

Topology Control Protocol Using Sectorized Antennas in Dense 802.11 Wireless Networks

Anand Prabhu Subramanian*

Bell Labs, Alcatel-Lucent

Murray Hill, NJ, USA

Anand_Prabhu.Subramanian@alcatel-lucent.com

Henrik Lundgren, Theodoros Salonidis

Thomson

Paris, France

firstname.lastname@thomson.net

Don Towsley

University of Massachusetts

Amherst, MA, USA

towsley@cs.umass.edu

Abstract—We introduce a measurement-based optimization framework for topology control in dense 802.11 networks using sectorized antennas. We first formulate a topology control optimization problem, where nodes activate their sectorized antenna patterns to minimize network interference and maximize network capacity. In contrast to previous approaches, our formulation is based on a physical interference model that relies on Received Signal Strength (RSS) measurements on multiple antenna patterns of each link. We then introduce a distributed measurement protocol to measure RSS of these antenna patterns and a greedy distributed topology control protocol that uses this information to achieve topologies of minimal interference. The protocols are experimentally evaluated on a dense 802.11 wireless testbed. Extensive measurements show that the protocols operate very close to optimal and yield significant increase in network throughput compared to omni-directional antennas. These results hold for various traffic scenarios in both wireless LAN and mesh network configurations.

I. INTRODUCTION

Recently, 802.11 wireless networks have gained significant momentum due to potential high data rate connectivity, very low deployment cost, and free or low-cost network access to users. 802.11 networks, however, suffer from serious interference problems limiting their capacity due to the broadcast nature of the wireless medium and their use of the unlicensed spectrum band. Sectorized antennas is an attractive solution to reduce wireless interference through higher spatial reuse. Such antennas can concentrate radio signal energy in a particular direction, instead of radiating it in all directions like their omni-directional counterpart.

In current literature, there exist primarily two approaches to exploit directional communication: directional MAC protocols and topology control protocols. Directional MAC protocols [3], [9] aim to achieve higher spatial reuse by switching between different sectors to communicate to different neighbors at a very fast (per-packet) time scale. They also require modifications to the 802.11 MAC protocol to combat deafness and directional hidden terminal problem [4]. On the other hand, topology control protocols [5], [6], [10], [14] activate multiple sectors simultaneously at slower time scales (in the order of minutes or more) at the potential expense of performance, but do not require complex modifications to

the standard 802.11 MAC protocol. These topology control solutions use simplistic interference model assumptions (such as pairwise interference) and use adaptations of graph theoretic algorithms. They ignore several practical considerations and are evaluated only through simulation studies.

On the other hand, recent measurement studies on densely deployed 802.11 networks [2], [13] have shown that common design assumptions such as usage of geographically pointing sectors and sector independence on multi-sector activation are sub-optimal in reflection-rich dense deployments. These studies have also investigated the effect of physical layer capture on reducing directional hidden terminal problem and have shown that sectorized antennas can provide higher spatial reuse opportunities. However, these observations have not been translated to practical protocols that can extract the capacity gains of antenna sectorization in real-world dense wireless environments.

In this paper, we introduce an optimization framework and develop a practical and efficient topology control protocol using sectorized antennas that leverages on previous experimental observations. The main distinction between our topology control approach and earlier proposals is that we carry out a measurement based approach for topology control using low-cost sectorized antennas that are simple to design and can work with commodity 802.11 wireless hardware.

We first formulate a topology control optimization problem as a quadratic integer program that activates multiple sectors on each node to minimize overall network interference and thereby maximize capacity. We then introduce a technique to reduce it to an equivalent linear integer program. The resulting linear integer program can be relaxed to obtain bounds on the optimal solution using standard LP solvers. Our formulation uses a physical interference model based on measurements of Received Signal Strength (RSS) values, which have been observed to capture the performance of different sector activations well [2], [13]. This formulation also avoids modeling the detailed operation of the 802.11 MAC protocol which is hard and cannot be easily incorporated in optimization procedures [8], [12].

We then propose a greedy distributed topology control protocol and a distributed measurement protocol. In the topology control protocol, nodes use a greedy heuristic to iteratively solve the topology control optimization problem based on local

*This work was undertaken while the first author was interning at Thomson, Paris; at that time, he was a PhD student at Stony Brook University, NY, USA.

one-hop measurements of RSS values of the antenna patterns on their adjacent links. The measurement protocol ensures that the most up to date RSS values on all antenna patterns are measured and provided to the topology control protocol, by properly coordinating node transmissions and reception on different antenna patterns. Both protocols can be readily implemented using commodity 802.11 hardware and operate transparently with the 802.11 MAC protocol and network layer link selection techniques, such as routing protocols in multi-hop wireless mesh networks or association mechanisms in WLANs.

Finally, we perform a detailed experimental evaluation of our greedy topology control protocol in an indoor 802.11 wireless testbed in a variety of network and traffic scenarios and show that it yields a near-optimal solution in all cases. The topology control protocol reduces both sender and receiver side interference and results in significant improvement in spatial reuse compared to using omni-directional antennas. We also evaluate the tradeoff between RSS measurement complexity of the measurement protocol and throughput gain achieved by the topology control protocol due to spatial reuse. We finally show that network density plays a critical role in the amount of spatial reuse that can be leveraged using antenna sectorization. The results collectively show that carefully engineered topology control protocols based on antenna sectorization without complex modifications to the 802.11 MAC protocol can provide substantial throughput gains in real world dense 802.11 wireless networks.

The paper is structured as follows. Section II surveys state of the art. Section III introduces the formulation of the topology control optimization problem and Section IV provides a Linear Programming (LP) technique to obtain a lower bound of the optimal solution. Section V and Section VI are devoted to the distributed topology control protocol and distributed measurement protocol, respectively. In Section VII, we experimentally evaluate the topology control optimization framework composed by these two protocols in a 802.11 wireless testbed. Section VIII concludes.

II. RELATED WORK

A majority of past work on the use of sectorized antennas and topology control in 802.11 based wireless networks has focused on analytical studies, and protocol design and evaluation through simulations. Kumar et al. [10] uses an antenna model consisting of multiple single-sector antennas and propose a centralized algorithm to construct a minimum degree spanning tree in a mesh network. Huang et al. [5], [6] propose two different topology control approaches which use transmission power adjustments and multi-sector antennas. They assume simplified antenna and network interference models, and propose a graph coloring based approach. Finally, Sundaresan et al. [14] consider the problem of determining the (antenna) communication pattern to be used by nodes in a mesh network and formulate it as interference minimization when using “smart antennas” (digital adaptive arrays). All the current approaches for topology control uses idealized

antenna models, assume pair-wise interference model, propose solutions based on graph theoretic algorithms, and perform simulation-based evaluations.

In contrast to previous work, we take an experimental approach to topology control using multi-sector antennas in real world dense deployments. We formulate a topology control optimization problem based on insights from experimental studies in current literature [2], [13]. Similar to Sundaresan et al. [14], we aim to minimize the overall interference in the network. On the contrary, our formulation uses physical interference models based on measured received signal strength values, rather than pairwise interference models. Furthermore, we do not rely on idealized antenna models with perfect radiation patterns. Instead, we adopt a measurement-based approach for characterizing the antenna sectors. We also contrast previous work as we do not rely on advanced antennas or digital signal processing, but instead use a simple and low-cost multi-sector antenna that can be readily used with regular off-the-shelf IEEE 802.11a hardware.

III. PROBLEM FORMULATION

In this section, we present our system model and formulate the topology control problem using multi-sector antennas as a combinatorial optimization problem.

A. Network Model

We represent the network by a directed graph $G = (V, E)$ where V is the set of N nodes in the network and $E \subseteq V^2$ is a set of links between pairs of nodes. More precisely, the link (u, v) exists between node pair u and v , $(u, v) \in E$ if data can be delivered from u to v with a delivery ratio above a threshold (PDR_{thresh}) when the link operates at the lowest bit rate supported¹ by the underlying physical layer. This ensures that each link in the network can carry data using at least one of the possible bit rates.

We say a link $(u, v) \in E$ is *active*, if it is selected to carry data traffic. Link selection is performed by a network layer protocol, such as the association mechanism between an AP and client in a WLAN, or a multi-hop routing protocol in a mesh network (e.g., AODV, DSR). In this paper, we do not consider the link selection procedure itself, but assume the presence of such a mechanism. As we show later, our topology control algorithm does not impact link selections based on link quality metrics such as ETX and ETT. Let $E' \subseteq E$ denote the set of *active* links in the network. Set E' is an input parameter to our topology control optimization problem.

B. Antenna Model

We assume each node in the network is equipped with a multi-sector antenna consisting of s sectors. An *antenna pattern* is formed by activating any combination of these s sectors. The total number of antenna patterns that can be formed is $k \leq 2^s - 1$. Let $K = \{0, 1, \dots, k - 1\}$ denote the set of possible antenna patterns available at each node. Here antenna pattern

¹The lowest supported bit rate is 1 Mbps in 802.11b mode and 6 Mbps in 802.11a mode

0 denotes the omni-directional mode where all s sectors are enabled. An important property of our multi-sector antenna model is that the maximum gain in omni-directional mode equals the maximum gain in single or multi-sector mode. This property eliminates the negative effects caused by directional hidden terminal problem due to asymmetry in gain [4]. Current solutions to eliminate this problem use multi-hop RTS/CTS packets, which are hard to implement in commodity 802.11 hardware. The multi-sector antennas considered in this paper thus primarily target interference reduction along the side lobes rather than improving the link budget along the main lobe.

Let $G^{i\theta}$ denote the gain of antenna (in dBi) when using antenna pattern i along the direction θ . Under ideal conditions, in the absence of multi-path reflection, the received signal strength (RSS) at a node v from a node u when node u is using antenna pattern i to transmit packets and node v is using antenna pattern j to receive the packets is given by,

$$RSS_{uvij} = \frac{P_u^t G_u^{i\theta_{uv}} G_v^{j\theta_{vu}}}{d_{uv}^\alpha}$$

where P_u^t is the transmit power at node u , θ_{uv} is the direction from node u to v , d_{uv} is the distance between u and v , and α is the path loss coefficient. In practice, RSS is an estimate of the signal energy level at the receiver during packet reception and is measured when receiving the PLCP packet header. Whenever a new antenna pattern is selected, the antenna gain changes with a resulting change in the energy level at the receiver. RSS is thus the fundamental physical layer quantity affected by this change. Using commodity wireless cards, the received signal strength of any link (u, v) can be directly obtained from the prism monitoring header added to each received packet. It has been shown in previous studies [2], [13] that RSS correlates well with higher layer metrics like packet delivery ratio and data rate of a link. RSS also forms the basis of the physical interference model. In the next section, we discuss how to use link RSS to model interference in the network and define the objective function of our topology control optimization problem.

C. Interference Model and Objective Function

The obvious approach to topology control is to maximize network capacity. However, we are concerned with networks that use CSMA and in contrast to a TDMA-based network, modeling network capacity of such a network is difficult. However, recently, several measurement-based capacity models have been proposed [8], [12]. These models are complex and are not easily incorporated in optimization frameworks.

Our approach to topology control is to minimize the *aggregate interference* in the network (defined later). This has been used [7] as an indirect way to optimize network capacity. In CSMA based wireless networks, a transmission suffers from interference both at the sender and receiver side. Consider a transmission from node u to node v . At the sender side, u senses the medium free only when the following condition is

satisfied.

$$\sum_{w \in I(u)} RSS_{wu} \leq CS_{thres} \quad (1)$$

where RSS_{wu} is the received signal strength at node u from node w , $I(u)$ is the set of nodes that can cause interference at node u (denoted as the *interference set*) and CS_{thres} is the carrier sense threshold. On the receiver side, a packet transmitted from v to u is successfully received only if the following condition is satisfied.

$$\frac{RSS_{vu}}{\sum_{w \in I(u)} RSS_{wu} + N} \geq \beta \quad (2)$$

Here N is the background noise and β is the receive threshold.

We define the interference set $I(u)$ for a node u as follows. A node $w \in I(u)$ iff (u, w) and $(w, u) \notin E'$. This implies that all nodes in the network that do not transmit to or receive from node u according to the links selected by the network layer protocol are part of the interference set. The interference at node u then is $\sum_{w \in I(u)} RSS_{wu}$. This affects node u when it is either the sender node of a link (Equation 1) or the receiver node of a link (Equation 2).

The focus of our approach is to minimize the aggregate interference in the network, which is defined as follows:

$$\sum_{u \in V} \sum_{w \in I(u)} RSS_{wu} = \sum_{(w, u) \notin E'} RSS_{wu} \quad (3)$$

D. Optimization Framework

Informally, our approach to topology control is to minimize the aggregate interference in the network as defined in Equation (3) by assigning appropriate antenna patterns to each node subject to a connectivity constraint. Our connectivity constraint is defined for each link in E' (the links that carry data traffic) as follows. *For each link (u, v) in the set E' , the RSS of that link after the antenna pattern assignment should be comparable (within a threshold of a few dB) to the RSS of the link when both node u and node v use omni-directional mode.* We define the optimization problem formally below.

Let X_{ui} be a decision variable that denotes whether or not antenna pattern i is assigned to node u . $X_{ui} = 1$ if antenna pattern i is assigned to node u and $X_{ui} = 0$ otherwise. For each link (u, v) in the set E' , let C_{uvij} denote whether assigning antenna pattern i to node u and j to node v satisfies the connectivity constraint.

$$C_{uvij} = \begin{cases} 1, & \text{if } RSS_{uv00} - RSS_{uvij} < C_{th} \\ 0, & \text{otherwise} \end{cases}$$

where RSS_{uv00} denotes the RSS of link (u, v) when both nodes use omni-directional mode and RSS_{uvij} denotes the RSS when node u uses antenna pattern i and node v uses antenna pattern j . C_{th} is a threshold chosen to be of the order of a few dB. Note that in our antenna model, omni-directional mode has the largest gain among all antenna patterns. Choosing antenna patterns i and j such that $C_{uvij} = 1$ ensures that the data rate of the link (u, v) does not suffer because of antenna pattern assignment. This ensures the links

selected by the network layer protocol do not suffer poor performance due to our antenna pattern assignment. In a high level, the goal of our topology control approach is to maintain the quality of the active links (links that carry traffic) and reduces the interference between the active links by efficient antenna pattern assignment.

Our topology control problem is defined as follows:

$$\begin{aligned} \text{minimize} \quad & \sum_{(w,u) \notin E'} \sum_{i,j \in K} RSS_{wuij} \cdot X_{wi} \cdot X_{uj} \\ & \sum_{i \in K} X_{ui} = 1, \forall u \in V \end{aligned} \quad (4)$$

$$\sum_{i,j \in K} C_{uvij} \cdot X_{ui} \cdot X_{vj} \geq 1, \forall (u,v) \in E' \quad (5)$$

$$X_{ui} \in \{0, 1\}, \forall u \in V, \forall i \in K \quad (6)$$

Equation (4) ensures that each node in the network is assigned exactly one antenna pattern. Equation (5) ensures that the RSS, using the assigned antenna patterns, of each link carrying data traffic is within C_{th} of the RSS when using omni-modes. Finally, Equation 6, ensures that X_{ui} 's take values 0 or 1.

The above optimization problem is a quadratically constrained quadratic optimization problem and is known to be NP-hard. In the next section, we reduce this problem to a integer linear program and then relax it to a linear program to obtain a lower bound on the optimal solution.

IV. LINEAR PROGRAM FORMULATION

Our goal in this section is to replace the quadratic objective and quadratic constraint in our formulation with a linear objective and constraint. Note that the objective function only includes links that do not belong to the set E' and constraint (5) only applies to links only belonging to set E' . These two sets of links are disjoint. We replace the product $X_{ui} \cdot X_{vj}$ by a new variable Z_{uvij} in equation (5) and introduce two new constraints,

$$Z_{uvij} \leq X_{ui}, \forall u, v \in E', \forall i, j \in K \quad (7)$$

$$Z_{uvij} \leq X_{vj}, \forall u, v \in E', \forall i, j \in K \quad (8)$$

In the original optimization formulation, $X_{ui} \cdot X_{vj} = 1$ when both X_{ui} and X_{vj} are 1 and 0 otherwise. Equations (7) and (8) ensure that Z_{uvij} is less than both X_{ui} and X_{vj} which ensures that Z_{uvij} must be 0 when either X_{ui} or X_{vj} is 0. When both X_{ui} and X_{vj} are 1, it can either be 0 or 1. Since the left side of equation (5) should be greater than 1, at least one of the Z_{uvij} will be forced to be 1. This ensures that the connectivity constraint is satisfied for every link carrying data.

We then replace the product $X_{ui} \cdot X_{vj}$ by the variable Y_{uvij} in the objective function and add the following additional constraints.

$$Y_{uvij} \geq X_{ui} + X_{vj} - 1, \forall u, v \notin E', \forall i, j \in K \quad (9)$$

In the original optimization formulation, $X_{ui} \cdot X_{vj} = 1$ in the objective function when both X_{ui} and X_{vj} are 1 and 0 otherwise. Equation (9) ensures that when both X_{ui} and X_{vj} are 1, then Y_{uvij} is forced to be 1. In the remaining cases, Y_{uvij} can either be 0 or 1. Since the optimization problem aims to minimize the objective, when given a choice of 0 or 1 to the variable Y_{uvij} , 0 is chosen.

Using the above mentioned reductions, the original quadratically constraint quadratic optimization problem transforms into the following integer linear program.

$$\text{minimize} \quad \sum_{(w,u) \notin E'} \sum_{i,j \in K} RSS_{wuij} \cdot Y_{wuij}$$

$$\sum_{i \in K} X_{ui} = 1, \quad \forall u \in V$$

$$\sum_{i,j \in K} C_{uvij} \cdot Z_{uvij} \geq 1, \quad \forall (u,v) \in E'$$

$$Y_{uvij} \geq X_{ui} + X_{vj} - 1, \forall u, v \notin E', \forall i, j \in K$$

$$Z_{uvij} \leq X_{ui}, \forall u, v \in E', \forall i, j \in K$$

$$Z_{uvij} \leq X_{vj}, \forall u, v \in E', \forall i, j \in K$$

$$X_{ui} \in \{0, 1\}, \forall u \in V, \forall i \in K$$

$$Y_{uvij} \in \{0, 1\}, \forall u, v \notin E', \forall i, j \in K$$

$$Z_{uvij} \in \{0, 1\}, \forall u, v \in E', \forall i, j \in K$$

We can obtain lower bounds on the optimal solution in polynomial time by relaxing the integrality constraints in the above integer linear program using standard solvers such as CPLEX [1].

V. GREEDY DISTRIBUTED TOPOLOGY CONTROL PROTOCOL

We present a distributed topology control protocol that aims to solve the topology control optimization problem. The protocol works transparently with existing link selection mechanisms (routing protocols such as AODV and DSR in mesh networks or association mechanisms in WLANs) and can be implemented on top of the 802.11 MAC protocol using commodity 802.11 wireless hardware. The protocol uses as input the active links E' from the link selection mechanism and the RSS_{wuij} values from a distributed measurement protocol (Section VI). We first present the algorithm that contains the idea behind the protocol.

A. Distributed Iterative Algorithm

The topology control optimization problem described in the previous section is extremely complex to solve even with centralized information. We present a distributed heuristic that aims to iteratively solve this problem using only local information. Consider an arbitrary assignment where $x_u \in K$ denotes the antenna pattern assigned to each node u in the network. For this assignment, the interference at node u is:

$$\sum_{(w,u) \notin E'} RSS_{wux_u x_u} \quad (10)$$

This objective function can also be written as $\sum_{u \in V} S_u$ where:

$$S_u = \sum_{w:w \in I(u)} RSS_{uw x_w x_u} + \sum_{w:u \in I(w)} RSS_{uw x_u x_w} \quad (11)$$

The first term of S_u corresponds to the interference other nodes w cause to node u and the second term corresponds to the interference node u causes to other nodes w . The iterative algorithm aims to minimize the objective function by iteratively minimizing each S_u .

At each iteration n , the algorithm operates as follows. Let $(x_u^{(n)}, \forall u \in V)$ denote the antenna pattern assignments of all nodes at iteration n . A node u is selected to modify its antenna pattern assignment and all its one-hop neighbors $\{w : w \in I(u), u \in I(w)\}$ and $\{v : (u, v) \in E'\}$ keep their antenna pattern assignments $x_w^{(n)}$ and $x_v^{(n)}$ fixed. The node selection order can be with respect to different criteria but for simplicity we assume it is arbitrary. Node u computes the sector x_u^* that minimizes $S_u(x_u)$ in equation (11) subject to the constraints:

$$\begin{aligned} RSS_{uv00} - RSS_{uvx_u^*x_v^{(n)}} &\leq C_{th}, \quad (u, v) \in E' \\ RSS_{vu00} - RSS_{vux_v^{(n)}x_u^*} &\leq C_{th}, \quad (v, u) \in E' \end{aligned}$$

This minimization includes (i) selecting all antenna patterns x_u that satisfy connectivity constraints of links (u, v) in the set E' , i.e., $C_{uvx_u x_v} = 1$, and (ii) for these x_u , compute and compare the corresponding values for S_u based on Equation (11). Finally, node u selects x_u^* as its new antenna pattern assignment.

It is easy to show that the above algorithm is guaranteed to converge to a local minimum of the objective function of Equation (10), after a finite number of iterations, regardless of initial state. The reason is that during each iteration the objective function decreases and there is a finite number of antenna pattern combinations in the network.

B. Greedy Distributed Protocol

We now provide a distributed protocol that realizes a greedy version of the distributed iterative algorithm. The protocol is executed continuously and asynchronously at each node u and relies only on local computations and one-hop message exchanges. It operates on top of the 802.11 MAC and its messages are implemented with broadcast or unicast network layer packets. All messages are sent using the omni-directional antenna pattern and at the lowest data rate (6 Mbps in our 802.11a implementation) to increase coverage and reliability, respectively.

Each node u maintains a table for each neighbor w that contains (i) current antenna pattern x_w and (ii) $RSS_{uw x_w x_u}$ values for all antenna sectors on link (u, w) , measured by a separate distributed measurement protocol described in the next section. Each node u periodically sends a broadcast packet containing its current antenna pattern assignment x_u so that neighbors update their tables.

Each node u is in two states with respect to its participation to the topology control protocol. During an antenna

pattern modification it is BUSY and uses the omni-directional antenna pattern, otherwise it is FREE and uses its assigned antenna pattern x_u . Upon entering the BUSY state, a node attempts to modify its antenna pattern assignment. Since the attempts among nodes are not synchronized, multiple nodes may attempt to modify their antenna patterns simultaneously. Recall that each iteration of the distributed algorithm requires a node to modify its assignment while its one-hop neighbors keep their assignments fixed. This issue is addressed using a distributed locking mechanism described next.

We now describe the protocol actions during the attempt of node u . In the following description, when we say that a node u (or w) *becomes BUSY*, it performs the following actions: it switches to BUSY state, switches to omni-antenna pattern 0, and sets a timeout equal to BUSY_TIMEOUT which is the maximum duration nodes participate in each attempt. When a node u (or w) *aborts*, it switches to FREE state and returns to its antenna pattern assignment x_u (or x_w).

Step 1. Node u first *becomes BUSY*. Then it minimizes S_u in Equation (11) based on the x_w values in its local table and obtains the minimum S_u^* and corresponding assignment x_u^* . Then, it sends a broadcast TOPO_REQ packet that contains S_u^* . It also sets a timeout TOPO_REQ_TIMEOUT during which it waits for TOPO_REQ_ACK responses from all neighbors.

Step 2. When a node w receives a TOPO_REQ packet from u , it returns a TOPO_REQ_NACK packet to u if it is BUSY or if $S_w^* \leq S_u^*$. Otherwise, if it is FREE and $S_w^* > S_u^*$, it *becomes BUSY* and returns a (unicast) TOPO_REQ_ACK packet to u . The comparison between S_u^* and S_w^* serves both to break ties and to implement a greedy iteration where only the node with minimum S^* within its one-hop neighborhood modifies its sector assignment. This results in maximum reduction of the global objective within each one-hop neighborhood.

Step 3. If node u receives at least one TOPO_REQ_NACK or its TOPO_REQ_TIMEOUT expires, it broadcasts a TOPO_CANCEL packet and *aborts*. Otherwise, if it receives TOPO_REQ_ACK from all neighbors within TOPO_REQ_TIMEOUT, it broadcasts a TOPO_CONFIRM packet. Then, it switches to FREE state and switches to the new antenna pattern assignment x_u^* .

Step 4. If node w receives a TOPO_CONFIRM packet assignment from u , it switches to FREE state and switches back to its antenna pattern assignment x_w .

At this point the antenna pattern reassignment at node u is complete. If during the above steps the BUSY_TIMEOUT timer of any BUSY node expires, or a node receives a TOPO_CANCEL packet, this node *aborts*.

VI. DISTRIBUTED MEASUREMENT PROTOCOL

We present a distributed measurement protocol that aims to measure the RSS values $RSS_{u i_m j_m}$ for all links (u, v) in the network, given a pre-specified set of M link antenna patterns $\mathcal{M} = \{(i_m, j_m), m = 1, \dots, M\}$, where i_m and j_m take values in $\{0, 1, \dots, 2^s - 1\}$. In the special case where $M = K^2 = (2^s - 1)^2$, the measurement protocol seeks to measure all

possible antenna patterns for all links in the network. Before describing the protocol, we first illustrate the idea behind it using an idealized algorithm.

A. Algorithm

Time is organized in N cycles, where N is the number of nodes in the network. In each cycle a single node transmits u broadcast probe packets and all other nodes v act as receivers. All nodes are perfectly synchronized and know the transmission times of the probe packets and which antenna patterns to use during each transmission. Each cycle consists of M rounds where in the m -th round node u transmits M broadcast packets on antenna patterns i_1, \dots, i_M while all other nodes are tuned in antenna pattern j_m . Thus, under ideal conditions of perfect coordination and no packet losses, the RSS_{uv, i_m, j_m} of all links can be measured in NM^2 measurements.

Translation of this algorithm to an operational protocol entails several challenges. Since the protocol should operate above the 802.11 MAC protocol, no tight synchronization on packet transmission and sector activation timings is available. This imposes limitations on coordinating nodes to send probes at the right time and on the right antenna patterns. In addition, the protocol should account for maintenance of coordination under packet losses and simultaneous transmissions and receptions among different nodes.

B. Protocol

The measurement protocol is executed continuously at each node u and aims to measure the RSS values on all antenna patterns $\mathcal{M} = \{(i_m, j_m), m = 1, \dots, M\}$. It operates on top of the 802.11 MAC and its probes are implemented with broadcast network layer packets. All probes are sent at the lowest data rate (6 Mbps in our 802.11a implementation) to increase reliability. During the protocol, each node u performs different transmitter and receiver actions, described next.

Transmitter operation. As transmitter, each node u continuously executes measurement cycles. Each measurement cycle consists of M broadcast probe packets, each sent to the MAC protocol buffer every T_M seconds at a pre-defined antenna pattern sequence $TXSEQ = (i_1, i_2, \dots, i_M)$. This sequence is an ordering of the TX antenna patterns $\{i_m\}$ of the set $\mathcal{M} = \{(i_m, j_m), m = 1, \dots, M\}$. The parameters M , T_M and $TXSEQ$ are fixed and known to all nodes in the network. They help remove the timing and antenna pattern uncertainty at the receivers.

Each probe of node u informs neighbor nodes about its own Tx antenna pattern and approximate time of arrival of the next probes. More specifically, the m -th probe contains (i) its Tx antenna pattern i_m (ii) its sequence number m in TXSEQ, and (iii) a timestamp t_m of the time it was buffered at the MAC queue of node u (in an implementation over 802.11 MAC we cannot know when a probe is actually transmitted over the air). This timestamp aids the neighbors of node u to coarsely synchronize their clocks with the clock of node u and estimate when it buffers its probe packets. The 802.11 MAC delay is random and consists of queuing delay and MAC

contention delay. In our MadWiFi implementation we removed the queuing delay by prioritizing the broadcast probe packets. The resulting MAC delay is in the order of a few milli-seconds, which determines the timestamp accuracy.

Receiver operation. As receiver, each node u must coordinate with the timing and antenna patterns of transmitted probes. A protocol parameter known to all nodes is the minimum dwell period $T_{D_{min}}$, which is the minimum required time to receive probe packets at a specific antenna pattern. $T_{D_{min}}$ includes (i) the delay for node u to switch the antenna pattern at the driver level, (ii) the packet transmission durations of two probe packets (one for receiving the probe of another node v and one for itself when it acts as transmitter) and (iii) the timestamp accuracy.

When node u receives a probe packet from each node v for the first time, it determines the antenna patterns and approximate start times of all future probe packets of v , based on M , T_M , $TXSEQ$ and the contents of the probe packet. Let t_m be the start time when node u enters dwell period m and activates antenna pattern j_m for a duration T_{D_m} . During this dwell period node u may receive probe packets from multiple neighbor nodes v . Let $i_m(v)$ be the antenna pattern in the probe sent by a node v during this dwell period. At the end of the dwell period, node u acts as follows.

Step 1: Node u extracts the RSS of the probe and uses it to update its moving average estimate of $RSS_{u,v, i_m(v), j_m}$.

Step 2: Node u adjusts its estimate of the start time t_{m+1} of the next dwell period using the earliest timestamp seen in the received probe packets. Then, the dwell period $T_{D_{m+1}}$ is determined as the maximum of $T_{D_{min}}$ (to account for the implementation constraints) and the difference of the latest timestamp and t_{m+1} (to account for all potentially received probes).

Step 3: Due to the global sequence TXSEQ, u already knows the next antenna patterns $i_{m+1}(v)$ of the nodes v . Node u selects the antenna pattern $j_{m+1} = j$ for its least recently received pattern $(i_{m+1}(v), j)$. This selection ensures that eventually all patterns will be measured.

Step 4: Node v schedules a timer to enter dwell period $m+1$ at sector j_{m+1} based on the computed t_{m+1} and $T_{D_{m+1}}$.

Node u may transmit or receive packets from any other nodes during the dwell period. The dwell period should be short enough to ensure that the measurement protocol will not operate using non-optimized sectors for extended periods of time but long enough to ensure that the probes will be received using the right sector configuration. In our implementation we use $T_{D_{min}} = 100\text{ms}$ and $T_M = 2\text{s}$, which corresponds to 5% operation in non-optimized antenna patterns.

C. Reducing measurement complexity

We consider two measurement methods. Method 1 measures all antenna patterns for each link and requires a number of measurements that grows exponentially with the number of sectors. Method 1 requires $M = (2^s - 1)^2$ measurements per link. Method 2 provides decreased measurement complexity at the potential cost of loss in accuracy. It predicts the

RSS of multi-sector patterns of each node based only on measurements of single-sector patterns. More specifically, for a given set of sectors, the RSS of their multi-sector pattern is approximated by the maximum RSS of the measured single-sector patterns. This method ignores multi-path reflections and assumes that sector combinations are superpositions of the individual sectors. Method 2 requires $M = s^2$ measurements per link.

Both methods work with the measurement protocol by replacing M with the appropriate number of measurements and having nodes being aware of which sector patterns they need to measure. We evaluate the accuracy of these methods in the experimental evaluation section that follows.

VII. EXPERIMENTAL EVALUATION

We first describe the experimental methodology, including testbed setup, sectorized antenna characteristics, and default experiment configurations. Next, we present the experimental results that quantify the performance gains of our approach in a real-world dense deployment under different scenarios.

A. Experimental Methodology

1) *Testbed Setup*: We deploy a six-node testbed in an indoor office environment as shown in Figure 2. This environment is rich in multi-path reflections, typical in dense home or enterprise WiFi deployments and community wireless mesh networks. Each node is a Dell D610 laptop running Linux (Fedora Core 6) and uses the Madwifi driver (version 0.9.4) to control an Atheros IEEE 802.11 a/b/g wireless card.




Fig. 2. A map of the testbed deployment. Dark squares show the node locations and the orientation of their antennas.

We equip the wireless card at each node with a Vivaldi multi-sector antenna [11], shown in Figure 1(a). This antenna can be used with any commodity IEEE 802.11 wireless card operating in the 5 GHz band. Each of the four sectors covers a different quarter of the azimuth plane. The sectors at each node can be simultaneously activated in any combination, thus allowing 15 different antenna patterns. Omni-mode is achieved by simultaneously activating all four sectors. Figures 1(b) and 1(c) show the radiation patterns in the azimuth plane for a single sector and omni-mode, respectively. Sector activation is controlled by a switch integrated on the antenna system, which we control through the parallel port of the host computer.

2) *Experiment Configurations*: We perform all our experiment using channel 52 in the 802.11a 5GHz band during night-time and without external interference. Unless specified otherwise, each node uses 1 dBm Tx power, 6 Mbps data rate and a 3 dB C_{th} threshold. We use a UDP traffic generator in backlogged mode to generate 1024-byte back-to-back UDP packets and collect measurement data using Tcpdump. In addition to the above default settings, in some experiments we vary the Tx power and C_{th} threshold, and switch between unicast and broadcast data packets to study the impact of different network densities, sector assignments, and data traffic.

Out of the 15 node pairs in our testbed, six have 100% delivery ratio in both directions. These links form part of the set E in our problem formulation. Most of our experimental configurations consist of 'two-link' topologies. In each 'two-link' topology we select two links from set E that do not share a node. We also perform experiments on both WLAN and wireless mesh network configurations, where each node acts as a client/AP or mesh router, respectively. All non-default and additional parameter settings and network configurations are detailed along with the results in the following section.


B. Experimental Results

We evaluate and compare three sector assignment schemes. In the 'Greedy' scheme, we assign antenna patterns to all the nodes using our greedy topology control protocol. In the 'Omni' scheme, all nodes use omni-directional mode. In the 'Optimal' scheme, we perform an exhaustive search (15^6) to obtain the optimal antenna pattern assignments; the 'Optimal' scheme is a useful benchmark for the efficiency of the 'Greedy' scheme.


We first study both sender side and receiver side interference reductions. Next, we study aggregate broadcast throughput performance under varying network density, and then unicast throughput. We then investigate the impact of varying C_{th} and different sector measurement schemes. Finally, we investigate the performance under WLAN and mesh network scenarios.

1) *Sender side interference reduction*: In a CSMA network, a sender node defers when it senses the carrier signal at a level above the carrier sense threshold. RSS reduction at the sender from interfering nodes ultimately leads to increased medium access probability and sending rate for these nodes.


For each of the 20 two-link topologies, we measure at each sender node the RSS from the interferer node (the sender of the other link), using the three sector assignment schemes. Figure 3 shows the RSS value at each sender node for 40 different links (two in each of the 20 topologies). In all cases, the Greedy and Optimal schemes generate less interference than the Omni scheme and the Greedy scheme performs close to the Optimal scheme. The carrier sense threshold in 802.11a mode at 6 Mbps is approximately -85 dBm [13]. When the RSS at the sender from the interferer is below this threshold, there is a high probability for the sender to completely eliminate the interferer and achieve higher spatial reuse. In the Greedy scheme 23 out of 40 links (58%) yield RSS less than -85 dBm, in the Optimal scheme 26 links



(a) Four-sector Vivaldi antenna.



(b) Radiation pattern in azimuth plane for a single sector.



(c) Radiation pattern in azimuth plane for omni-mode.

Fig. 1. Sectorized antenna used in our study. The concentric circles in the radiation patterns are 2 dB per division.

(65%) and in the Omni scheme no link. This clearly shows that our topology control protocol improves spatial reuse by reducing sender side interference. While higher spatial reuse leads to increased sending rates at the sender nodes, it may also lead to increased directional hidden terminal problems at the receivers. We therefore proceed to evaluate the receiver side interference reduction due to our topology control protocol.




Fig. 3. Sender interference reduction.





Fig. 4. Receiver interference reduction.


2) *Receiver side interference reduction:* The effect of the directional hidden terminal problem can be reduced by improving the signal to interference and noise ratio (SINR) at the receiver. When two packets arrive at the same time instant at a receiver node, the packet with a higher signal strength is captured if the difference in signal strength (SINR) is greater than the capture threshold. We evaluate the improvement in SINR using our topology control protocol. Note that the objective function aims to minimize the RSS from the interferers while maintaining the RSS from the sender at a high level, and thereby aims to improve the SINR at the receiver.

Figure 4 depicts the SINR in each of the 40 links using the three sector assignment schemes. The SINRs of the Greedy and Optimal schemes are higher compared to the SINR of the Omni scheme in most cases. The capture threshold in 802.11a mode at 6 Mbps is approximately 6 dB [13]. When the SINR of a link is more than this threshold, the signal is captured with high probability. Both Greedy and Optimal schemes result in 31 out of 40 (78%) with SINR greater than 6 dB, while the corresponding fraction using the Omni scheme is 19 out of 40 links (48%). Our topology control scheme thus reduces the receiver side interference compared to Omni-

mode. Therefore, it reduces the probability of occurrence of the directional hidden terminal problem and can thus leverage on the higher spatial reuse achieved at the sender node. Next, we study the impact of reducing both sender and receiver side interference on the aggregate network throughput.



(a) Objective value at 1 dBm Tx power. (b) Objective value at 4 dBm Tx power.



(c) Aggregate throughput at 1 dBm Tx power. (d) Aggregate throughput at 4 dBm Tx power.

Fig. 5. Objective value (aggregate interference) and aggregate throughput at two network densities, where nodes use 1 dBm and 4 dBm Tx power level.

3) *Objective value and aggregate throughput at different network densities:* Figures 5(a) and 5(b) show the objective value as defined in Equation (3) in our optimization framework for two network densities. In the first case, all nodes use the default 1 dBm Tx power. In the second case, all nodes use 4 dBm Tx power and the network is very dense—all 15 node pairs in our six-node network have 100% delivery ratio in both directions under Omni-mode. In both cases, on the average there is approximately 8 dB reduction in the objective value (aggregate network interference) when using the Greedy or the Optimal scheme compared to the Omni scheme.

Figure 5(c) and 5(d), depicts the aggregate network throughput using broadcast packets in the two network densities. We use broadcast packets rather than unicast packets, to




Fig. 6. Aggregate throughput using unicast traffic.

understand the maximum increase in spatial reuse in the absence of MAC level retransmissions and exponential backoff. Both Greedy and Optimal schemes result in higher aggregate throughput compared to the Omni scheme. At 1 dBm Tx power, the Greedy and Optimal schemes both result in two-fold aggregate throughput improvement for half of the topologies. At 4 dBm Tx power only a few topologies achieve similar improvement. In such highly dense networks, finer sectorization (sectors with smaller width) and joint Tx power control may lead to higher spatial reuse.


4) *Performance using unicast traffic:* In this section, we present throughput results for unicast traffic sent at 6 Mbps data rate. Figure 6 depicts the aggregate throughput for the three sector assignment schemes. The Greedy and Optimal schemes outperform the Omni scheme, though the relative improvement is less compared to when using broadcast traffic (cf. Figure 5(c)). There are two main reasons for this reduced improvement. First, unicast traffic has ACK packets sent by the receiver nodes that can reduce spatial reuse. Second, ACK packets can collide with data packets on other links causing increased directional hidden terminal problems. We observe an average throughput improvement of more than 1.5 Mbps (33% improvement over 5 Mbps) when using the Greedy scheme compared to the Omni scheme.

5) *Impact of varying C_{th} :* In this section, we study the performance impact of different values of C_{th} . Recall that this threshold corresponds to the maximum allowed reduction in RSS compared to Omni-mode in order to consider a specific sector assignment (i.e., antenna patterns leading to an RSS reduction greater than this threshold are not considered).

Figure 7 shows the aggregate throughput of the three sector assignment schemes for two different threshold values of C_{th} . As C_{th} increases from 0 dB to 5 dB, the fraction of links that achieve throughput improvements increases from 25% to 50%. Among those links, the average throughput improvement is around 60%. The performance when using 3 dB threshold (default value in our experiments) is intermediate between these two cases (See Figure 5(c)).

This performance improvement is explained as follows. A slight increase of the C_{th} leads to considering more sector combinations with fewer sectors (including single sectors). This, in turn, translates into less generated interference and thus a higher sending rate at the senders and less directional


hidden terminals at the receivers. An exception is topology 10 in Figure 7(b). In this case, there was higher spatial reuse at the senders but high collision at the receivers due to hidden terminals.



(a) $C_{th} = 0dB$, 1 dBm Tx power. (b) $C_{th} = 5dB$, 1 dBm Tx power.

Fig. 7. Comparison of aggregate throughput with varying threshold C_{th} .

6) *Performance using different prediction schemes:* In this section, we evaluate the performance impact of letting our measurement protocol reduce measurement complexity by applying a measurement method where multi-sector RSS is predicted from single sector RSS measurements ('Method 2' in Section VI). We compare this method with that obtained from measuring in all sector combinations ('Method 1' in Section VI) and against the Omni scheme. Figure 8 shows the aggregate throughput when using these three different schemes. We observe that for both the Greedy and Optimal schemes, Method 2 results in performance degradation compared to Method 1. However, Method 2 still outperforms the Omni scheme for several topologies in both cases.



(a) Greedy scheme.

(b) Optimal scheme.

Fig. 8. Comparison of two prediction methods for Greedy and Optimal schemes with the Omni scheme.

In the next two sections, we present results from a WLAN and a multi-hop wireless mesh network configuration.

7) *WLAN scenario:* We consider a WLAN scenario where the access points (APs) are equipped with sectorized antennas to reduce interference, while the clients have only omni-directional antennas. We use eight different topologies, each consisting of two APs with two clients each. We use *iperf* to send downlink backlogged traffic, from APs toward the clients.

We expect the relative improvement of our topology control protocol to be slightly less than in previous experiments. This is because APs need to use wider sector selection to communicate with multiple clients. In addition, since the clients use omni-directional antennas they cannot contribute to spatial reuse improvement and, furthermore, they are more exposed to directional hidden terminals. Despite these constraints, we observe in Figure 9 that the Greedy scheme results in 15-25%

aggregate throughput improvement over the Omni scheme except in topologies 1 and 3. In these cases, both schemes resulted in the same sector assignments.




Fig. 9. Aggregate throughput improvement in WLAN scenario.

8) *Multi-hop scenario*: We now consider a small multi-hop mesh network configuration. We configure nodes as mesh routers and select eight different topologies where backlogged unicast UDP traffic is transmitted simultaneously over a two-hop path and a single hop path. Mesh configurations pose more restrictions to topology control since mesh routers that forward traffic must have a wider sector selection enabled to communicate with both their predecessor and successor nodes in the multi-hop path. Figure 10 depicts the aggregate




Fig. 10. Aggregate throughput improvement in mesh scenario.

network throughput for the Greedy scheme and Omni scheme in these eight topologies. The Greedy scheme yields 40-60% improvement in four topologies, about 15% improvement in two topologies, and performs similar to the Omni scheme in the remaining topologies.

VIII. CONCLUSIONS

We introduced, implemented and evaluated a distributed topology control optimization framework and attendant protocols for dense 802.11 wireless networks using sectorized antennas. We first formulated the topology control problem as an optimization problem that minimizes aggregate network interference and introduced techniques to find lower bounds to the optimal solution. The optimization framework is based on a physical interference model with measured RSS values of different sector activations and leverages practical insights from previous experimental studies using sectorized antennas.

We then introduced a greedy distributed topology control protocol and an associated RSS measurement protocol that can be implemented in commodity 802.11 hardware and operate transparently with the 802.11 MAC and network layer link selection protocols. Our experimental evaluation on a dense wireless network testbed showed that our approach achieves both near-optimal performance and significant throughput increase over omni-directional antennas in various network and traffic scenarios. The results also show that network density affects the performance gains due to sectorization. Finally, we quantified the tradeoff between RSS measurement complexity and throughput performance when predicted RSS values are used for topology control.

We conclude that carefully engineered topology control protocols based on sectorized antennas without complex modifications to the 802.11 MAC protocol can provide substantial throughput gains in real world 802.11 wireless networks. In our future work, we plan to use finely tunable sectorized antennas along with dynamic transmit power and data rate adaptation to achieve maximum spatial reuse in dense wireless deployments.

REFERENCES

- [1] ILOG CPLEX: Mathematical Programming Optimizers. <http://www.ilog.com/products/cplex/>.
- [2] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan. On the Effectiveness of Switched Beam Antennas in Indoor Environments. In *Proc. PAM*, 2008.
- [3] R. Choudhury and N.H. Vaidya. Deafness: A Mac Problem in Ad Hoc Networks when using Directional Antennas. In *Proc. IEEE ICNP*, October 2004.
- [4] R. Choudhury, X. Yang, R. Ramanathan, and N.H. Vaidya. Using Directional Antennas for Medium Access Control in Ad Hoc Networks. In *Proc. ACM MobiCom*, September 2002.
- [5] Z. Huang, C. Shen, C. Srisathapornphat, and C. Jaikao. Topology Control for Ad hoc Networks with Directional Antennas. In *Proc. IEEE ICCCN*, 2002.
- [6] Zhuochuan Huang and Chien-Chung Shen. Multibeam Antenna-based Topology Control with Directional Power Intensity for Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 5(5), 2006.
- [7] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proc. ACM MobiCom*, 2003.
- [8] A. Kashyap, S. Ganguly, and S. R. Das. A Measurement-Based Approach to Modeling Link Capacity in 802.11-Based Wireless Networks. In *Proc. ACM MobiCom*, 2007.
- [9] Y. Ko, V. Shankarkumar, and N.H. Vaidya. Medium Access Control Protocols Using Directional Antennas in Ad Hoc Networks. In *Proc. IEEE INFOCOM*, March 2000.
- [10] U. Kumar, H. Gupta, and S.R. Das. A Topology Control Approach to using Directional Antennas in Wireless Mesh Networks. In *Proc. ICC*, 2006.
- [11] P. Minard, J.L. Robert, and A. Louzir. Multi-beam Antenna System for Mesh Network with Omnidirectional Mode. In *In Proc. IEEE Antennas and Propagation Society International Symposium (AP-S)*, 2008.
- [12] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan. A General Model of Wireless Interference. In *Proc. ACM MobiCom*, 2007.
- [13] A. P. Subramanian, H. Lundgren, and T. Salonidis. Experimental Characterization of Sectorized Antennas in Dense 802.11 Wireless Mesh Networks. In *Proc. ACM MobiHoc*, 2009.
- [14] K. Sundaresan, W. Wang, and S. Eidenbenz. Algorithmic Aspects of Communication in Ad-hoc Networks with Smart Antennas. In *Proc. ACM MobiHoc*, 2006.