Passive Learning of the Interference Graph of a Wireless Network

Jing Yang, Stark C. Draper, and Robert Nowak

Department of Electrical and Computer Engineering
University of Wisconsin, Madison, WI 53706
Email: {*yangjing, sdraper, nowak*}@*ece.wisc.edu*

***Abstract*—A key challenge in wireless networking is the management of interference between transmissions. Identifying which transmitters interfere with each other is crucial. Complicating this task is the fact that the topology of wireless networks can change from time to time, and so the identification process may need to be carried out on a regular basis. Injecting active probing traffic to assess interference can lead to unacceptable overhead, and so this paper focuses on interference estimation based on passive traffic monitoring in networks that use the CSMA/CA (Carrier Sense Multiple Access/Collision Avoidance) protocol. A graph is used to represent the interference in the network, where the nodes represent transmitters and edges represent interference between pairs of transmitters. We investigate the problem of learning the graph structure based on passive observations of network traffic transmission patterns and information about successes or failures in transmissions. Previous work has focused on algorithms and validations in small testbed networks. This paper focuses on the scaling behavior of such methods which is unaddressed in prior work. In particular we establish bounds on the minimum observation period required to identify the interference graph reliably. The main results are expressed in terms of the total number of nodes n and the maximum number of interfering transmitters per node (i.e., maximum node degree) d. The effects of hidden terminal interference (i.e., interference not detectable via carrier sensing) on the observation time requirement are also quantified. We show that it is necessary and sufficient that the observation period grows like $d^2 \log n$, and we propose a practical algorithm that reliably identifies the graph from this length of observation. We conclude that the observation requirements scale quite mildly with network size, and that the networks with sparse interference patterns can be more rapidly identified than those with dense interference patterns.**

I. Introduction

Due to the broadcast nature of wireless communications, simultaneous transmissions in the same frequency band and time slot may interfere with each other, thereby limiting system throughput. Interference estimation is thus an essential part of wireless network operation. Knowledge of interference among nodes is an important input in many wireless network configuration tasks, such as channel assignments, transmit power control, and scheduling. A number of recent efforts have made significant progress towards the goal of identifying in real time the interference in the network. Some of the recent approaches (e.g., Interference maps [1] and Micro-probing [2]) inject active traffic into the network to infer occurrences of interference. While such approaches can be quite accurate in determining interference, the overhead of making active measurements across a large network could limit their practicality. In addition, network topology can change over time, meaning that using active probing methods periodically to identify interference can place a significant burden on the network.

Inspired by two passive WLAN (wireless local area network) monitoring approaches Jigsaw [3], [4] and WIT [5], in [6], [7], a different approach to online interference estimation is explored, termed Passive Interference Estimation (PIE) algorithm. Rather than introducing active probe traffic, interference estimates are completely based on passively collected data. PIE estimates the actual interference in the network in real-time in an on-line fashion. All interference estimates are obtained by simply observing ongoing traffic and feedback information at the different transmitters. Experimental studies in small testbeds show that PIE is quite promising, but very little is understood about how the method scales-up to large complex networks, which is the focus of this paper.

The main contribution of this paper is to quantify the required observation time as a function of network size and topological connectivity. This provides insight into the time-scales over which network interference patterns can be identified, and potentially tracked, in a dynamically changing wireless network.

We formulate passive interference estimation as a statistical learning problem. Given an arbitrary WLAN that consists of n access points (APs), our goal is to recover a conflict graph among these APs with as few measurements as possible. Denoting the network size (number of APs/nodes) by n, and the maximum number of interfering APs per AP by d, we show that to identify the network one must collect a number of measurements proportional to $d^2 \log n$. This is quite mild dependence on network size and indicates that interference graph inference is scalable to large networks. The dependence on d tells us that sparser patterns of interference are easier to identify than denser patterns. We also quantify the effects of hidden terminal interference (i.e., interference not detectable via carrier sensing), which introduce an additional factor to the observation time requirement. We give lower bounds showing that this time-complexity cannot be improved by any possible scheme, and we give simple and easy-to-implement algorithms that achieve the bound (up to constants).

978-1-4673-2579-0/12/$31.00

II. Problem Formulation

A. Interference Graph

We use a graph $G = (\mathcal{V}, \mathcal{E})$ to represent the interference among APs in the network, where the node set $\mathcal{V}$ represents the APs, and edge set $\mathcal{E}$ represents the pairwise interference level among APs. An example of such of graph is depicted in Figure 1.

Because of the CSMA/CA protocol, we divide $\mathcal{E}$ into two subsets: *direct* interference $\mathcal{E}_D$ and *hidden* interference $\mathcal{E}_H$. Direct interference occurs when two APs are within each other's carrier sensing range. Since they can hear each other, under the protocol they avoid simultaneous transmissions. Under the assumption that the carrier sensing range is the same for every AP, the directly interfering edges in $\mathcal{E}_D$ are *undirected*. However, CSMA/CA cannot resolve all of the collisions in the network. If two APs cannot hear each other, and the transmission from one AP is corrupted at certain clients by the other AP, this type of interference is referred as hidden interference. The edges in $\mathcal{E}_H$ represent this type of hidden interference. Such interferences may be asymmetric and so the edges in $\mathcal{E}_H$ are *directed*. Additionally, it is possible to associate a probability or weight with each interference edge, reflecting the level (severity or persistence) of the interference. Our main interest is to identify the presence of interference, since once this is accomplished it is straightforward to estimate the level of the interference. The sparsity of the interference graph is parameterized by d, the maximum node degree. This parameter is a bound on the number of nodes that interfere with any node.

B. CSMA/CA protocol and ACK/NACK mechanism

The system is assumed to use the CSMA/CA protocol at the medium control layer [8]. When the transmitter has a packet to send, it listens to the desired channel. If the channel is idle, it sends the packet. If the channel is busy, i.e., there exists an active transmitter within the listener's carrier sensing range, the transmitter holds its transmission until the current transmission ends. The following backoff mechanism is then invoked. The transmitter randomly chooses an integer number, τ, uniformly in the range $[0, W-1]$, and waits for τ unit slots. W is a positive integer and represent the current backoff window size. If the channel is idle at the end of the node's backoff period, the node transmits its packet; otherwise it repeats the process until it gets a free channel. Statistically, the random backoff mechanism allows every node equal access to the medium.

In some communication situations, even though two APs cannot hear each other, when they transmit simultaneously, a collision happens. Such failures occur because one or both clients communicating with them lie within the other AP's transmission range. This is the so-called "hidden terminal" problem. We assume the ACK/NACK mechanism is adopted in the system. Specifically, we assume that whenever an AP successfully delivers a packet to its destination, an ACK will be fed back to acknowledge the successful transmission. If the AP does not receive the ACK after a period of time, it assumes that the packet is lost. Thus, the ACK/NACK mechanism enables the APs to detect collisions whenever they occur.

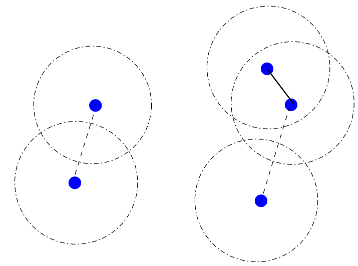

Fig. 1. System model with n APs. The edges connecting transmitters represent the interference between them. Solid line represents direct interference, and dotted lines represent hidden interference.

C. Inference Problem and Algorithm

Our objective is to determine all of the edges in the interference graph G based on passive observations. We assume that there exist a central controller that collects the transmission status of all APs, and the transmission success information (ACKs) fed back to them during a period of time. The collected information during this *observation period* is the *dataset* for us to infer the interference graph.

Based on the CSMA/CA protocol, if two APs are within each other's carrier sensing range, they don't transmit simultaneously. Therefore, for any pair of APs that are transmitting simultaneously at any time instance during the observation period, the transmission status indicates that they cannot hear each other, i.e., there is no *direct* interference between them. In the interference graph, all of the corresponding edges in $\mathcal{E}_D$ can thus be removed. When the observation period is sufficiently long, the correct *direct* interference edge set $\mathcal{E}_D$ can be successfully recovered.

The collected ACK information can be used to identify the *hidden* interference edge set $\mathcal{E}_H$. Once a collision is detected at an AP, it implies that one of the APs transmitting at the same time should be the interferer causing the collision. The subset of APs transmitting at that time instance becomes a candidate set of hidden interferes for that AP. When multiple collisions are detected for an AP, multiple candidate subsets are formed, one for each collision. Then, the true set of hidden interferers for that AP intersects with all of those candidate subsets. When the observation period is sufficiently long, the *minimum hitting set* which intersects with all of the candidate subsets is the true hidden interferer set. Therefore, edge set $\mathcal{E}_H$ can be recovered.

Our algorithm adopts the idea of PIE proposed in [6], [7]. For the direct interference, both approaches rely on simultaneous transmissions to infer the interference. For the hidden interference estimation, PIE uses a correlation based approach. Our approach is different, and attains more accurate results theoretically.

D. Statistical Model

For the simplicity of analysis, we partition the time axis into sessions. We assume that the APs are synchronized, so that they share the same sessions. At the beginning of each session, APs with traffic contend for the channel. In the following sections, we use a binary variable $Q_i(t)$ to indicate whether AP i has any traffic to send at a given session t. AP i contends for the channel in session t if and only if $Q_i(t) = 1$.

We let a graph $G_D = (\mathcal{V}, \mathcal{E}_D)$ represent the carrier sensing relationships among the APs. Specifically, if AP i and AP j are within each other's carrier sensing range, there is an edge between i, j, denoted as (i, j). We define $\mathcal{N}_i$ to be the set of neighbors of AP i in $G_D = (\mathcal{V}, \mathcal{E}_D)$, and let $d_i = |\mathcal{N}_i|$.

Define $T_i(t)$ as the backoff time of AP i at the beginning of session t. We denote $X_i(t)$ as the activation state of node i , where $X_i(t) = 1$ means that node i wins the channel in session t, and transmits in that session. $X_i(t) = 0$ means that node i is not active. Thus, $X_i(t) = 1$ if and only if $Q_i(t) = 1$ and $X_j(t) = 0$ for all $j \in \mathcal{N}_i$ with $T_j(t) < T_i(t)$.

Define $Y_i(t) \in \{0, 1, \varnothing\}$ to be the feedback information information received at AP i at the end of session t. $Y_i(t) = 1$ means that an ACK is received at AP i, indicating that the transmission in session t is successful; $Y_i(t) = 0$ means that transmission failed, caused by some simultaneous transmissions; $Y_i(t) = \varnothing$ means that node i does not transmit in that session, i.e., $X_i(t) = 0$.

We let a graph $G_H = (\mathcal{V}, \mathcal{E}_H)$ represent the hidden interference among the APs that cannot hear each other. Since such interference usually depends on the locations of clients when the transmission takes place, and thus is not symmetric. Therefore, the edges are *directed*, e.g., $i \to j$. We define p_{ij} as the probability that, when $X_i(t) = X_j(t) = 1$, AP i interferes with AP j, causing transmission failure of AP j. Define $\mathcal{S}_j$ to be the hidden interferer set for AP j, the set of APs with $p_{ij} > 0$. We let $s_j = |\mathcal{S}_j|$. We point out that in general

$$p_{ij} \neq \mathbb{P}(Y_j(t) = 0 | X_i(t) = 1, X_j(t) = 1),$$

i.e., p_{ij} is not equal to the probability that a collision occur at AP j when both APs i and j are transmitting. This is because the collision at AP j may be caused by an active transmitter other than i, and transmitter i may just happen to be transmitting at the same time.

Assumptions 1 *We make the following assumptions:*

1) *$Q_i(t)$s are i.i.d. Bernoulli random variables with common parameter p, $0 < p < 1$, i.e., $\mathbb{P}(Q_i(t) = 1) = p$ for all i, t.*
2) *$T_i(t)$s are continuous i.i.d. random variables uniformly distributed over $[0, W]$.*
3) *p_{ij}s are lower bounded by p_{min}.*
4) *d_is are upper bounded by the integer $d - 1$.*
5) *s_js are upper bounded by the integer s.*

The first assumption ignores the time dependency and coupling effect of the traffic queue status of APs. It is a good approximation of traffic for a system operating in light traffic regime. For the second assumption, we assume $T_i(t)$s are continuous instead of discrete quantities in CSMA/CA protocol to ensure that two adjacent nodes in $G_D = (\mathcal{V}, \mathcal{E}_D)$ don't have the same backoff time. This is a valid assumption when W is a large integer. The first two assumptions guarantee that the joint distribution of $X_i(t)$s and $Y_i(t)$s is independent and identical across t. Therefore, in our analysis here afterwards, we ignore the time index and focus on the distribution of X_is and Y_is in one observation.

III. Main Results

Under assumptions of the statistical model, we break the overall problem into four subproblems. The first problem is to upper bound the number of observations required to infer the directed interference edge set $\mathcal{E}_D$. Only the transmission patterns of APs are necessary for this inference. The second problem is to provide a lower bound on the number of observations required to correctly infer $\mathcal{E}_D$. In contrast to $\mathcal{E}_D$, inferring the hidden interference edge set $\mathcal{E}_H$ requires feedback information about transmission success as well as the transmission patterns. The third and fourth subproblems are to provide upper bound and lower bound on the number of observations required to infer $\mathcal{E}_H$.

A. An Upper Bound for Inference of $G_D = (\mathcal{V}, \mathcal{E}_D)$

In this section, we aim to estimate $G_D = (\mathcal{V}, \mathcal{E}_D)$ based on a sequence of transmission patterns, e.g., $\mathbf{X}(1), \mathbf{X}(2), \ldots, \mathbf{X}(k)$. Due to the use of the CSMA/CA protocol and the continuous backoff time of Assumption 2), any pair of APs active in the same slot must not be able to hear each other. Thus between any two of APs active in the same slot, there is no edge in $G_D = (\mathcal{V}, \mathcal{E}_D)$. In other words, given an observation $\mathbf{X}$, for any i, j with $X_i = X_j = 1$, we know that $(i, j) \notin \mathcal{E}_D$.

Based on this observation, we propose the following algorithm. Start with a fully connected graph connecting the n APs with $n(n-1)/2$ edges. For each transmission pattern $\mathbf{X}$ observed, remove all edges (i, j) s.t. $X_i = X_j = 1$. Our main question is how many observations are required in order to eliminate all edges not in $\mathcal{E}_D$, thereby recovering the underlying interference graph G_D. Furthermore, how close to optimal is this algorithm. Let k be the number of transmission patterns observed. The following theorem provides an upper bound on the required k to recover G_D with high probability.

Theorem 1 *Let $\delta > 0$, and let*

$$k \geq \frac{1}{-\log\left(1 - p^2/d^2\right)} \left(\log \binom{n}{2} + \log \frac{1}{\delta} \right) \tag{1}$$

then, with probability at least $1 - \delta$, the estimated interference graph $\hat{G}_D = (\mathcal{V}, \hat{\mathcal{E}}_D)$ is equal to G_D after k observations.

The proof of this theorem is provided in Appendix A. The idea is first to lower bound the probability that two nonadjacent APs i, j never transmit simultaneously in k observations. Then, by taking a union bound, an upper bound on the required k is obtained.

Remark: We note that in the theorem above, the upper bound $k = O(d^2 \log n)$ when $p^2/d^2 \to 0$. Intuitively, this order is the best we hope for in terms of p and d through passive observations. This is because if two non-interfering

APs never transmit simultaneously, their behavior is the same as if they were within each other's carrier sensing range. Thus, we cannot determine whether there is an edge between them or not. Since each transmitter competes with its neighbors, the probability that it gets the channel is of the order of $1/d$. So, roughly speaking, it takes about d^2 snapshots to observe two non-interfering APs active at the same time. Since there are about n^2 such pairs in the network, a union bound gives us the factor $\log n$.

B. A Minimax Lower Bound for Inference of $G_D = (\mathcal{V}, \mathcal{E}_D)$

We provide a minimax lower bound on the number of observations needed to recover the underlying direct interference graph $G_D = (\mathcal{V}, \mathcal{E}_D)$. Denote the estimated direct interference graph as $\hat{G}_D = (\mathcal{V}, \hat{\mathcal{E}}_D)$.

Theorem 2 *If $d - 1 \leq \frac{3n - \sqrt{n^2+16n}}{4}, n \geq 3$, and the number of observations satisfies $k \leq \frac{\alpha(d-1)^2}{\left(2+\frac{1}{1-p}\right)} \log n$, then,*

$$\min_{\hat{G}_D} \max_{G_D} \mathbb{P}(\hat{G}_D \neq G_D; G_D) \geq \frac{\sqrt{n}}{1+\sqrt{n}} \left(1 - 2\alpha - \sqrt{\frac{2\alpha}{\log n}}\right)$$

for $0 < \alpha < 1/8$.

The proof of this theorem, and the proofs of the following theorems, can be found in [9]. The outline of our proof is as follows: We construct a class of M graphs which is a subset of maximum-degree d graphs. We construct the graphs in such a way that they are very similar to each other. Thus it is hard to distinguish between them. Next, we reduce the original estimation problem to an M-ary hypothesis test. We lower bound the probability of error for this hypothesis test based on the Kullback-Leibler divergence between transmission pattern distributions under a pair of graphs. We prove that for this special collection of graphs, the number of observations required to detect the correct underlying conflict graph, is at least of the order $\Omega(d^2 \log n)$. Since the M-ary hypothesis test is easier than the original problem, this is also a lower bound for the original estimation problem.

Remark: We note that the lower bound is of the order $\Omega(d^2 \log n)$. It is in the same order of the upper bound in Theorem 1. Therefore, the estimation method based on pairwise comparison is optimal.

C. An Upper Bound for Inference of $G_H = (\mathcal{V}, \mathcal{E}_H)$

In the previous subsection, we fully characterize the direct interference in the interference graph. In this section, we aim to infer the hidden interference graph $G_H = (\mathcal{V}, \mathcal{E}_H)$ through observations of $\mathbf{X}$ and $\mathbf{Y}$.

We assume that there exists s_j hidden interferers for transmitter, $0 < s_j \leq s$. Because of this, when a collision happens to AP j, it may have been caused by any of the active APs in $\mathcal{S}_j$. In other words, there may not be a unique AP consistently showing up when $Y_j = 0$. We also note that an AP showing up frequently when $Y_j = 1$ may not be a hidden interferer. The cause can be that the transmission status of an AP is highly positively correlated (because of CSMA/CA) with one or more hidden interferers. Thus, although not the case, it transmits with high probability when the collision caused by the true hidden interferer occurs. This implies that correlation based approach may not accurately distinguish true interferers from such non-interferers. In the following, we propose a minimum hitting-set-based approach to infer the hidden interference graph $G_H = (\mathcal{V}, \mathcal{E}_H)$.

Define $\mathcal{K}_j$ as the time index set $\{t\}$ of the observations with $Y_j(t) = 0$. Then, for every index $t \in \mathcal{K}_j$, we define $\mathcal{S}_j^t$ as the set of transmitters with $X_i(t) = 1$ in $\mathbf{X}(t)$.

Definition 1 *Given a collection of subsets of some alphabet, a set which intersects all subsets in the collection in at least one element is called a hitting set. A "minimum" hitting set is a hitting set of the smallest size.*

Lemma 1 *As $k \to \infty$, the hidden interferer set $\mathcal{S}_j$ for transmitter j is the unique minimum hitting set such that $\mathcal{S}_j \cap \mathcal{S}_j^t \neq \varnothing$ for every $t \in \mathcal{K}_j$.*

Therefore, given k observations, in order to recover $G_H = (\mathcal{V}, \mathcal{E}_H)$, we look for the minimum hitting set $\mathcal{S}_j$ for every j. The following theorem provides an upper bound on the minimum number of observations required to correctly identify the minimum hitting set $\mathcal{S}_j$ for every AP $j \in \mathcal{V}$. Once the estimated minimum hitting set $\hat{\mathcal{S}}_j$ is obtained, the estimated hidden interference graph $\hat{G}_H$ can be constructed by adding a directed edge from any AP in $\hat{\mathcal{S}}_j$ to AP j.

Theorem 3 *Let $\delta > 0$, and let*

$$k \geq \frac{1}{-\log\left(1 - \frac{p^2(1-p)^s p_{min}}{d^2}\right)} \left(\log(ns) + \log\frac{1}{\delta}\right)$$

Then, with probability at least $1 - \delta$, $\hat{G}_H$ equals G_H.

For every AP j, we first upper bound the probability that the minimum hitting set obtained after k observations is not equal to the true minimum hitting set. This is equal to the probability that there exists at least one AP i in $\mathcal{S}_j$ that is not included in $\hat{\mathcal{S}}_j$. By taking the union bound across all possible i and j we obtain an upper bound on k.

In general, the minimum hitting set problem is NP-hard. However, under the assumption that $s \ll n$, the problem can be solved in polynomial time.

One straightforward approach is to test every subset of nonadjacent APs in G_D to determine whether it is a hitting set of the sets $\{\mathcal{S}_j^t\}$s. We start from the subset with one element, and gradually increase the size of the testing subset by one, until a hitting set is achieved. Since this is a smallest size we can achieve, it is a minimum hitting set for the given $\{\mathcal{S}_j^t\}$s.

For the worst case, the number of subsets we need to test is upper bounded by $O(n^s)$ [10].

D. A Minimax Lower Bound for Inference of $G_H = (\mathcal{V}, \mathcal{E}_H)$

In this subsection, we develop a lower bound on the number of observations required to recover the underlying hidden interference graph $G_H = (\mathcal{V}, \mathcal{E}_H)$ with high probability based on passive observations.

Theorem 4 *If $d \leq c_1 n, \quad s-1 \leq c_2 n, \quad 2c_1 + c_2 < 1$, and the number of observations satisfies*

$$k \leq -\frac{\log(2c_1(\frac{1}{2c_1+c_2}-1)n)}{\left(\frac{1-(1-p)^d}{d}\right)^2 (1-p)^{s-1}\log(1-p_{min})}$$

then

$$\min_{\hat{G}_H}\max_{G}\mathbb{P}(\hat{G}_H \neq G_H; G)$$
$$\geq \frac{\sqrt{2c_1(\frac{1}{2c_1+c_2}-1)n}}{1+\sqrt{2c_1(\frac{1}{2c_1+c_2}-1)n}}\left(1-2\alpha-\sqrt{\frac{2\alpha}{\log(2c_1(\frac{1}{2c_1+c_2}-1)n)}}\right)$$

for $0 < \alpha < 1/8$.

We follow the same approach in the proof of Theorem 2. We reduce the original problem to an M-ary hypothesis test and show that asymptotically the lower bound is of the same order as the upper bound.

Remarks: Since the distribution of $\mathbf{Y}$ depends on the underlying direct interference graph as well as on the hidden interference graph, the lower bound is over all possible interference graphs G. As d increases, the lower bound in Theorem 4 is of order $\Omega\left(-\frac{d^2}{\log(1-p_{min})(1-p)^{s-1}}\log n\right)$. Since $\log(1-p_{min})$ can be approximated as p_{min} when p_{min} is small, the lower bound is of the same order as the upper bound provided in Theorem 3. Therefore, the bounds are tight and our algorithm is optimal.

Appendix A
Proof of Theorem 1

Consider any two nonadjacent nodes i, j in $G_D = (\mathcal{V}, \mathcal{E}_D)$. Let $\mathcal{N}_{ij} = \mathcal{N}_i \cup \mathcal{N}_j$, and $\mathcal{N}_{i\setminus j} = N_i \cap N_j^c$.

Under the CSMA/CA protocol, transmitter i only contends for the channel when $Q_i = 1$. For ease of analysis, in this proof, we assume that every transmitter competes for the channel at the beginning of each session no matter whether or not $Q_i = 1$. Then, if transmitter i gets the channel, it starts transmit if $Q_i = 1$; otherwise, even it gets the opportunity to transmit, it keeps silent. The latter event leaves the opportunity open for one of its neighbors with a longer backoff time to transmit. Since this does not change the outcome of transmission pattern given that the backoff times of those transmitters with nonempty queues are fixed, the statistics of $\mathbf{X} := (X_1, X_2, \ldots, X_n)$ stay the same under this interpretation.

Define $T_{\mathcal{N}_i}$ as the minimum backoff time of the nodes in the set $\mathcal{N}_i$. Then,

$$\mathbb{P}\left(T_i < T_{\mathcal{N}_i}, T_j < T_{\mathcal{N}_{j\setminus i}}\right) = \frac{\binom{|\mathcal{N}_{ij}|+2}{|\mathcal{N}_i|+1}\cdot|\mathcal{N}_i|!\cdot|\mathcal{N}_{j\setminus i}|!}{(|\mathcal{N}_{ij}|+2)!} \geq \frac{1}{d^2}$$

The probability is obtained in this way: Considering nodes i, j and their neighbors, there are $|\mathcal{N}_{ij}|+2$ nodes in total, and there are $(|\mathcal{N}_{ij}|+2)!$ orderings of their backoff times. Among these ranks, $\binom{|\mathcal{N}_{ij}|+2}{|\mathcal{N}_i|+1}\cdot|\mathcal{N}_i|!\cdot|\mathcal{N}_{j\setminus i}|!$ ranks correspond to $T_i < T_{\mathcal{N}_i}, T_j < T_{\mathcal{N}_{j\setminus i}}$. Such rank can be obtained in this way: suppose these $|\mathcal{N}_{ij}|+2$ nodes are ordered according to their backoff times. Then, node i and its neighbors take $|\mathcal{N}_i|+1$ positions of them. There are $\binom{|\mathcal{N}_{ij}|+2}{|\mathcal{N}_i|+1}$ such combinations in total. Since $T_i < T_{\mathcal{N}_i}$, node i takes the first position out of the chosen $|\mathcal{N}_i|+1$ positions, and the rest $|\mathcal{N}_i|$ positions are for its neighbors. This results in $|\mathcal{N}_i|!$ ranks for each fixed position list. Node j and nodes in $\mathcal{N}_{j\setminus i}$ take the rest of the positions, and node j takes the first. This gives us the factor $|\mathcal{N}_{j\setminus i}|!$.

When $T_i < T_{\mathcal{N}_i}$, and $Q_i = 1$, based on the protocol, node i gets the channel thus $X_i = 1$. At the same time, transmission from all nodes in $\mathcal{N}_i$ are suppressed. Therefore, for node j, if $T_j < T_{\mathcal{N}_{j\setminus i}}$ and $Q_j = 1$, node j also gets a channel. Thus, we have $X_i = X_j = 1$. Since other scenarios result in $X_i = X_j = 1$, we have the following

$$\mathbb{P}(X_i = 1, X_j = 1)$$
$$\geq \mathbb{P}\left(T_i < T_{\mathcal{N}_i}, T_j < T_{\mathcal{N}_{j\setminus i}}, Q_i > 0, Q_j > 0\right) \geq \frac{p^2}{d^2},$$
$$\mathbb{P}(\text{edge } (i,j) \text{ is not removed with one observation } \mathbf{X})$$
$$= 1 - \mathbb{P}(X_i = 1, X_j = 1) \leq 1 - \frac{p^2}{d^2}.$$

Define $\mathcal{A}_{ij}$ as the event that edge (i, j) is not removed after k observations. Then, the probability that, after k observations, the graph cannot be identified successfully is

$$\mathbb{P}(\hat{G}_D \neq G_D) = \mathbb{P}(\cup_{(i,j)\notin\mathcal{E}_D}\mathcal{A}_{ij}) \leq \binom{n}{2}\left(1-\frac{p^2}{d^2}\right)^k.$$

The inequality follows from the fact that the number of nonadjacent pairs in $G_D = (\mathcal{V}, \mathcal{E}_D)$ is upper bounded by $\binom{n}{2}$. Under the assumption that $d \ll n$, this is a good approximation for the nonadjacent pairs in G_D.

References

[1] D. Niculescu, "Interference map for 802.11 networks," *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007.

[2] N. Ahmed and S. K. Smarta, "A self-managing architecture for thin access points," *ACM CoNEXT*, 2006.

[3] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage, "Jigsaw: solving the puzzle of enterprise 802.11 analysis," *ACM SIGCOMM*, 2006.

[4] Y.-C. Cheng, M. Afanasyev, P. Verkaik, P. Benko, J. Chiang, A. C. Snoeren, S. Savage, and G. M. Voelker, "Automating cross-layer diagnosis of enterprise wireless networks," *ACM SIGCOMM*, 2007.

[5] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Analyzing the mac-level behavior of wireless networks in the wild," *ACM SIGCOMM*, 2006.

[6] V. V. Shrivastava, "Optimizing enterprise wireless networks through centralization," Ph.D. dissertation, University of Wisconsin–Madison, 2010.

[7] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Pagaginnaki, "PIE in the sky: Online passive interference estimation for enterprise wlans," *NSDI*, 2011.

[8] *IEEE 802.11 standard*, http://www.ieee802.org/11/.

[9] J. Yang, S. Draper, and R. Nowak, "Passive interference graph learning in a wireless network," *to be submitted to IEEE Transactions on Information Theory*, 2012.

[10] M. Kalisch and P. Büuhlmann, "Estimating high-dimensional directed acyclic graphs with the pc-algorithm," *Journal of Machine Learning Research*, March 2007.