Topology Control with Limited Geometric Information

Kevin Lillis and Sriram V. Pemmaraju

Department of Computer Science
The University of Iowa, Iowa City, IA 52242-1419, USA,
[lillis, sriram]@cs.uiowa.edu

Abstract. Topology control is the problem of selecting neighbors for each node in a wireless network, so that the resulting network has a number of useful properties. More precisely, a topology control protocol P takes as input a network G and aims to construct a spanning subgraph G_P, that is sparse, "energy minimizing" and has sufficient connectivity so as to guarantee multiple short paths between pairs of nodes in G. Currently, topology control protocols assume that nodes in G reside in some Euclidean (usually, 2-dimensional) space and rely on geometric information such as node locations and pairwise distances between nodes to produce G_P with appropriate properties. However, these protocols are extremely sensitive to errors in location information and this feature makes them impractical because errors in location and distance information are pervasive in practical systems. This paper presents and analyzes two randomized topology control protocols that are tolerant to errors in pairwise distance estimates. The first protocol, called RTC (short for randomized topology control) uses *no* geometric information, relying only on connectivity information and is therefore completely immune to errors in location or distance information. The second protocol, called ε-RTC, generalizes the first protocol. Allowing for errors in distance estimates, but assuming that relative errors are bounded above by ε, the second protocol produces an output network that is symmetric, connected, sparse, and has good spanner properties. As $\varepsilon \to 0$, ε-RTC behaves like the XTC protocol (R. Wattenhofer and A. Zollinger, "XTC: A practical topology control algorithm for ad-hoc networks", WMAN 2004) and for large values of ε, it behaves like RTC. Our results hold whenever the input network is a unit disk graph or even a quasi unit disk graph.

1 Introduction

An ad-hoc wireless network consists of a set of nodes, each equipped with a wireless radio. Each node u can send messages to all nodes within its radio range and all such nodes are potential neighbors of u. However, for reasons explained below, it is preferable for u to communicate with an appropriately chosen subset of these reachable nodes. Informally speaking, the *topology control* problem is one of selecting neighbors for each node so that the resulting network has a number of useful properties. More precisely, let V be a set of nodes that can communicate via wireless radios and for each $v \in V$, let $N(v)$ denote the set of all nodes that v can reach when transmitting at maximum power. The induced digraph $G = (V, E)$, where $E = \{(u, v) \mid v \in N(u)\}$, represents the network in which every node has chosen to transmit at maximum power and has designated every node it can reach, as its neighbor. The topology control problem is the problem of devising a protocol P for selecting a set of neighbors $N_P(v) \subseteq N(v)$ for each node $v \in V$. The induced digraph $G_P = (V, E_P)$, where $E_P = \{(u, v) \mid v \in N_P(u)\}$ is typically required the satisfy properties such as symmetry (if $v \in N_P(u)$ then $u \in N_P(v)$), sparseness ($|E_P| = O(|V|)$) or bounded degree ($|N_P(v)| \leq c$ for all nodes v and some constant c), connectivity, and the spanner property.

Current research on topology control. In the last few years, a generation of topology control protocols have been proposed that achieve many of the properties mentioned above by assuming that the nodes in V lie in some Euclidean (typically, 2-dimensional) space and each node knows its spatial location with respect to some global coordinate system [4–7, 10]. In our view, the main problem with these protocols is the reliance on node location information and the lack of robustness with respect to errors in this information. Node location information is typically available only if nodes are GPS enabled or if an expensive protocol called *localization* [2] is run. GPS enabled nodes are costly and more importantly, consume scarce energy resources. No matter which approach is used to find node locations, errors in location information are quite likely. Unfortunately, none of the topology control protocols mentioned above, make allowances for any errors in location information and are extremely sensitive to these errors. In other words, critical properties of the output network such as connectivity or bounded maximum degree are not guaranteed to hold even if there is a small amount of error in location information.

More recently, Wattenhofer and Zollinger [11] have proposed a topology control protocol called XTC that does not rely on specific location information for each node, but rather requires each node to only know the distance to each of its neighbors. Although this is an improvement over location-based topology control algorithms, XTC still suffers from lack of robustness to errors in distance information. As we show in Section 2, there are networks modeled as unit disk graphs (UDGs) such that when XTC is run on these, its output network is disconnected and contains vertices with unbounded degree, even in the presence of arbitrarily small errors in pairwise distance estimates. This is a significant problem because in general it seems hard to accurately estimate pairwise node distances. For example, [11] mentions the use of the strength of the received signal (RSSI) as a way to estimate distances. While this technique is relatively cheap and does not need additional hardware, it is known to have low accuracy. Alternate techniques such as the use of ultrasound hardware have been proposed [9]. These have better accuracy than RSSI, but are significantly costlier, both in terms of additional hardware and in terms of energy consumption.

Our results. We present two randomized topology control protocols. The first protocol, called RTC (short for "randomized topology control") uses *no* geometric information, relying only on connectivity information. As a result RTC is completely immune to errors in distance or location information. Our second protocol, called ε-RTC is parameterized by $\varepsilon > 0$, which stands for the maximum relative error on pairwise distance estimates. Allowing for errors in distance estimates, but assuming that relative errors are bounded above by ε, ε-RTC produces an output graph that is symmetric, connected, sparse, and has spanner properties. As $\varepsilon \to 0$, ε-RTC behaves like XTC and for large values of ε, ε-RTC behaves like RTC. In general, ε-RTC combines the advantages of XTC and RTC. Unlike XTC, ε-RTC is tolerant to errors bounded by ε and unlike RTC, ε-RTC uses distance information to the extent they are reliable and attempts to drop long links in favor of short links, thereby saving on energy consumption. Both RTC and ε-RTC are randomized variants of XTC and are therefore extremely light weight, needing only two rounds of communication. We prove properties of the output network of RTC and ε-RTC assuming that the input network G is a UDG. However, our results hold even when G is a quasi UDG [1].

2 The XTC protocol

We start this section by reproducing the XTC protocol from [11], followed by a description of properties of XTC.

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 [11], the XTC protocol (shown above) 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 $\prec$. An appropriate choice of $\prec$ is critical to the success of XTC.

It is shown in [11] that G_{XTC} is symmetric provided G is and this is independent of $\prec$. It is also shown that if G is a Euclidean graph and $\prec_u$ is defined as

$$v \prec_u w \Leftrightarrow (|uv|, \min\{id_u, id_v\}, \max\{id_u, id_v\}) < (|uw|, \min\{id_u, id_w\}, \max\{id_u, id_w\}), \quad (1)$$

then G_{XTC} is 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. Finally, it is shown that if $\prec$ is a distance-based ordering and G is a UDG, then the maximum degree in G_{XTC} is at most 6 and G_{XTC} is planar.

Even though XTC is fast and simple and its output graph has many desired properties, it is extremely sensitive to small perturbations in the neighborhood orderings. In the following subsections we show XTC's lack of robustness to small errors.

2.1 G_{XTC} may be disconnected

We now present a simple example of a 4-vertex unit disk graph that illustrates the lack of robustness of XTC. As shown in [11], if neighborhood orderings are distance-based (as in (1)), then G_{XTC} is connected. Note that according to this definition, each node u orders its neighbors in increasing order of distance, breaking ties using ids. We show that if the distance estimates are erroneous, even slightly, then the resulting neighborhood orderings are such that G_{XTC} becomes disconnected.

Consider the unit disk graph shown in Figure 1. Pick an ε, $0 < \varepsilon < 1 - \frac{1}{\sqrt{2}}$. Let the lengths of the edges be $|ab| = |dc| = (1-\varepsilon)/2$ and $|ac| = |bd| = 1/2$. Then the neighborhood orderings $\prec$ according to (1) are:

$$d \prec_a b \prec_a c \qquad c \prec_b a \prec_b d \qquad b \prec_c d \prec_c a \qquad a \prec_d c \prec_d b.$$

Now suppose that node b incorrectly estimates distance $|ba|$ as $(1+\varepsilon)/2$ and node c incorrectly estimates distance $|cd|$ as $(1+\varepsilon)/2$. The resulting neighborhood orderings $\tilde{\prec} = \{\tilde{\prec}_a, \tilde{\prec}_b, \tilde{\prec}_c, \tilde{\prec}_d\}$ are shown below.

$$d \mathrel{\tilde{\prec}_a} b \mathrel{\tilde{\prec}_a} c \qquad c \mathrel{\tilde{\prec}_b} d \mathrel{\tilde{\prec}_b} a \qquad b \mathrel{\tilde{\prec}_c} a \mathrel{\tilde{\prec}_c} d \qquad a \mathrel{\tilde{\prec}_d} c \mathrel{\tilde{\prec}_d} b.$$

If XTC is run with input $\tilde{\prec}$ then G_{XTC} contains just the two edges $\{a, d\}$ and $\{b, c\}$ and is therefore disconnected. Thus two incorrect estimates by an arbitrarily small amount ε is sufficient to break connectivity.

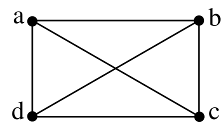

Fig. 1. A unit disk graph for showing the sensitivity of XTC to small perturbations. The lengths of the edges are $|ab| = |dc| = (1-\varepsilon)/2$ and $|ac| = |bd| = 1/2$.

2.2 G_{XTC} may have high degree

To start with, suppose that we have three nodes u, v_1, and v_2 (refer to Figure 2). Fix an $\varepsilon, 0 < \varepsilon < 1/2$ and let $|uv_1| = d$ and $|v_1v_2| = \frac{d}{1+\varepsilon}$. If the length of the third edge $|uv_2|$ is $\frac{\varepsilon d}{1+\varepsilon}$, then the three points u, v_2, and v_1 would be collinear. To make $\triangle uv_1v_2$ a non-trivial triangle, pick $|uv_2| = \frac{2\varepsilon d}{1+\varepsilon}$. If XTC were run, just on $\triangle uv_1v_2$, then vertex u would drop v_1 from its set of neighbors in favor of vertex v_2. Likewise vertex v_1 would drop u from its set of neighbors in favor of vertex v_2. This would eliminate the edge $\{u, v_1\}$[1] from the output graph. Now suppose that there are errors in distance estimates such that lengths $|uv_1|$ and $|uv_2|$ are unchanged, but the length $|v_1v_2|$ is overestimated by a factor of $(1+\varepsilon)$, thus making $|v_1v_2| = |uv_1|$. Now, because ties are broken by node ids, we can assume that $id_u < id_{v_2}$, thus forcing XTC to drop edge $\{v_1, v_2\}$, but retain $\{u, v_1\}$. This phenomenon can be forced to repeat arbitrarily. Specifically, u could have a third neighbor v_3 that plays the role relative to v_2 that v_2 played relative to v_1. Let the actual distances to the new vertex v_3 satisfy

$$|v_2v_3| = \frac{2\varepsilon d}{(1+\varepsilon)^2}, \qquad |uv_3| = \left(\frac{2\varepsilon}{1+\varepsilon}\right)^2 d, \qquad |v_1v_3| > \frac{d}{(1+\varepsilon)},$$

and perturb $|v_1v_3|$ and $|v_2v_3|$ by a factor of $(1+\varepsilon)$ such that it would appear that $|v_2v_3| = |uv_2|$ and $|v_1v_3| > d$. In addition, if we assume that $id_u < id_{v_3}$, then $\{v_1, v_3\}$ and $\{v_2, v_3\}$ would be dropped and $\{u, v_2\}$ would be retained. By continuing this construction, we can force u to have arbitrarily high degree. Note that in this example, some distances are unchanged while some are increased by a factor of $(1+\varepsilon)$.

[1] XTC outputs a symmetric network independent of the neighborhood orderings $\prec$. This allows us to think of the edges of G_{XTC} as undirected edges.

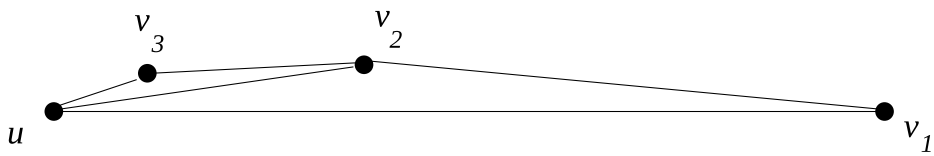

Fig. 2. A unit disk graph to illustrate how XTC may produce an output graph with unbounded node degree.

2.3 Experimental evidence

We have experimentally studied XTC's sensitivity to errors in distance information. To do this, we fixed a value for $\varepsilon > 0$ and generated a sequence of UDGs (a brief description of our experimental setup appears in the Appendix in Section A). For each instance G of a UDG and for each edge $\{u, v\}$ in G, we picked two distance estimates for $|uv|$, one to be assigned to u and the other to v. Both distance estimates are picked uniformly at random from the interval $[|uv| \cdot (1-\varepsilon), |uv| \cdot (1+\varepsilon)]$. Our experiments show that even for randomly generated instances of UDGs, XTC produces disconnected output networks a substantial fraction of the time. For example, for UDGs generated with node density of 3 nodes per unit square, for $\varepsilon = 0.4$, XTC produces a disconnected network 60% of the time. Even for UDGs of higher density, say 12 nodes per unit square, XTC produces a disconnected network 40% of the time at $\varepsilon = 0.4$. See [8] for more details.

The plots below shows the increase in the maximum degree of G_{XTC}, as $\varepsilon \to 1$. The plot on the left is for input UDGs with a density of 15 nodes per unit square and the plot on the right is for input UDGs with a density of 30 nodes per unit square. In each case, at a certain value of ε, the maximum degree of G_{XTC} exceeds beyond 5, which is the upper bound on $\Delta(G_{XTC})$ assuming completely accurate distances. Notice that in the case of the higher density graphs, $\Delta(G_{XTC})$ exceeds 5 at around $\varepsilon = 0.25$.

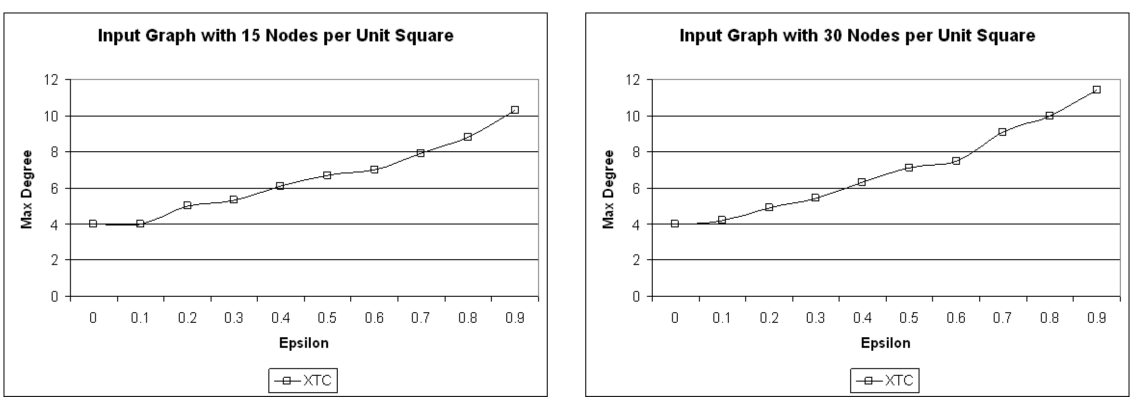

Fig. 3. These are plots showing the increase in the maximum degree of G_{XTC} as the error bound ε increases.

3 Randomized Topology Control

Here we describe two randomized topology control protocols, RTC and ε-RTC. Both are variants of XTC in that they first construct specific neighborhood orderings and then run XTC with these orderings as input. RTC ignores geometric information completely, relying on randomization alone to obtain good expected performance. ε-RTC does use the distance estimates, even though they may be erroneous. However, in an attempt to foil the adversary, ε-RTC first does a random perturbation of the distance estimates. The amount of perturbation is a function of ε.

3.1 The RTC protocol

RTC consists of two phases. In phase 1, a randomized edge labeling is constructed and this is used by each node u to define a neighborhood ordering $\prec_u$ on $N(u)$. In Phase 2, XTC is executed with these neighborhood orderings as input. Phase 1 of RTC, which we call `NeighborhoodOrdering`, is shown below. It simply consists of picking for each edge, a real number, uniformly at random from the range $[0, 1]$. The choices for different edges are independent. Each node u maintains an array $d_u[v]$, $v \in N(u)$, of local variables and for each edge $\{u, v\} \in E(G)$, one of the two endpoints u or v whichever has higher id, picks a real number $d \in [0, 1]$ to serve as the edge label for $\{u, v\}$ and this is assigned to both $d_u[v]$ and to $d_v[u]$. Finally, each node u constructs $\prec_u$ by ordering its neighbors in increasing order of the values $d_u[v]$, $v \in N(u)$. We call the output of RTC, G_{RTC}.

Algorithm `NeighborhoodOrdering`(u)

1. Node u sends to each neighbor $v \in N(u)$, the value id_u. It receives from each neighbor $v \in N(u)$, the value id_v.
2. For each neighbor v with $id_v < id_u$, node u picks $d \in [0, 1]$ uniformly at random and assigns $d_u[v] := d$ and sends $d_u[v]$ to v.
3. For each neighbor v with $id_v > id_u$, node u receives $d_v[u]$ from v and assigns $d_u[v] := d_v[u]$.
4. Node u computes an ordering $\prec_u$ of its neighborhood $N(u)$ such that for any pair of vertices $v_1, v_2 \in N(u)$:
$$v_1 \prec_u v_2 \Leftrightarrow (d_u[v_1], id_{v_1}) < (d_u[v_2], id_{v_2}).$$

3.2 The ε-RTC protocol

Let $\alpha(u, v)$ denote the distance between u and v, *as estimated by* u. It is possible that $\alpha(u, v) \neq \alpha(v, u)$. We assume that the errors in distance estimates are bounded. That is, there is an $\varepsilon > 0$, such that

$$(1 - \varepsilon) \cdot |uv| \leq \alpha(u, v) \leq (1 + \varepsilon) \cdot |uv|.$$

When the distance estimates satisfy this property, we say that they are *ε-error bounded.*

Like RTC, the ε-RTC protocol also consists of two phases, the first involves constructing neighborhood orderings and the second involves calling XTC. The first phase, called ε-`NeighborhoodOrdering`, is shown below. We assume that all nodes know the value of ε and each node u can compute distance estimates $\alpha(u, v)$ to all neighbors $v \in N(u)$. Unlike in RTC, ε-`NeighborhoodOrdering` does make explicit use of the estimated pairwise distances $\alpha(\cdot, \cdot)$. For each edge $\{u, v\}$, first the average of $\alpha(u, v)$ and $\alpha(v, u)$ is computed. Then, an interval around this average is constructed and a value is picked, uniformly at random from this interval. This value is assigned to both $d_u[v]$ and $d_v[u]$ as the final label of the edge $\{u, v\}$. We call the output of ε-RTC, $G_{\varepsilon\text{-}RTC}$.

Algorithm ε-`NeighborhoodOrdering`(u)

1. Node u sends to each neighbor $v \in N(u)$, the value id_u. It receives from each neighbor $v \in N(u)$, the value id_v.
2. Node u estimates the distance $\alpha(u, v)$ to each neighbor $v \in N(u)$. Then node u sends to each neighbor $v \in N(u)$ the estimate $\alpha(u, v)$ and receives from each neighbor $v \in N(u)$ the estimate $\alpha(v, u)$.
3. For each neighbor v with $id_v < id_u$, node u computes $a := (\alpha(u, v) + \alpha(v, u))/2$ and then picks $d \in [a(1-\delta_L), a(1+\delta_R)]$ uniformly at random and assigns $d_u[v] := d$ and sends $d_u[v]$ to v. Here, $\delta_L = 2\varepsilon/(1+\varepsilon)$ and $\delta_R = 2\varepsilon/(1-\varepsilon)$.
4. For each neighbor v with $id_v > id_u$, node u receives $d_v[u]$ from v and assigns $d_u[v] := d_v[u]$.
5. Node u computes an ordering $\prec_u$ of its neighborhood $N(u)$ such that for any pair of vertices $v_1, v_2 \in N(u)$:
$$v_1 \prec_u v_2 \Leftrightarrow (d_u[v_1], id_{v_1}) < (d_u[v_2], id_{v_2}).$$

4 Analysis of RTC and ε-RTC

In this section we show that the output networks produced by RTC and ε-RTC are sparse. Specifically, we show that for any node u, its expected degree in G_{RTC} is bounded above by $O(\log deg_G(u))$ and its expected degree in $G_{\varepsilon\text{-}RTC}$ is bounded above by $O(1)$. We use the notation $deg_G(u)$ to denote degree of node u in the graph G. Before we prove our sparsity results, we quickly show that both G_{RTC} and $G_{\varepsilon\text{-}RTC}$ are symmetric and connected.

It is observed in [11] that if G is symmetric, then so is G_{XTC} no matter what $\prec$ is. As a corollary we obtain that both G_{RTC} and $G_{\varepsilon\text{-}RTC}$ are symmetric since the algorithms RTC and ε-RTC are just implementations of XTC with specific choices of $\prec$. Connectivity of G_{RTC} and $G_{\varepsilon\text{-}RTC}$ follows from the following result due to [3, 11]. Before we state the result we need a definition.

Definition. Let S be an arbitrary set on which a total order $<_S$ is defined. The collection of neighborhood orderings $\prec = \{\prec_u \mid u \in V(G)\}$ is said to be *consistent* if there is a labeling of the edges $\ell : E \to S$ such that for any two neighbors $v_1, v_2 \in N(u)$,

$$v_1 \prec_u v_2 \Leftrightarrow (\ell\{u, v_1\}, id_{v_1}) < (\ell\{u, v_2\}, id_{v_2}).$$

In the above, $<$ denotes the lexicographic ordering on $S \times ID$, where ID is the space of all node ids. It is worth emphasizing that E is the set of *undirected* edges and therefore every edge $\{u, v\}$ gets a single label, that is, $\ell\{u, v\} = \ell\{v, u\}$.

Theorem 1. *[3, 11] Let $\prec$ be a collection of consistent neighborhood orderings of G. If XTC is executed with input $\prec$ and G is connected then G_{XTC} is connected.*

The above theorem essentially says that as long as some edge labels are agreed upon and used to construct the neighborhood orderings, connectivity is guaranteed. These edge labels need have nothing to do with actual pairwise distances. From the fact that RTC and ε-RTC both execute XTC with a consistent collections of neighborhood orderings, we obtain that both G_{RTC} and $G_{\varepsilon\text{-}RTC}$ are connected.

Theorem 2. *G_{RTC} and $G_{\varepsilon\text{-}RTC}$ are both symmetric and connected.*

4.1 Bound on vertex degrees

Here is a simple and useful fact about the probability that the neighbors of a node u are ordered in $\prec_u$ in a certain way. This follows from the fact that in RTC, for any node u, each ordering of its neighbors is equally likely to be $\prec_u$.

Lemma 1. *Let $v_1, v_2, \ldots, v_t$ be neighbors of u. Then, $Prob[\bigwedge_{i=2}^{t} v_1 \prec_u v_i] = \frac{1}{t}$.*

Theorem 3. *Let G be a UDG and let $H = G_{RTC}$. For any vertex u of G*

$$E[deg_H(u)] = O(\log deg_G(u)).$$

Proof. Since G is a UDG, the neighborhood $N(u)$ of u can be partitioned into at most 5 cliques. For some integer t, $1 \leq t \leq 5$, let $\{N^i(u) \mid i = 1, 2, \ldots, t\}$ be a partition of $N(u)$ into t cliques. Let $deg_G^i(u) = |N^i(u)|$, for each $i = 1, 2, \ldots, t$ and let $deg_H^i(u)$ be the number of nodes in $N^i(u)$ that continue to be neighbors of u in H. Then $deg_H(u) = \sum_{i=1}^{t} deg_H^i(u)$ and by linearity of expectation $E[deg_H(u)] = \sum_{i=1}^{t} E[deg_H^i(u)]$. We will now show that

$$E[deg_H^i(u)] = \Theta(\log deg_G^i(u)) = O(\log deg_G(u)).$$

Since $t \leq 5$, we have that $E[deg_H(u)] = O(\log deg_G(u))$.

Fix i, $1 \leq i \leq t$, and let $d = deg_G^i(u)$ and $N^i(u) = \{v_1, v_2, \ldots, v_d\}$ such that $v_1 \prec_u v_2 \prec_u \cdots \prec_u v_d$. For each $j = 1, 2, \ldots, d$, let X_j denote the indicator random variable that equals 1 if $\{u, v_j\} \in E(H)$ and 0 otherwise. Then, by linearity of expectation, $E[deg_H^i(u)] = \sum_{j=1}^{d} E[X_j] = \sum_{j=1}^{d} \text{Prob}[X_j = 1]$. Given that $v_\ell \prec_u v_j$ for each $\ell = 1, 2, \ldots, j-1$, for the edge $\{u, v_j\}$ to be present in H, it must be the case that $u \prec_{v_j} v_\ell$, for each $\ell = 1, 2, \ldots, j-1$. Therefore,

$$\text{Prob}[X_j = 1] \leq \text{Prob}[\wedge_{\ell=1}^{j-1} u \prec_{v_j} v_\ell].$$

By Lemma 1, $\text{Prob}[\wedge_{\ell=1}^{j-1} u \prec_{v_j} v_\ell] = 1/j$. Therefore,

$$E[deg_H^i(u)] = \sum_{j=1}^{d} \text{Prob}[X_j = 1] \leq \sum_{j=1}^{d} \frac{1}{j} = \Theta(\log d) = \Theta(\log deg_G^i(u)).$$

This completes the proof. □

Now we prove the sparseness of $G_{\varepsilon\text{-}RTC}$. Specifically, we show that the expected degree of each vertex u in $G_{\varepsilon\text{-}RTC}$ is bounded above by a constant. Before we embark on this proof, we state a simple inequality that expresses a connection between the actual distance $|uv|$ between a pair of neighbors u and v and the eventual edge label $d_u[v]$ assigned by ε-RTC.

Lemma 2.

$$|uv| \cdot \frac{(1-\varepsilon)^2}{(1+\varepsilon)} \leq d_u[v] \leq |uv| \cdot \frac{(1+\varepsilon)^2}{(1-\varepsilon)}.$$

Theorem 4. *Let G be a UDG and suppose that the pairwise distance estimates $\alpha : V \times V \to \Re^+$ are ε-error bounded. Let $H = G_{\varepsilon\text{-}RTC}$. Then there is a constant C such that for any vertex $u \in V$, $E[deg_H(u)] \leq C$.*

Proof. Set $\eta = (1+\varepsilon)^3/(1-\varepsilon)^3$. For this proof to go through, we require that $\eta < 2$. This happens whenever

$$\varepsilon < \frac{(2^{1/3} - 1)}{(2^{1/3} + 1)} = 0.115013....$$

For notational ease we denote the constant 0.115013... above by Ω. In other words, our proof goes through when ε is around 1/9 or smaller. A more complicated analysis can be used for larger values of ε between Ω and 1; we skip that for conciseness.

Set $\theta = \pi/2 - \sin^{-1}(\eta/2)$ and $t = \lceil \frac{2\pi}{\theta} \rceil$. Partition the unit disk centered at u into cones, $C_1, C_2, \ldots, C_t$ such that each cone C_i, $1 \leq i < t$, makes an angle θ at u, C_t makes an angle of at most θ at u. It is worth noting that as $\varepsilon \to 0$, we see that $\eta \to 1$ and therefore $\theta \to \pi/3$. Also, when $\varepsilon \to \Omega$ from below, we see that $\eta \to 2$ from below and $\theta \to 0$. In short, as ε increases, our cone partition becomes more fine.

Let $\{N^i(u) \mid i = 1, 2, \ldots, t\}$ be a partition of $N(u)$ into t subsets such that all nodes in $N^i(u)$ lie in cone C_i. Note that each $N^i(u)$ induces a clique in G because $\theta \leq \pi/3$ for all values of $\varepsilon < \Omega$. Let $deg^i_G(u) = |N^i(u)|$, for each $i = 1, 2, \ldots, t$ and let $deg^i_H(u)$ be the number of nodes in $N^i(u)$ that continue to be neighbors of u in H. Then $deg_H(u) = \sum_{i=1}^{t} deg^i_H(u)$ and by linearity of expectation $E[deg_H(u)] = \sum_{i=1}^{t} E[deg^i_H(u)]$. We will now show that $E[deg^i_H(u)] \leq c$, where c is a constant in the sense that it is independent of the size of the network, but does depend on ε. From this, it immediately follows that $E[deg_H(u)] \leq t \cdot c$. Note that t is also independent of the size of network and depends only on ε. Hence, we have that $E[deg_H(u)] \leq C$, for some constant C.

Fix i, $1 \leq i \leq t$, and let $d = deg^i_G(u)$ and $N^i(u) = \{v_1, v_2, \ldots, v_d\}$ such that $v_1 \prec_u v_2 \prec_u \cdots \prec_u v_d$. For each $j = 1, 2, \ldots, d$, let X_j denote the indicator random variable that equals 1 if $\{u, v_j\} \in E(H)$ and 0 otherwise. Then, by linearity of expectation, $E[deg^i_H(u)] = \sum_{j=1}^{d} E[X_j] = \sum_{j=1}^{d} \text{Prob}[X_j = 1]$. We now calculate $\text{Prob}[X_j = 1]$.

Given that $v_\ell \prec_u v_j$ for each $\ell = 1, 2, \ldots, j-1$, for the edge $\{u, v_j\}$ to be present in H, it must be the case that $u \prec_{v_j} v_\ell$, for each $\ell = 1, 2, \ldots, j-1$. Therefore,

$$\begin{aligned}
\text{Prob}[X_j = 1] &\leq \text{Prob}[\wedge_{\ell=1}^{j-1} u \prec_{v_j} v_\ell], \\
&= \text{Prob}[\wedge_{\ell=1}^{j-1} d_{v_j}[u] < d_{v_j}[v_\ell]], \\
&= \prod_{\ell=1}^{j-1} \text{Prob}[d_{v_j}[u] < d_{v_j}[v_\ell]]. \qquad (2)
\end{aligned}$$

The last equality follows from the mutual independence of the events $\{d_{v_j}[u] < d_{v_j}[v_\ell] \mid \ell = 1, 2, \ldots, j-1\}$. These events are mutually independent because in ε-RTC each edge label $d_u[v]$ is obtained by a random perturbation and these are done *independently*. We will now compute an upper bound on $\text{Prob}[d_{v_j}[u] < d_{v_j}[v_\ell]]$, where $1 \leq \ell \leq j-1$. Fix ℓ and for notational ease let $v \equiv v_\ell$. We now claim a geometric property that follows from the fact that $d_{v_j}[v] < d_{v_j}[u]$ and from our choice of θ. The proof appears in the appendix (Section B) due to lack of space.

Claim: Given that $\theta = \pi/2 - \sin^{-1}(\eta/2)$ and that $d_{v_j}[v] < d_{v_j}[u]$, it follows that $|v_j v| \leq |v_j u|$.

Let $b = |v_j v|$. Then, both $\alpha(v_j, v)$ and $\alpha(v, v_j)$ are bounded above by $b \cdot (1+\varepsilon)$ and therefore their mean, which we will denote by $\overline{\alpha}\{v_j, v\}$, is also bounded above by $b \cdot (1+\varepsilon)$. Since $b = |v_j v| \leq |v_j u|$, we obtain in a similar manner that $\overline{\alpha}\{v_j, u\}$, the mean of $\alpha(v_j, u)$ and $\alpha(u, v_j)$ is bounded below by $b \cdot (1-\varepsilon)$.

Now recall that $d_{v_j}[u]$ is chosen uniformly at random from the interval $[L_u, R_u]$, where $L_u = \overline{\alpha}\{v_j, u\} \cdot (1-\delta_L)$ and $R_u = \overline{\alpha}\{v_j, u\} \cdot (1+\delta_R)$. Note that R_u satisfies

$$R_u = \overline{\alpha}\{v_j, u\} \cdot (1+\delta_R) \geq b \cdot (1-\varepsilon) \cdot (1+\delta_R) = b \cdot (1+\varepsilon).$$

Thus, as shown in Figure 4, the right endpoint R_u lies to the right of $\overline{\alpha}\{v_j, v\}$. Similarly, recall that $d_{v_j}[v]$ is chosen uniformly at random from the interval $[L_v, R_v]$, where $L_v =$

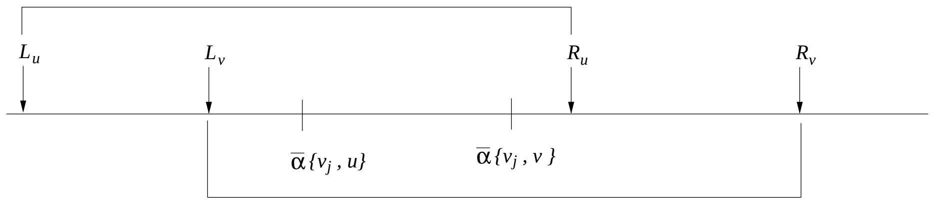

Fig. 4. This is the configuration of the point $\overline{\alpha}\{v_j, u\}$, the interval $[L_u, R_u]$ around it, the point $\overline{\alpha}\{v_j, v\}$, and the interval $[L_v, R_v]$ around it. Note that L_v is to the left of $\overline{\alpha}\{v_j, u\}$ and R_u is to the right of $\overline{\alpha}\{v_j, v\}$.

$\overline{\alpha}\{v_j, v\} \cdot (1 - \delta_L)$ and $R_v = \overline{\alpha}\{v_j, v\} \cdot (1 + \delta_R)$. Now note that L_v satisfies:

$$L_v = \overline{\alpha}\{v_j, v\} \cdot (1 - \delta_L) \leq b \cdot (1 + \varepsilon) \cdot (1 - \delta_L) \leq b \cdot (1 - \varepsilon).$$

Thus, as shown in Figure 4, the left endpoint L_v lies to the left of $\overline{\alpha}\{v_j, u\}$.

We now show that the interval $[L_v, R_u]$ is fairly large, compared to both $[L_u, R_u]$ and $[L_v, R_v]$. Recalling that $R_u \geq b \cdot (1 + \varepsilon)$ and $L_v \leq b \cdot (1 - \varepsilon)$, we see that $R_u - L_v \geq b \cdot (2\varepsilon)$. Furthermore, we have that

$$R_v - L_v = \overline{\alpha}\{v_j, v\} \cdot (1 + \delta_R) - \overline{\alpha}\{v_j, v\} \cdot (1 - \delta_L) \leq b \cdot (1 + \varepsilon) \cdot (\delta_R + \delta_L) = b \cdot (2\varepsilon) \cdot \frac{2}{(1 - \varepsilon)}.$$

Therefore,

$$\frac{R_u - L_v}{R_v - L_v} \geq \frac{1 - \varepsilon}{2} \qquad \text{and} \qquad \frac{R_u - L_v}{R_u - L_u} \geq \frac{1 - \varepsilon}{2}.$$

The latter follows from the fact that $R_u - L_u \leq R_v - L_v$. From these bounds, it follows that

$$\text{Prob}\Big[d_{v_j}[u] \in [L_v, R_u] \text{ and } d_{v_j}[v_\ell] \in [L_v, R_u]\Big] \geq \frac{1 - \varepsilon}{4}.$$

Given that both $d_{v_j}[u]$ and $d_{v_j}[v]$ are in $[L_v, R_u]$, either the event $d_{v_j}[u] < d_{v_j}[v]$ or the event $d_{v_j}[u] > d_{v_j}[v]$ occurs. Because of symmetry, the likelihood of these two possibilities is the same and therefore,

$$\text{Prob}\Big[d_{v_j}[u] > d_{v_j}[v]\Big] \geq \frac{1 - \varepsilon}{8},$$

implying that

$$\text{Prob}\Big[d_{v_j}[u] < d_{v_j}[v]\Big] \leq \frac{7 + \varepsilon}{8}.$$

Plugging this upper bound in (2), we get that $\text{Prob}[X_j = 1] \leq (\frac{7+\varepsilon}{8})^{j-1}$. Therefore,

$$E[deg_H^i(u)] \leq \sum_{j=1}^{d} \Big(\frac{7 + \varepsilon}{8}\Big)^{j-1} \leq \frac{8}{1 - \varepsilon}.$$

□

4.2 Experimental evidence for sparseness

In this subsection we report on experiment results related to the maximum degree of G_{RTC} and $G_{\varepsilon\text{-}RTC}$. Figure 5 shows the ratio $\Delta(G_{RTC})/\log\Delta(G)$ as the size of G increases. In the plot on the left, the increase in the size of G is due to an increase in the area in which the nodes of G are distributed. In this case, the density of G remains fixed. In the plot on the right, the size of G is increased by increasing the density of the graph, while keeping the area of the graph fixed.

Theorem 3 claims that for any vertex u, the expected value of $\deg_{G_{RTC}}(u)$ is bounded above by $O(\log deg_G(u))$. This does not imply that $E[\Delta(G_{RTC})] = O(log\Delta(G))$. To see this consider a graph G in which there are many nodes whose degrees are equal to $\Delta(G)$. While it is true that for each of these high degree nodes u, we would *expect* the degree of u in G_{RTC} to be equal to $O(\log deg_G(u))$, it is also true that with significant probability at least one of these nodes will have a degree in the output graph that is much larger than $\log deg_G(u)$. As a result, $\Delta(G_{RTC})$ may drift above $O(\log\Delta(G))$ and this drift is more pronounced as the number of high degree nodes increases. This phenomenon is illustrated by the two plots. In the plot on the left the ratio $\Delta(G_{RTC})/\log\Delta(G)$ tends to a constant because even though the size of G is increasing, the density is not. The plot on the right shows a steady increase because as the density of G increases, there are more and more high degree vertices.

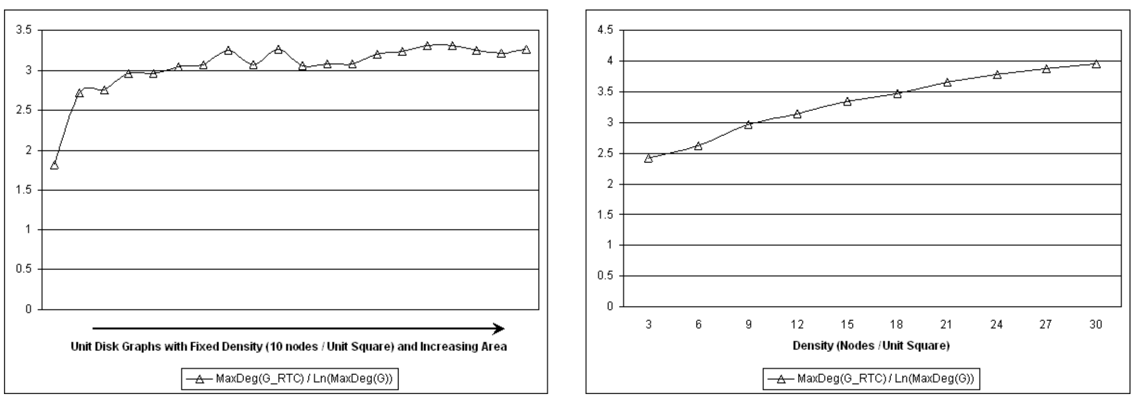

Fig. 5. Although maximum degree of G_{RTC} is not bounded, it is comparable to $log(\Delta)$.

Below (Figure 6) we show a 3-dimensional plot of the behavior of $\Delta(G_{\varepsilon\text{-}RTC})$ with respect to varying values of ε and the density of the input UDG. For $\varepsilon = 0$ the plot shows that $\Delta(G_{\varepsilon\text{-}RTC})$ is bounded above by a constant. This is to be expected because ε-RTC is the same as XTC for $\varepsilon = 0$. For large values of ε (say, 0.8 or more) the behavior of $\Delta(G_{\varepsilon\text{-}RTC})$ is similar to the behavior of $\Delta(G_{RTC})$. This is also to be expected. What is more interesting is the behavior of $\Delta(G_{RTC})$ for small, positive values of ε. Our plot shows that for small, positive ε, the value of $\Delta(G_{\varepsilon\text{-}RTC})$ remains, more or less, a constant even though the density of G increases. Theorem 4 claims that for any node u, the degree of u in $G_{\varepsilon\text{-}RTC}$ is bounded above by a constant C (whose value depends on ε). While this implies that the average degree of $G_{\varepsilon\text{-}RTC}$ is bounded above by $C/2$, it does not imply that $\Delta(G_{\varepsilon\text{-}RTC})$ is bounded above by a constant. Hence, it is a pleasant surprise to see $\Delta(G_{\varepsilon\text{-}RTC})$ so well behaved for small values of ε.

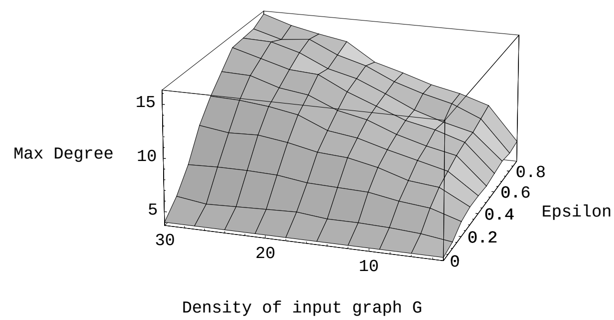

Fig. 6. The behavior of $\Delta(G_{\varepsilon\text{-}RTC})$, as ε increases and as the density of G increases.

4.3 Spanner properties

Due to lack of space we do not present a detailed analysis of the spanner properties of G_{RTC} and $G_{\varepsilon\text{-}RTC}$ here. We simply state our results without proof and postpone the proofs and a discussion of the implications of these results to the full version of the paper. Note that our results hold for arbitrary graphs and not just UDGs.

Definition: Let H be a spanning subgraph of the graph G. H is said to be a *t-hop spanner* of G if for any u and v in G, $c_H(u,v)/c_G(u,v) \leq t$, where $c_G(u,v)$ (respectively, $c_H(u,v)$) is the number of hops in a shortest u,v-path in G (respectively, H).

Theorem 5. *Let G be an arbitrary graph. G_{RTC} is a t-spanner of G for $t = O(1)$ and $G_{\varepsilon\text{-}RTC}$ is a t-spanner of G for $t = \log(1/\delta)$, where δ is the smallest distance $|uv|$ between any pair of nodes in G.*

References

1. L. Barriére, P. Fraigniaud, and L. Narayanan. Robust position-based routing in wireless ad hoc networks with unstable transmission ranges. In *Proceedings of the 5th international workshop on Discrete algorithms and methods for mobile computing and communications (DIALM)*, pages 19–27, 2001.
2. N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Wireless Communications*, 7(5):27–34, 2000.
3. S. Ghosh, K. Lillis, S. Pandit, and S. Pemmaraju. Robust topology control algorithms. In *Proceedings of the 8th International Conference on Principles of Distributed Systems (OPODIS)*, 2004.
4. 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 *Twentieth ACM Symposium on Principles of Distributed Computing (PODC)*, 2001.
5. X.-Y. Li, G. Calinescu, and P. Wan. Distributed construction of planar spanner and routing for ad hoc wireless networks. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.
6. X.-Y. Li, G. Calinescu, P.-J. Wan, and Y. Wang. Localized delaunay triangulation with application in ad hoc wireless networks. *IEEE Trans. Parallel Distrib. Syst.*, 14(10):1035–1047, 2003.
7. Xiang-Yang Li and Yu Wang. Efficient construction of low weight bounded degree planar spanner. In *Proceedings of the 9th International Computing and Combinatorics Conference (COCOON)*, 2003.
8. K. Lillis and S. Pemmaraju. Topology control with limited geometric information. Full paper: `www.cs.uiowa.edu/~sriram/randomTopControl.pdf`, 2005.
9. N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Mobile Computing and Networking*, pages 32–43, 2000.
10. W.-Z. Song, Y. Wang, X.-Y. Li, and O. Frieder. Localized algorithms for energy efficient topology in wireless ad hoc networks. In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 98–108. ACM Press, 2004.
11. R. Wattenhofer and A. Zollinger. XTC: A practical topology control algorithm for ad-hoc networks. In *4th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN)*, 2004.

APPENDIX

A Experimental Setup

We have implemented XTC, RTC, and ε-RTC on a Pentium 4 PC running Mathematica 5.0, release 1, with the Combinatorica add-on package.

All algorithms were run on UDGs and quasi UDGs that were generated by distributing nodes in the plane in a variety of ways. One series of input graphs was constructed by distributing nodes uniformly at random, at a density of 10 nodes per unit square, and increasing the area of the square linearly from 1 to 77 square units. Another series of input graphs was constructed by fixing the dimensions of the graph at 5×5 and increasing the density of nodes, placed uniformly at random, linearly form 3 to 30 nodes per unit square. A final series of input graphs was constructed by distributing nodes uniformly at random on a 5×5 square, at a density of 3 nodes per unit square, and then adding small regions $(\frac{1}{2} \times \frac{1}{2})$ with node density of 30 nodes per unit square.

The paper reports on a very small subset of our experimental results. The complete results can be viewed at `web.sau.edu/lilliskevinm/uofi/rtc/`.

B Proof of geometric claim

Claim: Given that $\theta = \pi/2 - \sin^{-1}(\eta/2)$ and that $d_{v_j}[v] < d_{v_j}[u]$, it follows that $|v_j v| \leq |v_j u|$.

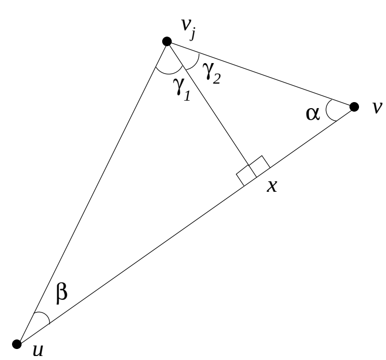

Fig. 7. This shows the triangle uv_jv.

Proof. Refer to Figure 7. To show that $|v_j v| \leq |v_j u|$, we will show that $\beta \leq \alpha$. This is equivalent to showing that $\gamma_1 \geq \gamma_2$. To obtain a contradiction suppose that $\gamma_1 < \gamma_2$. This implies that $|ux| < |xv|$.

Since $v \prec_u v_j$, we have that $d_u[v] < d_u[v_j]$. Combining this with Lemma 2 we obtain

$$|uv| \cdot \frac{(1-\varepsilon)^2}{(1+\varepsilon)} \leq d_u[v] < d_u[v_j] \leq |uv_j| \cdot \frac{(1+\varepsilon)^2}{(1-\varepsilon)}.$$

This yields

$$|uv| < \frac{(1+\varepsilon)^3}{(1-\varepsilon)^3} \cdot |uv_j| = \eta \cdot |uv_j|.$$

Combining the inequalities $|ux| < |xv|$, $|uv| < \eta \cdot |uv_j|$ with the fact that $|ux| + |xv| = |uv|$, we get that $|ux| < \eta/2 \cdot |uv_j|$. Now notice that $\sin \gamma_1 = |ux|/|uv_j| < \eta/2$, implying that $\gamma_1 < \sin^{-1}(\eta/2)$. Using the fact that $\beta \leq \theta$ and that θ was chosen to satisfy $\theta = \pi/2 - \sin^{-1}(\eta/2)$, we get that $\beta < \pi/2 - \gamma_1$. In other words, $\beta + \gamma_1 < \pi/2$, which is a contradiction. $\square$