the results were obtained by simulation of the GOAFR⁺ algorithm [11] on G_{XTC} and the Gabriel Graph of 2000 randomly generated networks for each considered network density. In particular, the performance measure for the algorithm A routing from a source s to a destination t on a network G (in our case s and t having been chosen randomly from G) is defined as

$$perf_A(G, s, t) := \frac{|p_A(G, s, t)|_d}{|p^*(G, s, t)|_d},$$

that is the number of steps taken by the algorithm normalized by the hop length of the shortest path from s to t on G .




Figure 7. Node degree of G_{XTC} (solid), Gabriel Graph (dashed), and Unit Disk Graph (dotted) over network density spectrum. Mean values are plotted in black, maximum values in gray.

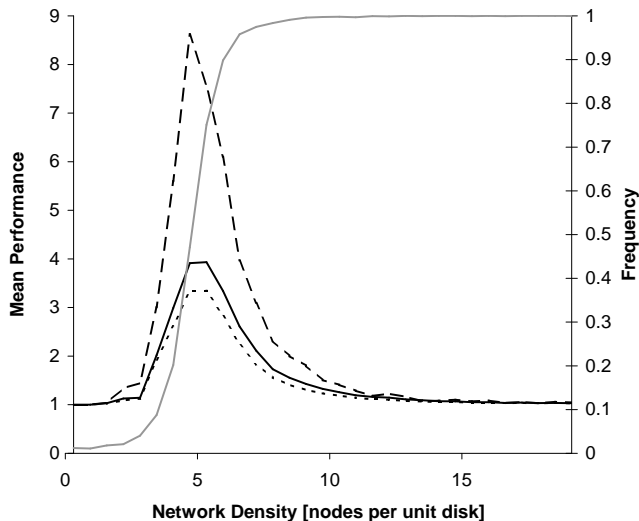


Figure 8. Mean performance of the GOAFR⁺ routing algorithm on G_{XTC} (solid), on the Gabriel Graph (dotted), and of the GFG/GPSR algorithm on the Gabriel Graph (dashed). For reference, the network connectivity rate is plotted against the right y axis.

Figure 8 shows the typical bell-shaped characteristic of the mean algorithm performance curves around the critical density range. The critical network density range is additionally identified by the sharp increase of the network connectivity rate, that is the frequency with which a randomly chosen node pair at a given network density is connected. The figure depicts that above all in the critical density range the performance of the GOAFR⁺ algorithm on G_{XTC} is slightly worse than on the Gabriel Graph. In light of the performance curve of the well-known GFG/GPSR routing algorithm, however, the degradation of GOAFR⁺ employed on G_{XTC} compared to the Gabriel Graph can be considered negligible.

7 Conclusion

In this paper we introduce the XTC topology control algorithm. Compared to previous proposals for topology control XTC has three main advantages. First, it is not only simple, but also local: Every node communicates with its neighbors in the network not more than twice. Second, unlike many other topology control algorithms, XTC does not require the network graph to be a Euclidean Graph, let alone a Unit Disk Graph. Also for the global case of the network graph being a general weighted graph, XTC proves correct and computes a resulting subgraph maintaining connectivity. Third, while previously proposed topology control algorithms commonly assumed that exact node and neighbor position information is available, XTC does not require this assumption. The algorithm works with a general notion of a quality order over a node's neighbors. Whereas correctness of the algorithm can be shown even without any strict assumptions on this neighbor order, the resulting topology features the bounded degree property provided that the neighbor order corresponds to Euclidean distances and that the network is a Unit Disk Graph. On average-case random Unit Disk Graphs the resulting graph also

shows good spanner properties, above all with respect to the energy metric. Being planar, the proposed topology is finally also suitable for geometric routing.

References

- [1] K. Alzoubi, P.-J. Wan, and O. Frieder. Message-Optimal Connected Dominating Sets in Mobile Ad Hoc Networks. In *Proc. of the 3rd ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2002.
- [2] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with Guaranteed Delivery in ad hoc Wireless Networks. In *Proc. of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, pages 48–55, 1999.
- [3] O. Dousse, P. Thiran, and M. Hasler. Connectivity in ad-hoc and hybrid networks. In *Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.
- [4] K. Gabriel and R. Sokal. A New Statistical Approach to Geographic Variation Analysis. *Systematic Zoology*, 18:259–278, 1969.
- [5] T. Hou and V. Li. Transmission Range Control in Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 34(1):38–44, 1986.
- [6] L. Hu. Topology Control for Multihop Packet Radio Networks. *IEEE Trans. on Communications*, 41(10), October 1993.
- [7] L. Jia, R. Rajaraman, and C. Scheideler. On Local Algorithms for Topology Control and Routing in Ad Hoc Networks. In *Proc. of the 15th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, 2003.
- [8] L. Jia, R. Rajaraman, and R. Suel. An Efficient Distributed Algorithm for Constructing Small Dominating Sets. In *Proc. of the 20th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 33–42, 2001.
- [9] B. Karp and H. Kung. GPSR: Greedy Perimeter Stateless Routing for Wireless Networks. In *Proc. of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 243–254, 2000.
- [10] F. Kuhn and R. Wattenhofer. Constant-Time Distributed Dominating Set Approximation. In *Proc. of the 22nd ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.
- [11] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric Routing: Of Theory and Practice. In *Proc. of the 22nd ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.
- [12] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-Hoc Networks Beyond Unit Disk Graphs. In *Proc. of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing*, 2003.
- [13] F. Kuhn, R. Wattenhofer, and A. Zollinger. Worst-Case Optimal and Average-Case Efficient Geometric Ad-Hoc Routing. In *Proc. of the 4th ACM Int. Symposium on Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, 2003.
- [14] L. Li, J. Halpern, V. Bahl, Y.-M. Wang, and R. Wattenhofer. Analysis of a Cone-Based Distributed Topology Control Algorithm for Wireless Multihop Networks. In *Proc. of the 20th ACM Symposium on Principles of Distributed Computing (PODC)*, 2001.
- [15] X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed Construction of Planar Spanner and Routing for Ad Hoc Wireless Networks. In *Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.
- [16] R. Prakash. Unidirectional Links Prove Costly in Wireless Ad-Hoc Networks. In *Proc. of the 3rd International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM)*, 1999.
- [17] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment. In *Proc. of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2000.
- [18] V. Rodoplu and T. H. Meng. Minimum Energy Mobile Wireless Networks. *IEEE J. Selected Areas in Communications*, 17(8), August 1999.
- [19] H. Takagi and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transactions on Communications*, 32(3):246–257, 1984.
- [20] G. Toussaint. The Relative Neighborhood Graph of a Finite Planar Set. *Pattern Recognition*, 12(4):261–268, 1980.
- [21] Y. Wang and X.-Y. Li. Distributed Spanner with Bounded Degree for Wireless Ad Hoc Networks. In *2nd International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, 2002.
- [22] Y. Wang and X.-Y. Li. Localized Construction of Bounded Degree Planar Spanner. In *Proc. of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing*, 2003.
- [23] 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. of the 20th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2001.
- [24] A. C.-C. Yao. On Constructing Minimum Spanning Trees in k-Dimensional Spaces and Related Problems. *SIAM Journal on Computing*, 11(4), 1982.