Does Topology Control Reduce Interference?*

Martin Burkhart, Pascal von Rickenbach, Roger Wattenhofer, Aaron Zollinger
Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
{bumartin, pascalv}@student.ethz.ch, {wattenhofer, zollinger}@inf.ethz.ch

ABSTRACT

Topology control in ad-hoc networks tries to lower node energy consumption by reducing transmission power and by confining interference, collisions and consequently retransmissions. Commonly low interference is claimed to be a consequence to sparseness of the resulting topology. In this paper we disprove this implication. In contrast to most of the related work—claiming to solve the interference issue by graph sparseness without providing clear argumentation or proofs—, we provide a concise and intuitive definition of interference. Based on this definition we show that most currently proposed topology control algorithms do not effectively constrain interference. Furthermore we propose connectivity-preserving and spanner constructions that are interference-minimal.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*wireless communication*;
F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*

General Terms

Algorithms, Theory

Keywords

Ad-hoc networks, interference, network connectivity, network spanners, topology control

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

MobiHoc'04, May 24–26, 2004, Roppongi, Japan.

1. INTRODUCTION

In mobile wireless ad-hoc networks—formed by autonomous devices communicating by radio—energy is one of the most critical resources. The main goal of *topology control* is to reduce node power consumption in order to extend network lifetime. Since the energy required to transmit a message increases at least quadratically with distance, it makes sense to replace a long link by a sequence of short links. On the one hand, energy can therefore be conserved by abandoning energy-expensive long-range connections, thereby allowing the nodes to reduce their transmission power levels. On the other hand, reducing transmission power also confines interference, which in turn lowers node energy consumption by reducing the number of collisions and consequently packet retransmissions on the media access layer.[1] Dropping communication links however clearly takes place at the cost of network connectivity: If too many edges are abandoned, connecting paths can grow unacceptably long or the network can even become completely disconnected. As illustrated in Figure 1, topology control can therefore be considered a trade-off between energy conservation and interference reduction on the one hand and connectivity on the other hand.

Figure 1: Topology control constitutes a trade-off between node energy conservation and network connectivity.

The interference aspect is often maintained by developers of topology control algorithms to be solved by sparseness or low node degree of the resulting topology graph, without providing rigorous motivation or proofs. The foremost contribution of this paper is to disprove this assertion.

In contrast to most of the related work—where the interference issue is seemingly solved by sparseness arguments—, we start out by precisely defining our notion of interference. The definition of interference is based on the natural question, how many nodes are affected by communication over a

[1]Sometimes also the construction of node clusters and dominating sets of nodes is considered topology control. In this paper we restrict ourselves to the study of topology control based on transmission power reduction.

certain link. By prohibiting specific network edges, the potential for communication over high-interference links can then be confined.

We employ this interference definition to formulate the trade-off between energy conservation and network connectivity. In particular we state certain requirements that need to be met by the resulting topology. Among these requirements are connectivity (if two nodes are—possibly indirectly—connected in the given network, they should also be connected in the resulting topology) and the spanner property (the shortest path between any pair of nodes on the resulting topology should be longer at most by a constant factor than the shortest path connecting the same pair of nodes in the given network). After stating such requirements, an optimization problem can be formulated to find the topology meeting the given requirements with minimum interference.

For the requirement that the resulting topology should retain connectivity of the given network, we show that most currently proposed topology control algorithms—already by having every node connect to its nearest neighbor—commit a substantial mistake: Although certain proposed topologies are guaranteed to have low degree yielding a sparse graph, interference becomes asymptotically incomparable with the interference-minimal topology. We also show that there exist graphs for which no local algorithm can approximate the optimum. With respect to the sometimes desirable requirement that the resulting topology should be planar, we show that planarity can increase interference.

Furthermore we propose a centralized algorithm (LIFE) that computes an interference-minimal connectivity-preserving topology. For the requirement that the resulting topology should be a spanner with a given stretch factor, we present (based on a centralized variant of the algorithm) a distributed local algorithm (LLISE) that computes a provably interference-optimal spanner topology.

Our results are not confined to worst-case considerations; we also show by simulation that on average-case graphs traditional topology control algorithms—in particular the Gabriel Graph and the Relative Neighborhood Graph—fail to effectively reduce interference. Moreover these constructions are shown to be outperformed by the LLISE algorithm, which therefore proves to be average-case effective in addition to its worst-case optimality.

After discussing related work in the following section, we state the model for this paper in Section 3. Focusing on the drawbacks of currently proposed topology control algorithms with respect to interference in Section 4, we present interference-optimal algorithms in the subsequent section. Section 6 assesses our algorithms as well as previously proposed topologies regarding their interference on average-case graphs and the subsequent section concludes the paper.

2. RELATED WORK

In this section we discuss related work in the field of topology control with special focus on the issue of interference.

2.1 Topology Control

The assumption that nodes are distributed randomly in the plane according to a uniform probability distribution formed the basis of pioneering work in the field of topology control [7, 21].

Later proposals adopted constructions originally studied in computational geometry, such as the Delaunay Triangulation [8], the minimum spanning tree [19], the Relative Neighborhood Graph [10], or the Gabriel Graph [20]. Most of these contributions mainly considered energy-efficiency of paths preserved by the resulting topology, whereas others exploited the planarity property of the proposed constructions for geometric routing [4, 13].

The Delaunay Triangulation and the minimum spanning tree not being computable locally and thus not being practicable, a next generation of topology control algorithms emphasized locality. The CBTC algorithm [24] was the first construction to focus on several desired properties, in particular being an energy spanner with bounded degree. This process of developing local algorithms featuring more and more properties was continued partly based on CBTC, partly based on local versions of classic geometric constructions such as the Delaunay Triangulation [15] or the minimum spanning tree [14]. One of the most recent such results is a locally computable planar distance (and energy) spanner with constant-bounded node degree [23]. Another thread of research takes up the average-graph perspective of early work in the field; [3] for instance shows that the simple algorithm choosing the k nearest neighbors works surprisingly well on such graphs.

Yet another aspect of topology control is considered by algorithms trying to form clusters of nodes. Most of these proposals are based on (connected) dominating sets [1, 2, 9] and focus on locality and provable properties, such as [12], which achieves a non-trivial approximation of the minimum dominating set in constant time. Cluster-based constructions are commonly regarded a variant of topology control in the sense that energy-consuming tasks can be shared among the members of a cluster.

Topology control having so far mainly been of interest to theoreticians, first promising steps are being made towards exploiting the benefit of such techniques also in practical networks [11].

2.2 Interference

As mentioned earlier, reducing interference—and its energy-saving effects on the medium access layer—is one of the main goals of topology control besides direct energy conservation by restriction of transmission power. Astonishingly however, all the above topology control algorithms at the most implicitly try to reduce interference. Where interference is mentioned as an issue at all, it is maintained to be confined at a low level as a consequence to sparseness or low degree of the resulting topology graph.

A notable exception to this is [16] defining an explicit notion of interference. Based on this interference model between edges, a time-step routing model and a concept of congestion is introduced. It is shown that there are inevitable trade-offs between congestion, power consumption and dilation. For some node sets, congestion and energy are even shown to be incompatible.

The interference model proposed in [16] is based on current network traffic. The amount and nature of network traffic however is highly dependent on the chosen application. Since usually no a priori information about the traffic in a network is available, a static model of interference depending solely on a node set is consequently desirable.

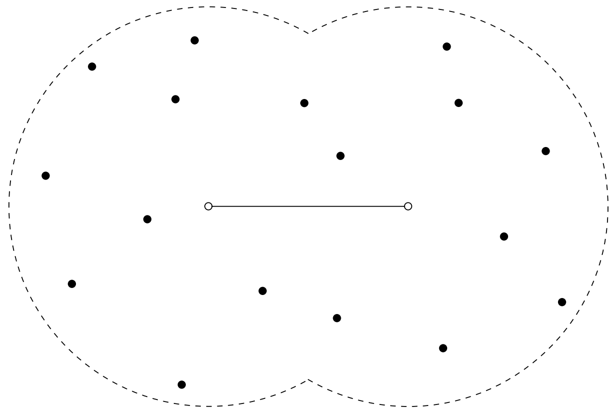

Figure 2: Nodes covered by a communication link.

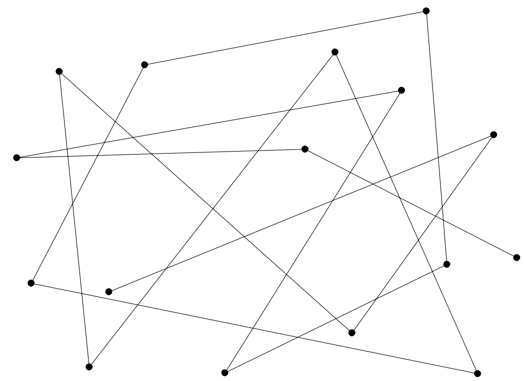

Figure 3: Low degree does not guarantee low interference.

3. MODEL

Mobile ad-hoc networks are commonly modeled by graphs. A graph $G = (V, E)$ consists of a set of nodes $V \subset \mathbb{R}^2$ in the Euclidean plane and a set of edges $E \subseteq V^2$. Nodes represent mobile hosts, whereas edges represent links between nodes. In order to prevent already basic communication between directly neighboring nodes from becoming unacceptably cumbersome [18], it is required that a message sent over a link can be acknowledged by sending a corresponding message over the same link in the opposite direction. In other words, only *undirected* (symmetric) edges are considered.

We assume that a node can adjust its transmission power to any value between zero and its maximum power level. The maximum power levels are not assumed to be equal for all nodes. An edge (u, v) may exist only if both incident nodes are capable of sending a message over (u, v), in particular if the maximum transmission radius of both u and v is at least $|u, v|$, their Euclidean distance. A pair of nodes u, v is considered *connectable in the given network* if there exists a path connecting u and v provided that all transmission radii are set to their respective maximum values. The task of a *topology control* algorithm is then to compute a subgraph of the given network graph with certain properties, reducing the transmission power levels and thereby attempting to reduce interference and energy consumption.

With a chosen transmission radius—for instance to reach a node v—a node u affects at least all nodes located within the circle centered at u and with radius $|u, v|$. $D(u, r)$ denoting the disk centered at node u with radius r and requiring edge symmetry, we consequently define the *coverage* of an (undirected) edge $e = (u, v)$ to be the cardinality of the set of nodes covered by the disks induced by u and v:

$$\begin{aligned} Cov(e) :=& \big| \{w \in V | w \text{ is covered by } D(u, |u, v|)\} \cup \\ & \{w \in V | w \text{ is covered by } D(v, |v, u|)\} \big|. \end{aligned}$$

In other words the coverage $Cov(e)$ represents the number of network nodes affected by nodes u and v communicating with their transmission powers chosen such that they exactly reach each other (cf. Figure 2).

The edge level interference defined so far is now extended to a graph interference measure as the maximum coverage occurring in a graph:

DEFINITION 1. *The interference of a graph G=(V,E) is defined as*

$$I(G) := \max_{e \in E} Cov(e).$$

Since interference reduction per se would be senseless (if all nodes simply set their transmission power to zero, interference will be reduced to a minimum), the formulation of additional requirements to be met by a resulting topology is necessary. A resulting topology can for instance be required

- to maintain connectivity of the given communication graph (if a pair of nodes is connectable in the given network, it should also be connected in the resulting topology graph),
- to be a spanner of the underlying graph (the shortest path connecting a pair of nodes u, v on the resulting topology is longer by a constant factor only than the shortest path between u and v on the given network), or
- to be planar (no two edges in the resulting graph intersect).

Finding a resulting topology which meets one or a combination of such requirements with minimum interference constitutes an optimization problem.

4. INTERFERENCE IN KNOWN TOPOLOGIES

It is often argued that sparse topologies with small or bounded degree are well suited to minimize interference. In this section we show that low degree does not necessarily imply low interference. Moreover we demonstrate that most currently known topology control algorithms can perform badly compared to the interference optimum, that is a topology which minimizes interference in the first place.

In particular we consider in this section the basic problem of constructing an interference-minimal topology maintaining connectivity of the given network.

The following basic observation states that—although often maintained—low degree alone does not guarantee low interference. Figure 3 for instance shows a topology graph with degree 2 whose interference is however roughly n, the number of network nodes. A node can interfere with other

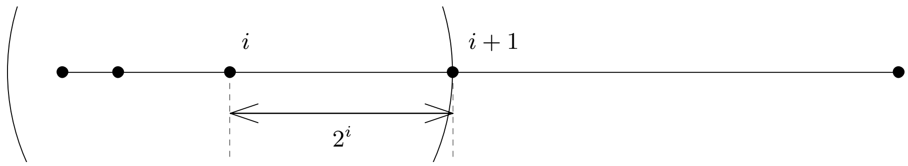

Figure 4: Exponential node chain with interference $\Omega(n)$.

nodes that are not direct neighbors in the chosen topology graph. Whereas the maximum degree of the underlying communication graph of the given network (with all nodes transmitting at full power) is an upper bound for interference, the degree of a resulting topology graph is only a lower bound.

There exist instances where also the optimum exhibits interference $\Omega(n)$, for instance a chain of nodes with exponentially growing distances (cf. Figure 4, proposed in [16]), whose large interference is caused as a consequence to the requirement that the resulting topology is to be connected. Every node u_i (except for the leftmost) is required to have an incident edge to the left, which covers all nodes left of u_i. Assessing the interference quality of a topology control algorithm, its interference on a given network therefore needs to be compared to the optimum interference topology for the same network.

To the best of our knowledge, all currently known topology control algorithms constructing only symmetric connections have in common that every node establishes a symmetric connection to at least its nearest neighbor. In other words all these topologies contain the Nearest Neighbor Forest constructed on the given network. In the following we show that by including the Nearest Neighbor Forest as a subgraph, the interference of a resulting topology can become incomparably bad with respect to a topology with optimum interference.

Theorem 1. *No currently proposed topology control algorithm establishing only symmetric connections—required to maintain connectivity of the given network—is guaranteed to yield a nontrivial interference approximation of the optimum solution. In particular, interference of any proposed topology can be $\Omega(n)$ times larger than the interference of the optimum connected topology, where n is the total number of network nodes.*

Proof. Figure 5 depicts an extension of the example graph shown in Figure 4. In addition to a horizontal exponential node chain, each of these nodes h_i has a corresponding node v_i vertically displaced by a little more than h_i's distance to its left neighbor. Denoting this vertical distance d_i, $d_i > 2^{i-1}$ holds. These additional nodes form a second (diagonal) exponential line. Between two of these diagonal nodes v_{i-1} and v_i, an additional helper node t_i is placed such that $|h_i, t_i| > |h_i, v_i|$.

The Nearest Neighbor Forest for this given network (with the additional assumption that the transmission radius of each node can be chosen sufficiently large) is shown in Figure 6. Roughly one third of all nodes being part of the horizontally connected exponential chain, interference of any topology containing the Nearest Neighbor Forest amounts to at least $\Omega(n)$. An interference-optimal topology, however, would connect the nodes as depicted in Figure 7 with constant interference. □

In other words, already by having each node connect to the nearest neighbor, a topology control algorithm makes an "irrevocable" error. Moreover, it commits an asymptotically worst possible error since the interference in any network cannot become larger than n.

Since roughly one third of all nodes are part of the horizontal exponential node chain in Figure 5, the observation stated in Theorem 1 would also hold for an average interference measure, averaging interference over all edges.

The following theorem even shows that for connectivity-preserving topologies no local algorithm can approximate optimum interference for every given network. Thereby the definition of a distributed *local* algorithm assumes that each network node is informed about its network neighborhood only up to a given constant distance.

Theorem 2. *For the requirement of maintaining connectivity of the given network, there exists a class of graphs for which there is no local algorithm that approximates optimum interference.*

Proof. In Figure 8 the maximum transmission radius of a node is $|u, v|$. Let n be the number of nodes in the graph. Then the shaded area contains $\Omega(n)$ evenly distributed nodes which can be connected with constant interference. For each such node i the inequalities $|i, v| < |u, v|$ and $|u, i| > |u, v|$ hold. It follows that edge (u, v) has $\Omega(n)$ interference since it covers all nodes in the shaded area. In addition there is a chain of nodes (dashed path) connecting node u with node v indirectly through the nodes located in the shaded area. The nodes in the chain are located in such a way that it is possible to connect them with constant interference. For such a graph $O(1)$ interference can be achieved by connecting u to the rest of the graph through the chain of nodes and not directly through edge (u, v), which would cause $\Omega(n)$ interference.

A local algorithm at node u has to decide if it can drop edge (u, v) or not. This is only possible if u knows about the existence of an alternative path from u to v in order to maintain connectivity. By elongating the chain sufficiently, the local algorithm can thus be forced to include edge (u, v), pushing up interference to $O(n)$ whereas the optimum is $\Omega(1)$. □

As mentioned in Section 3, another popular requirement for topology control algorithms besides bounded degree is planarity of the resulting topology, meaning that no two edges of the resulting graph intersecty. This is often desired because numerous well-understood routing algorithms exist

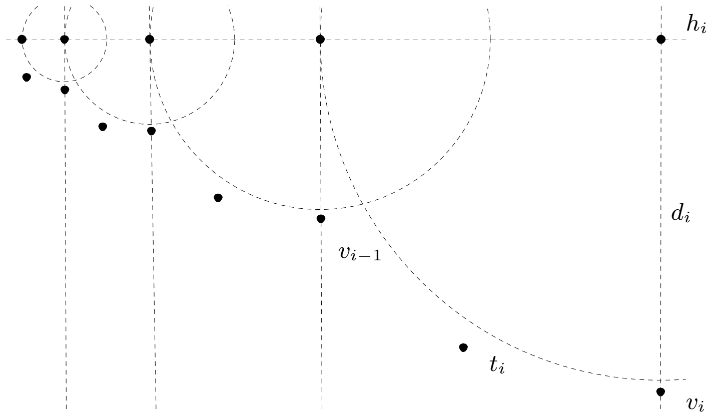

Figure 5: Two exponential node chains.

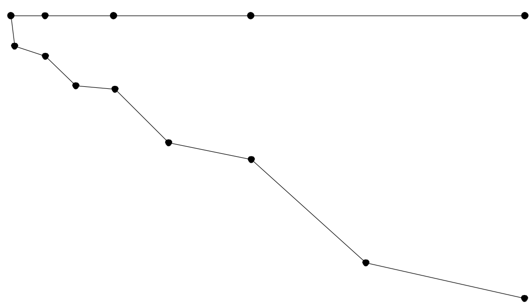

Figure 6: The Nearest Neighbor Forest yields interference $\Omega(n)$.

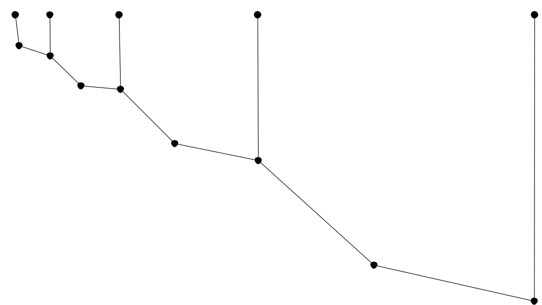

Figure 7: Optimal tree with constant interference.

that are only applicable to planar graphs. But topology control algorithms enforcing planarity are not optimal in terms of interference:

THEOREM 3. *There exist graphs on which interference-optimal topologies—required to maintain connectivity—are not planar.*

PROOF. In Figure 9 the maximum transmission radius of a node is $|a, b|$. All eligible edges are depicted together with the coverage area for edges whose incident nodes are both in $\{a, b, c, d\}$. The indicated weight of an edge e corresponds to its coverage $Cov(e)$. V and W represent sets of 3 and 4 nodes, respectively. The nodes in set V and W, respectively, can be connected among themselves with interference 3. A topology control algorithm can only reduce interference by removing all edges with maximum interference (here (a, c) and (b, c)) from the graph. Thereafter, no further edge can be removed without breaking connectivity since the graph without (a, c) and (b, c) is a tree. Thus the resulting tree is interference-optimal and non-planar since both edges (a, b) and (c, d) must remain in the resulting topology. □

5. LOW-INTERFERENCE TOPOLOGIES

In this section we present three algorithms that explicitly reduce interference of a given network. The first algorithm is capable of finding an interference-optimal topology maintaining connectivity of the given network. The other two algorithms compute an interference-optimal topology under the additional requirement of being a spanner of the given network. Whereas the first spanner algorithm assumes global knowledge of the network, the second can be computed locally.

5.1 Interference-Optimal Spanning Forest

In the following we again require the resulting topology to maintain connectivity of the given network. A topology graph meeting this requirement can therefore consist of a tree for each connected component of the given network since additional edges might unnecessarily increase interference. A *Minimum Interference Forest* is therefore a set of trees maintaining the connectivity of the given network with least possible interference.

Algorithm LIFE computes a Minimum Interference Forest.

Low Interference Forest Establisher (LIFE)

Input: a set of nodes V, each $v \in V$ having attributed a maximum transmission radius r_v^{max}

```
E = all eligible edges (u, v) (r_u^max ≥ |u, v| and r_v^max ≥ |u, v|) (* unprocessed edges *)
E_LIFE = ∅
G_LIFE = (V, E_LIFE)
while E ≠ ∅ do
    e = (u, v) ∈ E with minimum coverage
    if u, v are not connected in G_LIFE then
        E_LIFE = E_LIFE ∪ {e}
    end if
    E = E \ {e}
end while
```

Output: Graph G_{LIFE}

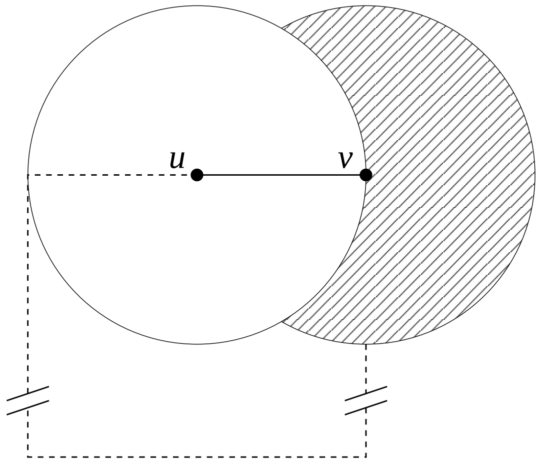

Figure 8: Worst case graph for which no local algorithm can approximate optimum interference.

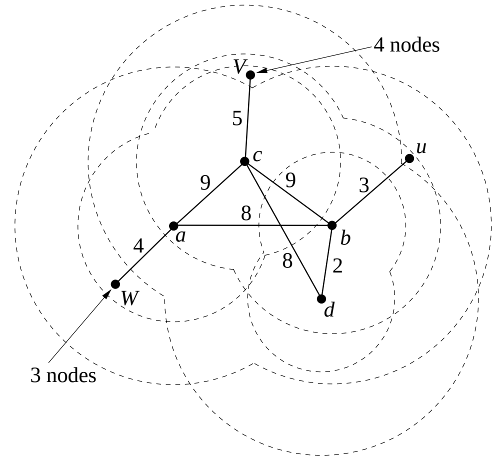

Figure 9: Node set whose interference-optimal topology is not planar.

THEOREM 4. *The forest constructed by LIFE is a Minimum Interference Forest.*

PROOF. The LIFE algorithm computes a minimum spanning forest (MSF) of the graph $G = (V, E)$, where E is the set of all eligible edges, if every edge $e \in E$ is attributed the weight $Cov(e)$. With its greedy strategy, it follows the lines of Kruskal's MSF algorithm [5]. To prove the theorem, it is therefore sufficient to show that the MSF is optimal with respect to interference. Optimality of LIFE then follows from the fact that a minimum spanning forest also minimizes the maximum edge weight in any spanning forest. (Assuming for contradiction that G^* is an MSF with maximum weight edge e^*, whereas G_{LIFE} is a spanning forest with lower maximum edge weight, e^* could be replaced by a corresponding edge from G_{LIFE}, yielding a spanning forest with total edge weight smaller than G^*'s, which contradicts the assumption that G^* is an MSF.) □

With an appropriate implementation of the connectivity query in Line 6 the running time of the algorithm LIFE is $O(n^2 \log n)$. If the given network is known to consist of one connected component only, Prim's minimum-spanning-tree algorithm can be employed with running time $O(n^2)$. Algorithms computing a minimum spanning tree in a distributed way—as particularly suitable for ad-hoc networks—are described in detail in [17].

5.2 Low-Interference Spanners

LIFE optimizes interference for the requirement that the resulting topology has to maintain connectivity. In addition to connectivity it is often desired that the resulting topology should be a spanner of the given network. A formal definition of a t-spanner follows:

DEFINITION 2. *[t-Spanner] A t-spanner of a graph $G = (V, E)$ is a subgraph $G' = (V, E')$ such that for each pair (u, v) of nodes $|p^*_{G'}(u, v)| \le t \cdot |p^*_G(u, v)|$, where $|p^*_{G'}(u, v)|$ and $|p^*_G(u, v)|$ denote the length of the shortest path between u and v in G' and G, respectively.*

In this paper we consider Euclidean spanners, that is, the length of a path is defined as the sum of the Euclidean lengths of all its edges. With slight modifications our results are however also extendable to hop spanners, where the length of a path corresponds to the number of its edges.

Algorithm LISE is a topology control algorithm that constructs a t-spanner with optimum interference. LISE starts with a graph $G_{LISE} = (V, E_{LISE})$ where E_{LISE} is initially the empty set. It processes all eligible edges of the given network $G = (V, E)$ in descending order of their coverage. For each edge $(u, v) \in E$ not already in E_{LISE}, LISE computes a shortest path from u to v in G_{LISE} provided that the Euclidean length of this path is less than or equal to $t\,|u, v|$. As long as no such path exists, the algorithm keeps inserting all unprocessed eligible edges with minimum coverage into E_{LISE}.

To prove the interference optimality of G_{LISE}, we introduce an additional lemma, which shows that G_{LISE} contains all eligible edges whose coverage is less than $I(G_{LISE})$.

LEMMA 5. *The graph $G_{LISE} = (V, E_{LISE})$ constructed by LISE from a given network $G = (V, E)$ contains all edges e in E whose coverage $Cov(e)$ is less than $I(G_{LISE})$.*

Low Interference Spanner Establisher (LISE)

Input: a set of nodes V, each $v \in V$ having attributed a maximum transmission radius r_v^{max}

```
E = all eligible edges (u,v) (r_u^max ≥ |u,v| and r_v^max ≥ |u,v|) (* unprocessed edges *)
E_LISE = ∅
G_LISE = (V, E_LISE)
while E ≠ ∅ do
    e = (u,v) ∈ E with maximum coverage
    while |p*(u,v) in G_LISE| > t |u,v| do
        f = edge ∈ E with minimum coverage
        move all edges ∈ E with coverage Cov(f) to E_LISE
    end while
    E = E \ {e}
end while
```

Output: Graph G_{LISE}

PROOF. We assume for the sake of contradiction that there exists an edge e in E with $Cov(e) < I(G_{LISE})$ which is not contained in E_{LISE}. Consequently, LISE never takes an edge with coverage $Cov(e)$ in line 7 since the algorithm would insert all edges with $Cov(e)$ into E_{LISE} in line 8 instantly (thus also e). There exists however an edge f in E_{LISE} with $Cov(f) = I(G_{LISE})$ eventually taken in line 7. Therefore the inequality $Cov(e) < Cov(f)$ holds. At the time the algorithm takes f in line 7, all edges taken in line 5 must have had coverage greater than or equal to $Cov(f)$ since the maximum of an ordered set can only be greater than or equal to the minimum of the same set. Hence e has never been taken in line 5 and therefore has never been removed from E in line 10. Consequently, e is still in E when f is taken as the edge with minimum coverage in E. Thus it holds that $Cov(f) \leq Cov(e)$ which leads to a contradiction. □

With Lemma 5 we are ready to prove that the resulting topology constructed by LISE is an interference-optimal t-spanner.

THEOREM 6. *The graph $G_{LISE} = (V, E_{LISE})$ constructed by LISE from a given network $G = (V, E)$ is an interference-optimal t-spanner of G.*

PROOF. To show that G_{LISE} meets the spanner property, it is sufficient to prove that for each edge $(u, v) \in E$ there exists a path in G_{LISE} with length not greater than $t\,|u, v|$. This holds since for a shortest path $p^*(u, v)$ in G a path $p'(u, v)$ in G_{LISE} with $|p'| \leq t\,|p|$ can be constructed by substituting for each edge on p the corresponding spanner path in G_{LISE}. For edges in E which also occur in E_{LISE} the spanner property is trivially true. On the other hand an edge (u, v) can only be in E but not in E_{LISE} if a path from u to v in G_{LISE} with length not greater than $t\,|u, v|$ exists (see if-condition in line 6). Thus G_{LISE} is a t-spanner of G.

Interference optimality of LISE can be proved by contradiction. We therefor assume, that G_{LISE} is not an interference-optimal t-spanner. Let $G^* = (V, E^*)$ be an interference-optimal t-spanner for G. Since G_{LISE} is not optimal, it follows that $I(G_{LISE}) > I(G^*)$. Thus all edges in E^* have coverage strictly less than $I(G_{LISE})$. From Lemma 5 follows that E^* is a nontrivial subset of E_{LISE}. Let T be the set of edges in E_{LISE} with coverage $I(G_{LISE})$ and $\tilde{G} = (V, \tilde{E})$ the graph with $\tilde{E} = E_{LISE} \setminus T$. $\tilde{G}$ is a t-spanner since E^* is still a subset of $\tilde{E}$, and $I(\tilde{G}) \leq I(G_{LISE}) - 1$ holds. Because T is eventually inserted into E_{LISE} in line 8, there exists an edge $(u, v) \in E$ that was taken in line 5 and for which no path $p(u, v)$ exists in $\tilde{G}$ with $|p| \leq t\,|u, v|$. Thus $\tilde{G}$ is no t-spanner (and therefore also G^*), which contradicts the assumption that G^* is an interference-optimal t-spanner. □

As regards the running time of LISE, it computes for each edge at most one shortest path. This holds since multiple shortest path computations for the same edge in line 6 cause at least as many edges to be inserted into E_{LISE} in line 8 without computing shortest paths for them. Since finding a shortest alternative path for an edge requires $O(n^2)$ time and as the network contains at most the same amount of edges, the overall running time of LISE is as well polynomial in the number of network nodes.

In contrast to the problem of finding a connected topology with optimum interference, the problem of finding an interference-optimal t-spanner is locally solvable. The reason for this is that finding an interference-optimal path $p(u, v)$ for an edge (u, v) with $|p| \leq t\,|u, v|$ can be restricted to a certain neighborhood of (u, v).

LLISE

```
collect (t/2)-neighborhood G_N = (V_N, E_N) of G = (V, E)
E = ∅
G' = (V_N, E')
repeat
    f = edge ∈ E_N with minimum coverage
    move all edges ∈ E_N with coverage Cov(f) to E'
    p = shortestPath(u − v) in G'
until |p| ≤ t |u, v|
inform all edges on p to remain in the resulting topology.
```

Note: $G_{LL} = (V, E_{LL})$ consists of all edges eventually informed to remain in the resulting topology.

In the following we describe a local algorithm similar to LISE that is executed at all eligible edges of the given network. In reality, algorithm LLISE (Local LISE) is executed for each edge by one of its incident nodes (for instance the one with the higher identifier). The description of LLISE assumes the point of view of an edge $e = (u, v)$. The algorithm consists of three main steps:

1) Collect $(\frac{t}{2})$-neighborhood,
2) compute minimum interference path for e, and
3) inform all edges on that path to remain in the resulting topology.

In the first step, e gains knowledge of its $(\frac{t}{2})$-neighborhood. For a Euclidean spanner, the k-neighborhood of e is defined as all edges that can be reached (or more precisely at least one of their incident nodes) over a path p starting at u or v, respectively, with $|p| \leq k\,|e|$. Knowledge of the $(\frac{t}{2})$-neighborhood at all edges can be achieved by local flooding.

THEOREM 7. *The graph $G_{LL} = (V, E_{LL})$ constructed by LLISE from a given network $G = (V, E)$ is an interference-optimal t-spanner of G.*

During the second step a minimum-interference path p from u to v with $|p| \leq t\,|e|$ is computed. LLISE starts with a graph $G_{LL} = (V, E_{LL})$ consisting of all nodes in the $(\frac{t}{2})$-neighborhood and an initially empty edge set. It inserts edges consecutively into E_{LL}—in ascending order according to their coverage—, until a shortest path $p^*(u, v)$ is found in G_{LL} with $|p^*| \leq t\,|e|$.

In the third step, e informs all edges on the path found in the second step to remain in the resulting topology. The resulting topology then consists of all edges receiving a corresponding message. In the following we show that it is sufficient for e to limit the search for an interference-optimal path $p(u, v)$ meeting the spanner property to the $(\frac{t}{2})$-neighborhood of e.

LEMMA 8. *Given an edge $e = (u, v)$, no path p from u to v with $|p| \leq t\,|e|$ contains an edge which is not in the $(\frac{t}{2})$-neighborhood of e.*

PROOF. For the sake of contradiction we assume that a path p from u to v with $|p| \leq t\,|e|$ containing at least one edge (w,x) not in the $(\frac{t}{2})$-neighborhood of e. Without loss of generality we further assume that, traversing p from u to v, we visit w before x. Since (w,x) is not in the $(\frac{t}{2})$-neighborhood, by definition, no path from u to w with length less than or equal to $(\frac{t}{2})|e|$ exists (the same holds for any path from v to x). Consequently, the inequality $|p| > t\,|e| + |(w,x)|$ holds, which contradicts the assumption that $|p| \leq t\,|e|$. □

With Lemma 8 we are now able to prove that the topology constructed by LLISE is a t-spanner with optimum interference.

PROOF. The spanner property of LLISE can be proven similar to the first part of the proof of Theorem 6, where LISE is shown to be a t-spanner.

To show interference optimality, it is sufficient to prove that the spanner path constructed for any edge $e = (u,v) \in G$ by LLISE is interference-optimal, where interference of a path is defined as the maximum interference of an edge on that path. The reason for this is that only edges that lie on one of these paths remain in the resulting topology; non-optimality of G_{LL} would therefore imply non-optimality of at least one of these spanner paths. In the following we look at the algorithm executed by $e = (u,v)$. In line 6 edges in E are consecutively inserted into E', starting with $E' = \emptyset$, until a spanner path p from u to v is found in line 8. Since LLISE inserts the edges into E' in ascending order according to their coverage and p is the first path meeting the spanner property, p is an interference-optimal t-spanner path from u to v in the $(\frac{t}{2})$-neighborhood. From Lemma 8 we know that the $(\frac{t}{2})$-neighborhood of e contains all spanner paths from u to v and therefore also the interference-optimal one. Thus it is not possible that LLISE does not see the global interference-optimal t-spanner path due to its local knowledge about G. Consequently, p is the global interference-optimal t-spanner path of e. □

6. AVERAGE-CASE INTERFERENCE

In this section we consider interference of topology control algorithms on average-case graphs, that is on graphs with randomly placed nodes.

In particular networks were constructed by placing nodes randomly and uniformly on a square field of size 20 by 20 units and subsequently computing for each node set the Unit Disk Graph—defined such that an edge exists if and only if its Euclidean length is at most one unit. The resulting Unit Disk Graphs were then employed as input networks for topology control. Since node density is a fundamental property of networks with randomly placed nodes, the networks were generated over a spectrum of node densities.

6.1 Connectivity-Preserving Topologies

To evaluate connectivity-preserving topologies on average-case graphs, two well-known topology control algorithms are considered, in particular the Gabriel Graph [6] and the Relative Neighborhood Graph [22]. The interference-reducing effect of these two constructions is considered by comparison with the interference value of the given Unit Disk Graph network on the one hand and with the interference-optimal connectivity-preserving topology on the other hand. The

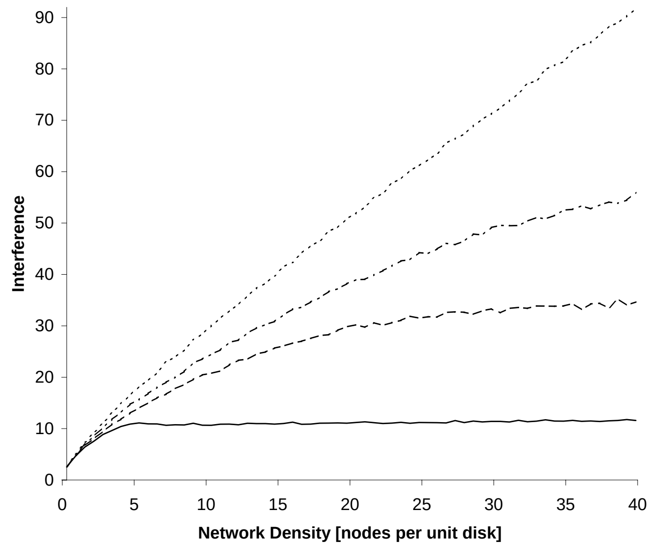

Figure 10: Interference values of the Unit Disk Graph without topology control (dotted), the Gabriel Graph (dash-dotted), the Relative Neighborhood Graph (dashed), and the interference-optimal connectivity-preserving topology (solid).

interference-optimal topology was constructed by means of the LIFE algorithm presented in Section 5.

Figure 10 shows the interference mean values over 1000 networks for each simulated network density. While the resulting interference curves behave similarly for very low network densities, they fall into three groups with increasing density: At a density of roughly 5 network nodes per unit disk the interference-optimal curve stagnates and remains at a value of approximately 11.5. On the other hand the interference curve of the Unit Disk Graph without topology control rises almost linearly. Between these two extremes the Gabriel Graph and Relative Neighborhood Graph values increase clearly more slowly than the Unit Disk Graph curve, but show significantly higher values than the interference-optimal topology.

The simulation results show that the edge reduction performed by the Gabriel Graph and Relative Neighborhood Graph constructions reduce interference of the given network; this effect is clearer with the Relative Neighborhood Graph due to its stricter edge inclusion criterion and consequently its being a subgraph of the Gabriel Graph. However, the interference values of these two constructions are considerably higher than the results of the interference-optimal connectivity-preserving topology. Furthermore, although (unless in special cases) the Relative Neighborhood Graph has degree at most 6, it is not even clear whether with increasing network density the respective interference curve remains around the maximum value found so far or whether it would increase further for densities beyond the simulated spectrum. It can therefore be concluded that also for average-case graphs sparseness does not imply low interference.

6.2 Low Interference Spanners

Going beyond connectivity-preserving topologies, we consider in this section spanners, that is topologies guaranteeing that the shortest paths on the resulting topology are only

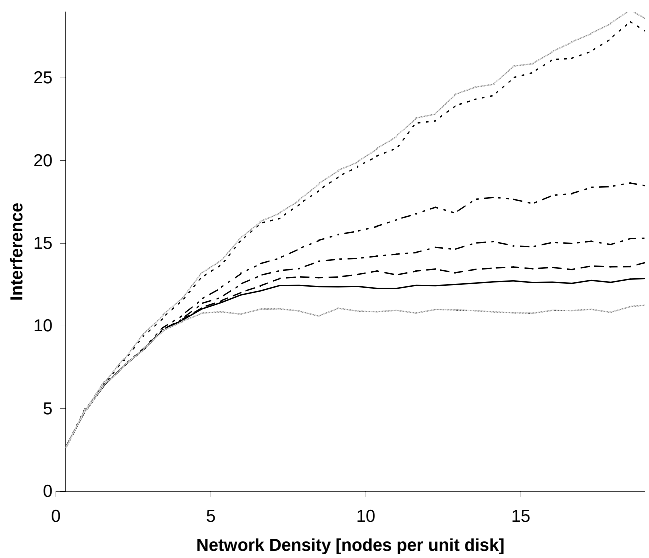

Figure 11: Interference values of LISE for stretch factors 2 (dotted), 4 (dash-dot-dotted), 6 (dash-dotted), 8 (dashed), and 10 (solid). Interference values of the Relative Neighborhood Graph (upper gray) and interference-optimal connectivity-preserving topology (lower gray) are plotted for reference.

by a constant factor longer than on the given network (cf. Section 5.2).

Figure 11 depicts simulation results—in particular the mean interference values over 100 networks at each simulated network density—of the topology constructed by the LISE algorithm introduced in Section 5 for different stretch factors t. The simulation results show that by increasing the requested stretch factor it is possible to achieve interference values close to the optimum interference values caused by connectivity-preserving topologies as described in the previous section. Moreover, even with a low stretch factor of 2, LISE does not perform worse than the Relative Neighborhood Graph, which is *not* a spanner. In summary, the simulation results show that the LLISE algorithm performs well with respect to interference also on average-case graphs. An illustration of the simulation graphs is provided in Figure 12.

7. CONCLUSION

In this paper we disprove the widely advocated assumption that sparse topologies automatically imply low interference. In contrast to most of the related work we provide an intuitive definition of interference. With this interference model we show that currently proposed topology control constructions—although claiming so—do not in the first place focus on reducing interference.

In addition we propose provenly interference-minimal connectivity-preserving and spanner constructions. A locally computable version of the interference-minimal spanner construction can even be considered practicable since it is shown to significantly outperform previously suggested topology control algorithms also on average-case graphs.

8. REFERENCES

[1] K. Alzoubi, P.-J. Wan, and O. Frieder. Message-Optimal Connected Dominating Sets in Mobile Ad Hoc Networks. In *Proc. of the 3^{rd} ACM Int. Symposium on Mobile ad hoc networking & computing (MOBIHOC)*, 2002.

[2] L. Bao and J. Garcia-Luna-Aceves. Topology Management in Ad Hoc Networks. In *Proc. of the 4^{th} ACM Int. Symposium on Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, 2003.

[3] D. Blough, M. Leoncini, G. Resta, and P. Santi. The K-Neigh Protocol for Symmetric Topology in Ad Hoc Networks. In *Proc. of the 4^{th} ACM Int. Symposium on Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, 2003.

[4] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with Guaranteed Delivery in ad hoc Wireless Networks. In *Proc. of the 3^{rd} International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, pages 48–55, 1999.

[5] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*, chapter 24.2. MIT Press, Cambridge, MA, USA, 1990.

[6] K. Gabriel and R. Sokal. A New Statistical Approach to Geographic Variation Analysis. *Systematic Zoology*, 18:259–278, 1969.

[7] T. Hou and V. Li. Transmission Range Control in Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 34(1):38–44, 1986.

[8] L. Hu. Topology Control for Multihop Packet Radio Networks. *IEEE Trans. on Communications*, 41(10), 1993.

[9] L. Jia, R. Rajaraman, and R. Suel. An Efficient Distributed Algorithm for Constructing Small Dominating Sets. In *Proc. of the 20^{th} ACM Symposium on Principles of Distributed Computing (PODC)*, pages 33–42, 2001.

[10] B. Karp and H. Kung. GPSR: Greedy Perimeter Stateless Routing for Wireless Networks. In *Proc. of the 6^{th} Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 243–254, 2000.

[11] V. Kawadia and P. Kumar. Power Control and Clustering in Ad Hoc Networks. In *Proc. of the 22^{nd} Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.

[12] F. Kuhn and R. Wattenhofer. Constant-Time Distributed Dominating Set Approximation. In *Proc. of the 22^{nd} ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.

[13] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric Routing: Of Theory and Practice. In *Proc. of the 22^{nd} ACM Symposium on the Principles of Distributed Computing (PODC)*, 2003.

[14] N. Li, C.-J. Hou, and L. Sha. Design and Analysis of an MST-Based Topology Control Algorithm. In *Proc. of the 22^{nd} Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.

[15] X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed Construction of Planar Spanner and Routing for Ad

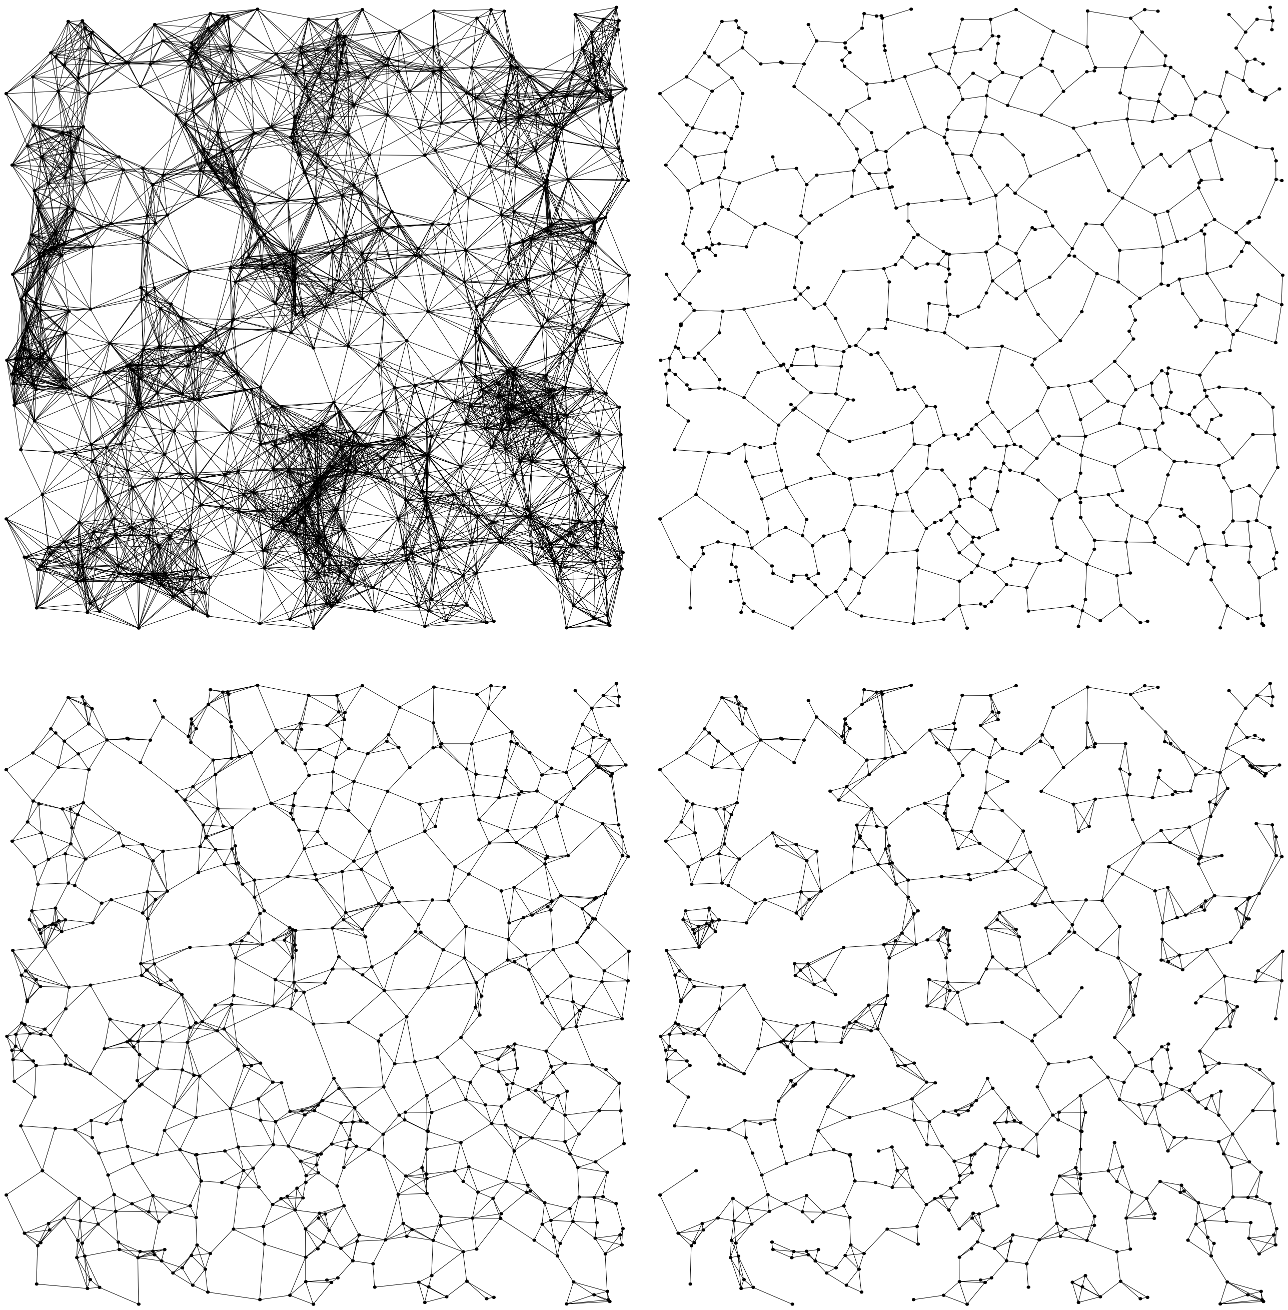

Figure 12: The Unit Disk Graph G (top left, interference 50), the Relative Neighborhood Graph of G (top right, interference 25), G_{LL} computed by LLISE with stretch factor 2 (bottom left, interference 23) and 10 (bottom right, interference 12) at a network density of 20 nodes per unit disk on a square field of 10 units side length. Note that, for instance in the western region of the graph, LLISE—depending on the chosen stretch factor—omits high-interference "bridge" edges if alternative spanning paths exist.

Hoc Wireless Networks. In *Proc. of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.

[16] F. Meyer auf der Heide, C. Schindelhauer, K. Volbert, and M. Grünewald. Energy, Congestion and Dilation in Radio Networks. In *Proceedings of the 14th ACM Symposium on Parallel Algorithms and Architectures*, 10 - 13 Aug. 2002.

[17] D. Peleg. *Distributed Computing, A Locality-Sensitive Approach*, chapter 5.5. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

[18] 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 (DIAL-M)*, 1999.

[19] 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.

[20] V. Rodoplu and T. H. Meng. Minimum energy mobile wireless networks. *IEEE J. Selected Areas in Communications*, 17(8), 1999.

[21] H. Takagi and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transactions on Communications*, 32(3):246–257, 1984.

[22] G. Toussaint. The Relative Neighborhood Graph of a Finite Planar Set. *Pattern Recognition*, 12(4):261–268, 1980.

[23] 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.

[24] 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)*, pages 1388–1397, 2001.