Topology Control for Wireless Networks with Highly-Directional Antennas

Thomas Stahlbuhk[†‡], Brooke Shrader[‡] and Eytan Modiano[†]

[†] Massachusetts Institute of Technology, Cambridge, MA

[‡] MIT Lincoln Laboratory, Lexington, MA

Abstract—In order to steer antenna beams towards one another for communication, wireless nodes with highly-directional antennas must track the channel state of their neighbors. To keep this overhead manageable, each node must limit the number of neighbors that it tracks. The subset of neighbors that each node chooses to track constitutes a network topology over which traffic can be routed. We consider this topology design problem, taking into account channel modeling, transmission scheduling, and traffic demand. We formulate the optimal topology design problem, with the objective of maximizing the scaling of traffic demand, and propose a distributed method, where each node rapidly builds a segment of the topology around itself by forming connections with its nearest neighbors in discretized angular regions. The method has low complexity and message passing overhead. The resulting topologies are shown to have desirable structural properties and approach the optimal solution in high path loss environments.

I. INTRODUCTION

In recent years, mesh networking using millimeter wave communications has emerged as a promising technology with application to the cellular backbone [1], broadband access [2], and data center networks [3]. Interest in this technology is driven by the ability of millimeter wave systems to support high throughput by accessing large portions of available spectrum [4]. However, access to this wide bandwidth comes at the expense of the high path loss found at millimeter wave frequencies [5]. Because of this severe path loss, millimeter wave systems rely on highly-directional antennas to construct reliable links [5].

In applications with node mobility and environmental motion, the use of directional antennas is complicated as the channel state frequently changes [6]. To effectively use highly-directional antennas in this setting, nodes must track the channel state of their neighbors so that they may steer beams towards one another for communication [4], [7]. This introduces tracking overhead that grows with the number of neighbors that a node tracks. To keep this overhead manageable, nodes must limit the subset of neighbors that they track and subsequently communicate with [8]. This introduces the need for topology control, the process of choosing the set of neighbors with which each node can communicate within the network.

Topology control for communication networks has previously been studied in both wireline and wireless settings. In optical communication networks, topology control has focused on building network structures that can support a specified end-to-end traffic demand where each node in the network is limited in the number of connections it may form [9]. In [10], [11], and [12], this problem was formulated as an integer linear program and branch exchange heuristics were explored for obtaining suboptimal solutions. In contrast, [13] explored fixed graphical structures for generating topologies with desirable properties.

In wireless networks, topology control has been studied in the context of power control. In this setting, nodes adjust their transmission range to balance power consumption, interference, and connectivity. In [14] centralized and distributed heuristics were explored for topology control. In [15] and [16] algorithms using the geometric position of nodes were used for constructing topologies with guaranteed performance, and in [17] algorithms for building topologies with given degree distributions were demonstrated.

In wireless networks with directional antennas, previous work in topology control has focused on ensuring network connectivity and reducing inter-link interference. In [18] and [19] heuristic algorithms for generating connected topologies were considered in beam steering and beam switching systems respectively. Similarly, [20] and [21] considered algorithms for generating topologies with low inter-link interference. The trade-off between node-degree and path-stretch was explored in [8].

In this work we formulate the problem of topology control with the objective of maximizing throughput as measured by the scaling of end-to-end traffic between users, subject to scheduling and routing constraints. This focus is distinct from the previous literature on directional-antenna topology control that focused on optimizing graphical properties of the topology such as connectivity or path length. Central to our work is the problem of overcoming the high path loss of millimeter wave communications. We model the wireless channel using a generalized path loss model, where path loss scales on the order of d^α with d the distance between transmitter and receive. Such path loss models have been shown to approximate well the dominate effects of wireless environments [22]. The impact of the path loss exponent, α, on the topology design problem is fundamentally explored.

In order to mitigate some of this severe path loss, future

millimeter wave systems are envisioned as relying on highly-directional antenna systems. Because of the channel tracking overhead incurred by highly-directional antenna systems, we consider topologies limited by node degree constraints. Using the results of [23], we show that the link scheduling of millimeter wave systems with highly-directional antennas is restricted by primary interference constraints. Thus, in order to form a topology with good throughput performance, the network must take into account the interacting effects of the channel model and wireless resource scheduling.

In this work, we formulate the topology control problem as a mixed-integer linear program. Although solutions to this problem are informative as to what characteristics make for good topologies, the problem is computationally intractable in general, and direct solution is, therefore, impractical for implementation. We thus present a distributed method for constructing topologies and show that under the above channel modeling and link scheduling constraints, the topologies generated by our algorithm approach the optimal solution when the path loss is sufficiently large. Moreover, we show through simulation our algorithm generates near optimal topologies for the range of path loss values reported for millimeter wave communications.

The outline of the paper is as follows. In Section II we discuss the impact of tracking overhead and resource scheduling on topology control and formulate the optimal topology control problem. We analyze the computational complexity of the problem and explore the structure of optimal topologies in Section III. In Section IV, we describe the distributed topology control method and derive optimality properties of the resulting topologies. In Section V, we give simulation results and conclude in Section VI.

II. PROBLEM FORMULATION

We consider a wireless mesh network modeled as a directed graph $G = (N, E)$ where N is the set of nodes in the network and E is the set of edges. A node is denoted $i \in N$ and an edge from node i to j as $(i, j) \in E$. Each edge (i, j) has an edge capacity $c_{(i,j)}$. Each node in the network is equipped with a directional antenna. At a given time a node may focus its antenna beam towards one neighbor, whose channel state information it is tracking, to transmit or receive data. In our model, data may be forwarded through multiple hops in the network from its source node s to its destination t. We define an end-to-end demand matrix Λ, where the element $\Lambda_{s,t}$ of the matrix specifies a desired rate of flow from node s to t.

The objective of the topology control problem is to construct a subgraph $G' = (N, E')$ with $E' \subseteq E$ such that the maximum degree of G' is no greater than T and there exists a scheduling and routing policy that maximizes a scaling factor of Λ (i.e., maximize ρ such that for all s, t node pair, $\rho \Lambda_{s,t}$ can be supported). Note that this objective is closely related to that of [9], [10], [11], and [12]. We now proceed to enumerate the physical constraints of directional antenna systems, leading to a concrete problem formulation in Section II-D.

A. Beam Steering and Tracking Overhead

In antenna array beam steering systems, nodes utilize multiple antennas to transmit and receive over the wireless spectrum. Through signal processing techniques applied over each element, nodes may steer transmitter power and receiver sensitivity towards desired spatial directions [6]. However, in order to effectively perform this processing, nodes must learn each other's channel state information. In antenna arrays that contain many elements, a necessity for high directionality [6], this is often accomplished through beam search procedures where a transmitter and receiver test multiple beam pattern combinations to find a suitable pair for communication (e.g., [7]). Due to the time required to search over the beam set, this process can introduce considerable overhead which grows with increasing antenna directionality [24]. As the channel state between nodes evolves over time, due to node mobility and movement in the environment, beam patterns must be continuously adapted. Thus, in order to keep this overhead tractable, nodes must limit the number of neighbors that they track and form connections with [8].

This constraint gives rise to a topology control problem where each node has a degree constraint. Specifically, we focus on networks with homogeneous nodes and thus require a topology $G' = (N, E')$ with $E' \subseteq E$ such that each node in the network is incident to at most T edges in E'. As E' consists of all the edges maintained in the topology, any subsequent scheduling and routing policy is limited to using only these links.

B. Resource Scheduling

Access to the wireless medium must be shared by all nodes in the network. Thus, in order to optimize the topology we consider restrictions on simultaneous transmissions. In [23] it was shown that highly-directional antenna networks operating at millimeter wavelengths have minimal inter-link interference between pairs of communicating nodes. However, each node is restricted to form only a single beam to transmit or receive over [23]. This imposes a primary interference scheduling constraint, where simultaneous edge activations are limited to matchings over the underlying graph.

In order to model the above we adopt the formulation of [25]. We specify an activation set A where each element of the set $a \in A$ is a unique matching in graph G. For each time unit that matching a is activated, edge (i, j) obtains a data rate

$$c_{(i,j)}^a = \begin{cases} c_{(i,j)}, & \text{if } (i,j) \text{ is in matching } a \\ 0, & \text{otherwise} \end{cases}$$

where $c_{(i,j)}$ is defined to be the capacity of edge $(i, j) \in E$.

C. Channel Modeling

Fundamental to the topology design problem is the impact of channel path loss. In this work we assume nodes are distributed on the two-dimensional plane with the distance between nodes i and j denoted as $d(i, j)$. Then the capacity between i and j is given by the following expression that approximates the capacity of an additive white Gaussian noise

channel in the power-limited regime under a generalized path loss model [22].

$$c_{(i,j)} = K \left(\frac{d_0}{d(i,j)} \right)^{\alpha} \tag{1}$$

In [27], path loss exponents (α) were empirically found to be roughly in the range of 3 to 4.5 for millimeter wave communications operating in urban environments. We explore the impact of α on topology performance below.

D. Model

With the above considerations, we specify the topology control problem as a mixed-integer linear program (2). We define the following variables:

1) Flow variables $x_{(i,j)}^{s,t}$ specifying the rate of commodity flow from source s to destination t transversing edge $(i,j) \in E$.
2) Scheduling variables y^a indicating the fraction of total time (normalized to 1) that matching $a \in A$ is activated.
3) Topology variables $z_{(i,j)}$ which take value 1 if $(i,j) \in E'$ and 0 otherwise.

As specified above, the objective is to generate an edge subset E' such that a feasible scheduling and routing policy can deliver $\rho\Lambda_{s,t}$ rate of flow between all source s and destination t pairs. The solution to the topology control problem specifies the $z_{(i,j)}$ variables in (2) which completely define the topology G'. The scheduling and routing policy can then be viewed as an assignment to variables $x_{(i,j)}^{s,t}$ and y^a such that flow only transverses over the generated topology and the scaling variable ρ is maximized. The following mixed-integer linear program is a mathematical formulation of the optimal topology design problem.

maximize ρ

subject to:

$$\sum_{j:(i,j)\in E} x_{(i,j)}^{s,t} - \sum_{j:(j,i)\in E} x_{(j,i)}^{s,t} = \begin{cases} \rho\Lambda_{s,t} & \text{if } i=s \\ -\rho\Lambda_{s,t} & \text{if } i=t \\ 0 & \text{if otherwise} \end{cases},$$
$$\forall i,s,t \in N \tag{2a}$$

$$\sum_{s,t\in N} x_{(i,j)}^{s,t} \leq \sum_{a\in A} c_{(i,j)}^a y^a, \qquad \forall (i,j)\in E \tag{2b}$$

$$\sum_{a\in A} y^a \leq 1 \tag{2c}$$

$$x_{(i,j)}^{s,t} \leq M z_{(i,j)}, \qquad \forall s,t \in N, \forall(i,j)\in E \tag{2d}$$

$$z_{(i,j)} = z_{(j,i)}, \qquad \forall (i,j)\in E \tag{2e}$$

$$\sum_{j:(i,j)\in E} z_{(i,j)} \leq T, \qquad \forall i \in N \tag{2f}$$

$$x_{(i,j)}^{s,t} \geq 0, \qquad \forall s,t\in N, \forall(i,j)\in E \tag{2g}$$

$$y^a \geq 0, \qquad \forall a \in A \tag{2h}$$

$$z_{(i,j)} \in \{0,\ 1\}, \qquad \forall(i,j)\in E \tag{2i}$$

In the above formulation, constraint (2a) forces conservation of flow through the network subject to the exogenous data input. Constraints (2b) and (2c) ensures that each edge receives enough activation time to support its rate of flow and that the total duration of all activations sums to at most one unit of time. Constraints (2d), (2e), and (2f) state that the topology must not violate the degree constraint requirements at each node and that data only transverses the topology subset. The value M in (2) is a sufficiently large number (e.g., $M = \max c_{(i,j)}^a$).

III. THE TOPOLOGY CONTROL PROBLEM

We now proceed to explore solutions to the topology control problem given above. After examining the structural features that make for good topologies, we propose a distributed algorithm in Section IV.

A. Approximating the Scheduling Solution

Solving the scheduling problem of Section II-B is computationally complex [28]. In order to reduce the complexity of the problem and evaluate the performance of topologies with relatively large numbers of nodes, throughout this section we use the approximation method of [28]. In [28] it was shown that for given $z_{(i,j)}$ values, replacing constraints (2b) and (2c) with the constraint,

$$\sum_{s,t\in N} \left(\sum_{j:(i,j)\in E} \frac{x_{(i,j)}^{s,t}}{c_{(i,j)}} + \sum_{j:(j,i)\in E} \frac{x_{(j,i)}^{s,t}}{c_{(j,i)}} \right) \leq \frac{2}{3}, \tag{3}$$

results in a solution that is at least $\frac{2}{3}$ of the optimal solution of (2). The approximation of [28] abstracts the scheduling problem as a multigraph coloring problem where each color represents a discrete time slot and each edge a unit of data transversing a link. The approximation then uses the well known result that the chromatic index of a multigraph is at most $\frac{3}{2}$ the maximum degree of the graph, and thus the schedule length sums to at most 1 if the total data entering and exiting any single node is limited to $\frac{2}{3}$ units. See [28] for details. Throughout this section we use (3) in place of (2b) and (2c) and interpret the resulting solutions as being closely related to the optimal solutions to (2).

B. Computation Time

In Fig. 1, we evaluate the average computation time over 25 random problem instantiations using CPLEX, a commonly used optimization toolbox, to directly solve the topology control problem for increasing node set size with $T = 6$. In each instant, nodes are uniformly distributed inside the unit disk with capacity given by (1), $\alpha = 3$, and each node sends or receives one flow with unit demand to or from one other node selected uniformly at random. We note that the time required to solve the problem grows rapidly with the number of nodes in the graph, with run times on the order of minutes for $|N| = 34$.

(a) Demands Between Nodes   (b) Optimal Topology $\alpha = 2$ (c) Optimal Topology $\alpha = 4$
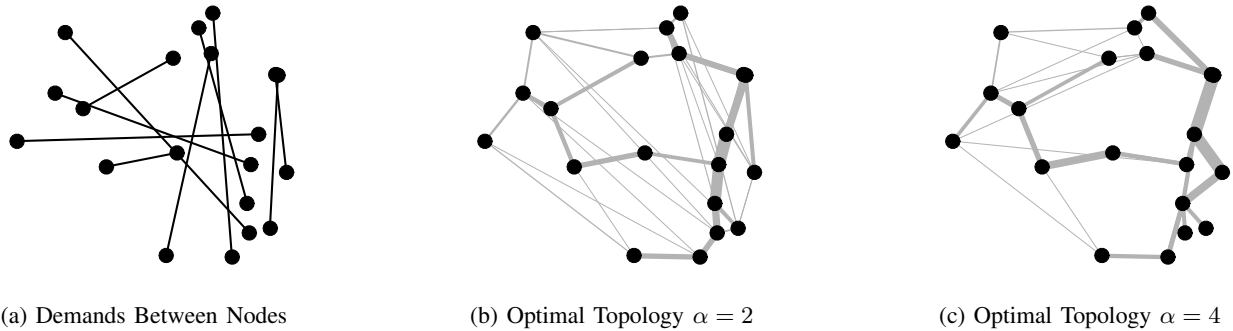
Fig. 2: Structure of optimal topologies for 20 nodes distributed on the unit disk with $T = 6$.
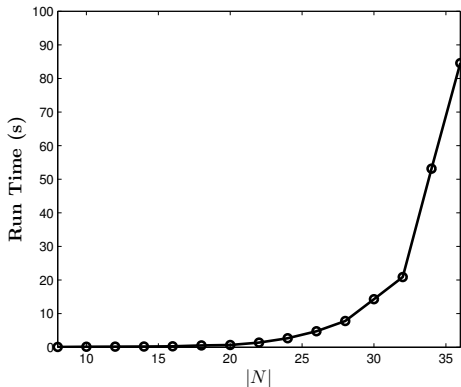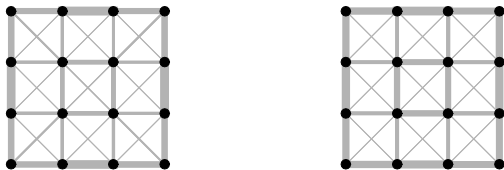

Fig. 1: Computation time for increasing number of nodes.

(a) Optimal Topology $\alpha = 2$ (b) Optimal Topology $\alpha = 4$

Fig. 3: Structure of optimal topologies for 16 nodes arranged on a square grid with $T = 8$.

C. Structure of Optimal Solutions

In order to gain insight into the properties of *good* topologies, we now explore the structure of optimal solutions to (2). In Fig. 2, 20 nodes are distributed inside the unit disk with each node the source or destination of one commodity with unit demand, as illustrated in Fig. 2a. Fig. 2b and Fig. 2c show the optimal topologies for the cases where α is 2 and 4 respectively. The width of each edge is proportional to the total amount of flow transversing that edge. It can be seen from the figure that the most utilized paths contain short edges that form corridors between sources and destinations. For $\alpha = 2$ the topology uses some long distance edges to help mitigate

congestion at centrally located nodes. However, for $\alpha = 4$ these long distance edges are rarely used.

The example of Fig. 2 considered relatively sparse graphs in both node density and traffic. In order to gain insight into more dense networks, in Fig. 3 we plot the optimal topologies when 16 nodes are arranged in a square grid and every node has unit demand destined to every other node in the network. Again the width of each edge is proportional to the amount of traffic. We see that the optimal topologies use only short edges. Although long edges could reduce the hop count between sources and destinations, because of path loss the additional time required to schedule these long edges undoes any benefit.

From the above observations we see that good topologies contain paths consisting of short edges that connect source nodes with their destinations. Because the path loss heavily penalizes long edges, short edges support a large fraction of the flow in the network. We next propose a sectorized-distributed method that constructs topologies that capture this general structure.

IV. DISTRIBUTED TOPOLOGY CONTROL

In this section we present a distributed method for rapidly constructing a network topology that satisfies the degree constraint T at every node. We assume all nodes exist on a two-dimensional plane. Each node partitions the plane into discrete cone-shaped sectors around itself where each cone has angular width $\frac{2\pi}{T}$, as shown in Fig. 4. The node then elects to form a topological edge with the nearest (minimum distance) neighbor in each cone. If the choice is reciprocal (i.e., the neighbor also elects to form an edge with that node) an edge is established in E'. It is easy to see that this method constructs a topology that satisfies the degree constraint T at every node. A key result, similar to that of [15], is that if the angular size of each cone is no greater than $60°$ (i.e., $T \geq 6$) then the resulting topology is connected for all node pairs. We present a proof of this below. The resulting algorithm has low computational complexity and requires only message passing with at most T neighbors.

A. The Sectorized-Distributed Algorithm

The sectorized-distributed algorithm is implemented at each node independently. The operations carried out at each node
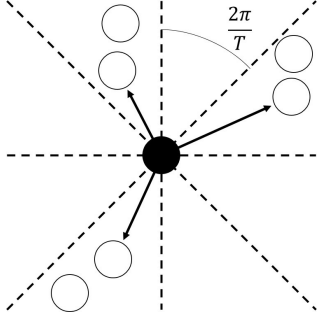
Fig. 4: Sector partitioning and neighbor selection in Algorithm 1

(a) Node placement of Fig. 2   (b) Node placement of Fig. 3

Fig. 5: Topologies generated by Algorithm 1

i are described in Algorithm 1. At each node, the algorithm takes as input the distance $d(i,j)$ and cone sector S_j of each node $j \in N$. The sectors about node i may be arbitrarily oriented and labeled. The algorithm uses the set $N'(i)$ to denote those nodes that i elects to be neighbors. Edge (i,j) is included in topology E' if and only if $j \in N'(i)$ *and* $i \in N'(j)$. Note that for proper function the algorithm does not require that the distance $d(i,j)$ of each neighbor is known precisely, only that the relative order of each neighbor's distance is known. Thus, $d(i,j)$ and S_j only need to be updated at the rate at which nodes interchange ordered distance from i and move between sectors. In general, this rate is orders of magnitude smaller than the rate at which the channel state information changes.
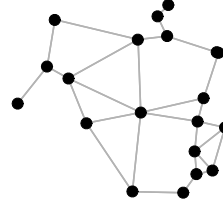

Algorithm 1 Sectorized-Distributed Method at Node i

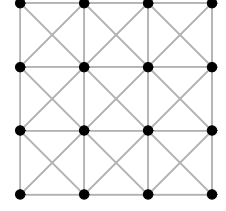**Input:** Distance $d(i,j)$ and Sector S_j for each $j \in N$
Output: The set of edges $(i,j) \in E'$ for all $j \in N$
  1: $N'(i) \leftarrow \emptyset$
  2: **for** $t = 1$ to T **do**
  3:    **if** There exists at least one j with $S_j = t$ **then**
  4:      Add to set $N'(i)$ the node j with minimum $d(i,j)$ such that $S_j = t$
  5:    **end if**
  6: **end for**
  7: **for** $j \in N'(i)$ **do**
  8:    Inform $j \in N'(i)$ that a connection is desired
  9:    If j replies that $i \in N'(j)$ then $(i,j) \in E'$
 10: **end for**

In Fig. 5a we plot a topology generated by Algorithm 1 for the same setting as Fig. 2. Note that the edges in this plot only indicate presence in the topology and are not weighted by their traffic flow. We see that many of the most utilized edges in Fig. 2 are present in the generated topology. In Fig. 5b we give the topology generated for the network of Fig. 3 again noting that the heavily utilized edges in Fig. 3 are captured by the topology. In Section V we compare throughput performance results of topologies generated by Algorithm 1 to optimal topologies.

As a final note, the sectorized-distributed algorithm does not maximize the number of edges in the topology and many nodes in the network could end up with far fewer than T adjacent edges after completion of Algorithm 1. (e.g., Fig. 5a). One technique that can improve performance is to apply additional algorithms that add edges to a topology generated by the sectorized-distributed algorithm until each node reaches its degree constraint. Thus, Algorithm 1 could be a first step in a larger topology construction process. We explore such extensions in Section V.

B. Connectivity of the Topology

We now proceed to show that topologies generated by Algorithm 1 are connected when $T \geq 6$. Our proof is based on a result from [15] on the connectivity of wireless networks with omni-directional antennas. In this section and Section IV-C we assume each edge length $d(i,j)$ is unique. In practice this constraint is nearly inconsequential as modifying any duplicate $d(i,j)$ by an arbitrary small value will fulfill the requirement.

Theorem 1. *For* $T \geq 6$*, Algorithm 1 generates a connected topology,* G'*.*

We first give the following lemma, due to [15], which will be used in the proof of Theorem 1.

Lemma 1. *If* $T \geq 6$*,* $\forall i,j \in N$ *either* $j \in N'(i)$ *or* $\exists k \in N'(i)$ *such that* $d(i,k) \leq d(i,j)$ *and* $d(k,j) \leq d(i,j)$*.*

Proof: Consider the $\frac{2\pi}{T}$ wide cone extending from node i that is occupied by node j. If j is the closest node to i, then $j \in N'(i)$ and the lemma's condition is met. Otherwise, there must exist another node k in the cone such that $d(i,k) \leq d(i,j)$.

Now, because the cone containing j and k has angular width no greater than $\frac{\pi}{3}$ (since $T \geq 6$) this implies that $\angle jik \leq \frac{\pi}{3}$. Furthermore, from basic trigonometry, $d(i,j) \geq d(i,k)$ implies that $\angle ikj \geq \angle ijk$. Since, $\angle ijk + \angle ikj + \angle jik = \pi$, this implies that $\angle ikj \geq \frac{\pi}{3} \geq \angle jik$.

Thus, since $\angle ikj \geq \angle jik$, $d(k,j) \leq d(i,j)$. ∎

Note that by the definition of Algorithm 1 the above implies that if $(i,j) \notin E'$ there must exist a node k such that $d(i,k) \leq d(i,j)$ and $d(k,j) \leq d(i,j)$. We now prove Theorem 1. The proof is similar to that in [15] but uses a simpler constructive argument.

Proof: Our proof is constructive. We show that any two arbitrary nodes i and j are connected in the topology G' by a path P, and thus topology G' is connected.

We begin by constructing a path $P = \{(i,j)\}$ (i.e., the path consisting of edge (i,j)). If edge $(i,j) \in E'$ then P clearly connects i and j in topology G' and the claim holds. Otherwise, if $(i,j) \notin E'$, Lemma 1 implies there exists a node k such that $d(i,k) < d(i,j)$ and $d(k,j) < d(i,j)$. (Note that the strict inequalities follow from the assumption of unique edge lengths.) We remove (i,j) from P and, in its place, insert (i,k) and (k,j). Clearly P is still a path from i to j in E.

We now recursively carry out the above operation. We choose the maximum length edge in P that is not in E' (i.e., $\max_{(l,m)\in P}\{d(l,m):(l,m)\notin E'\}$), remove it, and add two edges from E of strictly smaller length that connect the end points of the removed edge. Note that Lemma 1 guarantees these two edges must exist and P remains a path from i to j after the operation. Now, after each operation one of two things must occur: either $P \subseteq E'$, or the length of the maximum length edge in P that is not in E' must decrease. Therefore, in no more than $|E|$ iterations we must obtain a path $P \subseteq E'$. ∎

Note that an immediate result of the above proof is that either edge (i,j) is in E' or there exists a path from i to j in E' such that each edge has length less than $d(i,j)$.

C. Structural Properties of the Topology

We next explore properties of topologies generated by the above sectorized-distributed method. Throughout this subsection we assume that $G' = (N, E')$ is a topology generated by Algorithm 1 with $T \geq 6$.

Define the capacity of an arbitrary path P from node s to t (denoted $C(P)$) as the maximum flow that may transverse path P under the link activation constraints of Section II-B. Define the optimal path $P_{s,t}^* \subseteq E$ as the single path that obtains the maximum capacity between s and t (i.e., $C(P_{s,t}^*) \geq C(P)$ for all $P \subseteq E$ from s to t). We show that for every s,t node pair there exists a path $P_{s,t}' \subseteq E'$ such that $C(P_{s,t}') \geq \frac{1}{2}C(P_{s,t}^*)$, if the edge capacity $c_{(i,j)}$ is a monotonically decreasing function of $d_{(i,j)}$ (e.g., see (1)).

Theorem 2. $\exists P_{s,t}' \subseteq E'$ such that $C(P_{s,t}') \geq \frac{1}{2}C(P_{s,t}^*)$ if $c_{(i,j)}$ is a monotonically decreasing function of $d_{(i,j)}$.

To prove Theorem 2, we first establish the following two lemmas. We denote the set of all paths from node s to t in E by $\mathcal{P}^{s,t}$.

Lemma 2. For all s,t node pairs, $\exists P_{s,t}' \subseteq E'$ from s to t such that

$$\max\{d(i,j):(i,j)\in P_{s,t}'\} = \min_{P\in\mathcal{P}^{s,t}}\max\{d(i,j):(i,j)\in P\} \tag{4}$$

Lemma 2 states that E' contains a path that has the minimum *maximum edge length*.

Proof: Consider any path $P \subseteq E$ from s to t. Consider any edge $(i,j) \in P$. Assume $(i,j) \notin E'$. Then by the proof

of Theorem 1 there exists a path from i to j in E' such that all edges in the path have length less than $d(i,j)$. Since this is true for all $(i,j) \in P$, (4) follows. ∎

Lemma 3. *For any path P,*

$$\frac{1}{2}\min_{(i,j)\in P}c_{(i,j)} \leq C(P) \leq \min_{(i,j)\in P}c_{(i,j)}$$

Proof: The capacity of any path is upper bounded by its minimum cut, thus $C(P) \leq \min_{(i,j)\in P}c_{(i,j)}$. Now, number the edges along path P from source s to destination t using the natural numbers, $0,1,2,...,|P|-1$. Clearly all odd (even) edges form a valid activation. By activating the odd edges for half of the time and even edges for the other half, it can be observed that $C(P) \geq \frac{1}{2}\min_{(i,j)\in P}c_{(i,j)}$ is achieved and the result follows. ∎

We now prove Theorem 2.

Proof: From Lemma 2, there exists a $P_{s,t}' \subseteq E'$ such that:

$$\max_{(i,j)\in P_{s,t}'}\{d(i,j)\} \leq \max_{(i,j)\in P_{s,t}^*}\{d(i,j)\}$$

Since $c_{(i,j)}$ is a monotonically decreasing function of $d(i,j)$ this implies

$$\min_{(i,j)\in P_{s,t}'}\{c_{(i,j)}\} \geq \min_{(i,j)\in P_{s,t}^*}\{c_{(i,j)}\}$$

By Lemma 3,

$$C(P_{s,t}') \geq \frac{1}{2}\min_{(i,j)\in P_{s,t}'}\{c_{(i,j)}\}$$
$$C(P_{s,t}^*) \leq \min_{(i,j)\in P_{s,t}^*}\{c_{(i,j)}\}$$

It can then be observed that, $C(P_{s,t}') \geq \frac{1}{2}C(P_{s,t}^*)$. ∎

Theorem 2 shows that for any node pair there exists in E' a single *good* path consisting of short hops. As was discussed in Section III these corridors of short-hop paths form much of the backbone for communication through the network. However, in order to truly maximize ρ in (2) the scheduler must schedule edge activations for multiple paths that are competing for the wireless resources. In general, analyzing the interaction of these multiple paths is difficult. However, we will next show that in the presence of high path loss between nodes this interaction becomes increasingly less important as compared to the achievable link capacities.

In particular, assume we are given a placement of nodes on the two dimensional plane, traffic demand matrix Λ, and edge capacities defined by (1). Denote by ρ^* the optimal solution to the corresponding formulation of (2) and by ρ' the optimal solution when the topology variables $z_{(i,j)}$ are obtained using the sectorized-distributed method. We prove that there exists a value α_0 such that for all path loss exponents $\alpha > \alpha_0$, $\rho'/\rho^* \geq 1 - \epsilon$ for any $\epsilon > 0$.

Theorem 3. $\forall \epsilon > 0$, $\exists \alpha_0$ such that $\forall \alpha > \alpha_0$, $\frac{\rho'}{\rho^*} \geq 1 - \epsilon$.

Proof: The complete proof of Theorem 3 is omitted for brevity. Here we provide an outline of the proof that conveys the intuition.
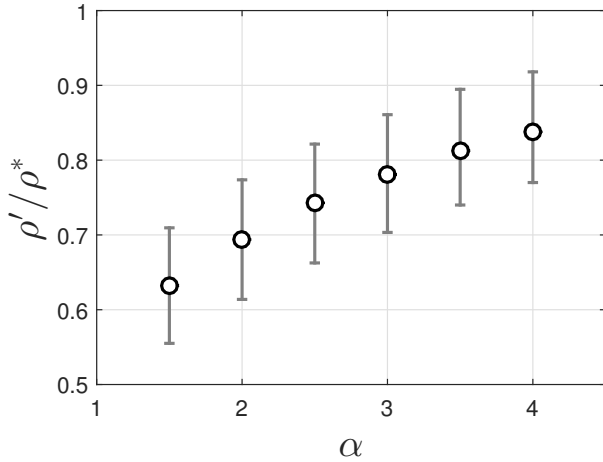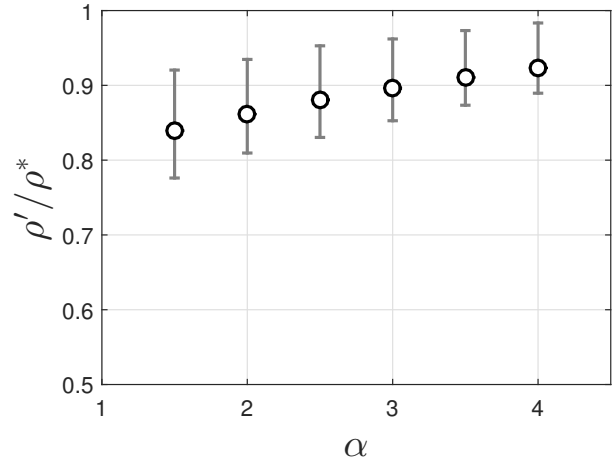
(a) Sectorized-Distributed Method   (b) Augmented Sectorized Method

Fig. 6: Ratio of achieved throughput versus optimal for varying path loss exponent with $T = 6$.

The capacity of edge $(i, j) \in E$ is given by (1). Therefore, for every unit of flow transversing (i, j), the schedule must activate the edge for a length of time $1/c_{(i,j)} = K^{-1} (d(i, j)/d_0)^{\alpha}$. Then, for increasing α, the edge with maximum length (and non-zero flow) requires infinitely more activation time to schedule its flow relative to the amount of time required by any other edge. This implies that as $\alpha \to \infty$, the schedule length required to schedule all units of flow in the network is determined by the length of time needed by this maximum length edge.

Now, let G' be the topology generated by Algorithm 1. Denote the optimal throughput scaling over topology G' as ρ', and let ρ^* be the optimal solution to (2). From Lemma 2, for any s,t pair, G' contains a path from s to t that minimizes the maximum edge length along the path. Suppose, that for all s,t we send $\Lambda_{s,t}$ units of flow over such a path. For traffic demand matrix Λ, such a strategy clearly minimizes the maximum edge length in the graph with non-zero flow. Thus, as $\alpha \to \infty$ the schedule length required to support Λ is minimized. Equivalently, if we restrict the schedule to be length 1 (c.f. constraint (2c)), the scaling of the traffic demand matrix, ρ', is maximized. Thus, ρ' approaches ρ^* as α becomes large. ∎

V. SIMULATION RESULTS

Theorem 3 indicates that for a sufficiently high path loss exponent, any topology generated by Algorithm 1 can support a throughput that achieves an arbitrarily close fraction of the optimal throughput of (2). We now explore the performance of the sectorized-distributed algorithm for varying path loss values, α, through simulation and empirically observe for different test cases and generated topologies the value of α_0 needed to get within $1 - \epsilon$ of optimal.

In particular, we consider a test scenario where there are 20 nodes distributed uniformly within a unit disk, edge capacity is given by (1) and $T = 6$. Each node is either a source

or a destination for one flow with demand $\Lambda_{s,t} = 1$. As in Section III we use the approximation method of [28] to approximate the capacity of the generated topologies. In Fig. 6a we plot statistics of the ratio of the observed throughput achieved via the sectorized-distributed method of Algorithm 1 versus the optimal throughput for 500 random node placements and demand matrix instantiations. The effect of α is tested at $\frac{1}{2}$ increments with the ∘ denoting the mean of the 500 tests and the bars showing the span of the 25^{th} to 75^{th} percentiles. The range of α shown covers path loss exponents in millimeter wave applications which have been empirically observed ranging in values from 3 to 4.5 [27].

The topologies generated by the sectorized-distributed algorithm are guaranteed to be connected, but are not necessarily dense. In order to improve the throughput performance of the generated topologies, we consider an augmented sectorized method that constructs a topology by first running Algorithm 1 and then greedily adds edges, in ascending order of edge length, until no edge may be added without violating the degree constraint at a node. [1] The results of Sections IV-B and IV-C clearly still hold for this method. In Fig. 6b, we plot the performance of this augmented algorithm. As can be seen from the figure, the augmented sectorized method significantly outperforms the plain sectorized algorithm and achieves a mean throughput of over 90% when $\alpha > 3$. The gap between the two methods narrows at high path loss, as both approach the throughput of the optimal solution.

In Fig. 7 we illustrate the effect of the degree constraint T on the achievable throughput of topologies generated by the augmented sectorized method. Specifically, we plot the achieved throughput of the topology relative to the maximum possible throughput of G (i.e., when there is no degree constraint and $T = \infty$ in (2)). We average over 100 test instantiations. The tests consist of 30 nodes distributed within

---

[1] Greedy algorithms similar to this may be obtained using distributed methods.
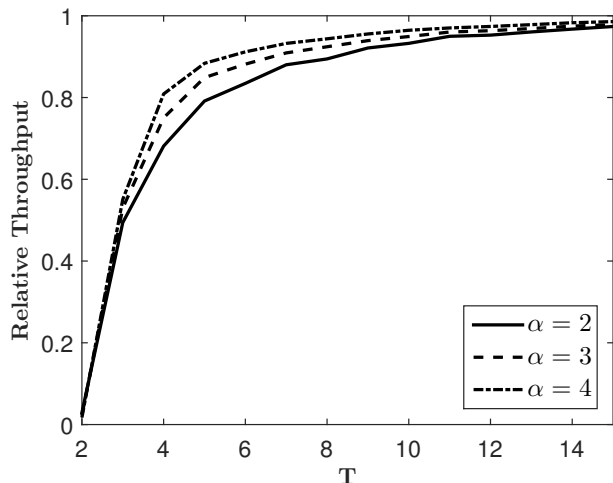
Fig. 7: Mean throughput of topologies generated by the augmented sectorized method, relative to the complete graph, as a function of T.

the unit circle, edge capacity given by (1), and with each node a source or destination for one commodity with unit demand. We see that for all α values the throughput increases most rapidly as T increases towards 6, the necessary condition for the results of Sections IV-B and IV-C. We further note that for larger α values, the benefit of increasing T beyond 6 is diminished. This observation is in line with Theorem 3 which indicates that for any $T \geq 6$, Algorithm 1 produces topologies that approach optimality for sufficiently large path loss exponents.

VI. CONCLUSION

In this paper, we examined the interacting effects of channel modeling, resource scheduling, and end-to-end traffic demands on mesh networks with highly-directional antenna systems. Because of the overhead introduced in tracking a neighbor's channel state, each node in the network can only maintain communication with a limited number of neighbors. This observation gives rise to a topology control problem, where each node must select its neighboring nodes subject to a degree constraint. We formulated the optimal topology design problem as a mixed-integer linear program, and developed a low complexity, distributed algorithm that produces connected topologies with good structural properties.

The problem formulation of Section II assumed a uniform tracking overhead per neighbor, and many of the results of Sections III and IV assumed channel models with capacity that decreased monotonically with edge length. In applications with non-uniform node mobility and shadowing environments these assumptions may not hold. Extensions to account for these effects are important directions for future research.

REFERENCES

[1] T. S. Rappaport, et al. "Millimeter wave mobile communications for 5G cellular: It will work!," *IEEE Access*, vol. 1, 2013, pp. 335-349.

[2] R. Mudumbai, S. Singh, and U. Madhow, "Medium access control for 60 GHz outdoor mesh networks with highly directional links." *IEEE INFOCOM*, 2009.

[3] Y. Zhu, et. al., "Cutting the Cord: a Robust Wireless Facilities Network for Data Centers," *Proc. MobiCom, ACM*, 2014.

[4] J.G. Andrews, et al. "What Will 5G Be?" *IEEE JSAC*, vol. 32, no. 6, May 2014, pp. 1065-1082.

[5] S. Rangan, T.S. Rappaport, and E. Erkip, "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proc. IEEE*, vol. 102, no. 3, 2014, pp. 366-385.

[6] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, Cambridge University Press, 2005.

[7] J. Wang, et. al., "Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems," *IEEE Selected Areas in Commun.*, vol. 27, no. 8, Oct. 2009, pp. 1390-1399

[8] E. Gelal, et. al., "Topology Management in Directional Antenna-Equipped Ad Hoc Networks," *IEEE Trans. Mobile Computing*, vol. 8, no. 5, May 2009, pp. 590-605.

[9] B. Mukherjee, "WDM-based local lightwave networks. II. Multihop systems," *IEEE Networks*, vol. 6, no. 4, July 1992, pp.20-32.

[10] J.-F. P. Labourdette and A. S. Acampora, "Logically rearrangeable multihop lightwave networks," *IEEE Trans. Commun.*, vol. 39, no. 8, Aug. 1991, pp.1223-1230.

[11] J.-F. P. Labourdette, "Traffic optimization and reconfiguration management of multiwavelength multihop broadcast lightwave networks," *Comput. Networks and ISDN Systems*, vol. 30, no. 9, 1998, pp.981-998.

[12] A. Narula-Tam, and E. Modiano. "Dynamic load balancing for WDM-based packet networks." *IEEE Proc. on INFOCOM*, 2000.

[13] M.G. Hluchyj and M. J. Karol, "Shuffle Net: An application of generalized perfect shuffles to multihop lightwave networks," *J. Lightw. Technol.*, vol. 9, no. 10, Oct. 1991, pp.1386-1397

[14] R. Ramanathan and R. Rosales-Hain. "Topology control of multihop wireless networks using transmit power adjustment." *IEEE Proc. on INFOCOM*, 2000.

[15] L. Li, et al.. "Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks." *ACM Proc. on Distributed Computing*, Aug. 2001, pp. 264-273.

[16] N. Li, J. C. Hou, and L. Sha. "Design and analysis of an MST-based topology control algorithm." *IEEE Trans. Wireless Commun.*, vol. 4, no. 3, May 2005, pp.1195-1206.

[17] D.M. Blough, et al. "The K-Neigh Protocol for Symmetric Topology Control in Ad Hoc Networks," *ACM Proc. on MobiHoc*, 2003.

[18] Z. Huang and C.-C. Shen, "Multibeam antenna-based topology control with directional power intensity for ad hoc networks," *IEEE Trans. Mobile Computing*, vol. 5, no. 5, May 2006, pp. 508-517.

[19] V. Namboodiri, L. Gao, and R. Janaswamy, "Power Efficient Topology Control for Wireless Networks with Switched Beam Directional Antennas," *IEEE Mobile Adhoc and Sensor Syst. Conference*, 2005.

[20] U. Kumar, H. Gupta, and S. R. Das, "A Topology Control Approach to Using Directional Antennas in Wireless Mesh Networks," *IEEE ICC 2006*.

[21] A. Subramanian, et. al., "Topology Control Protocol Using Sectorized Antennas in Dense 802.11 Wireless Networks," *IEEE ICPN*, 2009.

[22] A. Goldsmith, *Wireless Communications*, Cambridge University Press, 2005.

[23] S. Singh, R. Mudumbai, and U. Madhow, "Interference Analysis for Highly Directional 60-Ghz Mesh Networks: The Case for Rethinking Medium Access Control," *IEEE/ACM Trans. on Networking (TON)*, vol. 19, no. 5, Oct. 2011, pp.1513-1527.

[24] H. Shokri-Ghadikolaei, L. Gkatzikis, and C. Fischione, "Beam-searching and Transmission Scheduling in Millimeter Wave Communications." Preprint.

[25] K. Jain, et al. "Impact of interference on multi-hop wireless network performance," *Wireless Networks*, Springer, vol. 11, no.4, 2005, pp.471-487.

[26] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows. Theory, Algorithms, and Applications*, Prentice Hall, 1993.

[27] G. R. MacCartney Jr, et al., "Path Loss Models for 5G Millimeter Wave Propagation Channels in Urban Microcells." *IEEE 2013 Global Commun. Conference*, 2013.

[28] M. Kodialam and T. Nandagopal, "Characterizing achievable rates in multi-hop wireless networks: the joint routing and scheduling problem," *ACM Proc. of MobiCom*, 2003.