Improving Throughput Through Spatial Diversity in Wireless Mesh Networks: A PCA-Based Approach

ABSTRACT

In this paper, we consider the problem of mitigating interference and improving network capacity in wireless mesh networks from the angle of *spatial diversity*. In a nutshell, while the achievable throughput on a multihop wireless path is limited by intra-flow interference, the overall capacity of a multihop wireless network can be increased by exploiting spatial diversity that exists among a number of multihop paths. Connections that are routed along these paths can be scheduled to take place simultaneously if their transmissions do not interfere with each other (significantly).

To make a case of exploiting spatial diversity to improve network capacity, we focus on transporting *downstream* traffic at gateway nodes with Internet access. We propose to construct, based on measurements of received signal strengths, a *virtual coordinate system* that is used to determine the sets of paths along which transmissions can take place with the least inter-flow interference. Based on the sets of non-interfering paths, the gateway node then determines the order with which a gateway node schedules frames of different connections to be transmitted. Through extensive simulation (with real-life measurement traces on *Champaign Urbana Wireless Community Network*, we show that the downstream throughput of a gateway node in a wireless mesh network can be improved by 10-55% under a variety of network topologies and traffic distributions. This, coupled with the fact that the proposed approach requires only minor code change in the gateway nodes and does not require any additional hardware, makes it a viable option to improving network capacity in existing wireless mesh networks.

Keywords

Wireless mesh network, Intra-/inter-flow interference, spatial diversity, and topology discovery.

1. INTRODUCTION

Wireless mesh networks have emerged to be a new, cost-effective and performance-adaptive network paradigm for the next-generation wireless Internet. Targeting primarily for solving the well known last mile problem for broadband access [23,25], wireless mesh networks aim to offer high-speed coverage at a significantly lower deployment and maintenance cost. In such networks, most of the nodes are either stationary or less mobile. Only a fraction of nodes have direct access, and will serve as gateways, to the Internet. Several nodes serve as relays forwarding traffic from other nodes (as well as their own traffic) and maintain network-wide Internet connectivity, while the remaining nodes send packets along dynamically selected ad-hoc paths to gateway nodes with Internet access. Wireless mesh networks are preferable to existing cable/DSL based networks or wireless LANs (that provide WiFi access), due to the following potential advantages: (i) it is more cost effective, as service providers do not have to install a wired connection to each subscriber ($20-50$K per square mile to establish access, approximately 1/4 of the cost incurred in high speed cable access); (ii) it is inherently more reliable, as each node has redundant paths to reach the Internet; (iii) the throughput attained by a user can be increased through routing via multiple, bandwidth-abundant paths (in contrast, in WLANs the shared bandwidth decreases as the number of users within a HotSpot increases); and (iv) the wireless network can readily extend their coverage by installing additional ad-hoc hops.

Several cities have planned (and/or partially deployed) wireless mesh networks, such as Bay Area Wireless User Group (BAWUG) [1], Boston Roofnet [3], Champaign-Urbana Community Wireless Network (CUWiN) [2], SFLan [5], Seattle Wireless [4], Southampton Open Wireless Network (SOWN) [6], and Wireless Leiden (in Netherlands) [3]. Although initial success has been reported in these efforts, a number of performance related problems have also been identified. Excessive packet losses/collision [9,12,15], unpredictable channel behaviors [9,12], inability to achieve throughput as afforded by IEEE 802.11 PHY/MAC, inability to find stable and high-throughput paths [9,12], and lack of incentives to forward transit packets [22] are among those most cited to question whether or not the success of wireless mesh networks will sustain.

To deal with the problem of locating stable and high-throughput paths, several research efforts have been made to devise more sophisticated route metrics, e.g., expected transmission count (ETX) [13], round trip time (RTT) [8], and weighted cumulative expected transmission time (WCETT) [14]. Several of the aforementioned, throughput-related problems (e.g., excessive packet losses, unpredictable channel behaviors, and throughput degradation) have also been identified to be attributed, in part, by intra- and inter-flow interference [11,15,26]. Specifically, flows that are routed along different paths within the interference range compete for the channel bandwidth, resulting in inter-flow interference. On the other hand, consecutive packets in a single flow may be spread over the route to their destination and may interfere with one another, resulting in intra-flow interference. With the interference left uncontrolled, the operational range of a wireless mesh network would be limited to within a few hops, representing a *insignificant* stretch from the current wireless LANs or hotspots with respect to coverage.

To mitigate intra-/inter-flow interference, several complimentary approaches have been suggested (although they may not necessarily be proposed in the context of wireless mesh networks). For example, power control (a.k.a. topology control) aims to enable each node to transmit with the minimal possible transmit power, subject to network connectivity [18–21]. With each node transmitting with the minimal possible power, the interference due to concurrent transmission is mitigated. Another (orthogonal) approach is to control the carrier sense threshold (with which a node determines whether the shared wireless medium is busy or idle) [**?**]. By having each node use a large carrier sense threshold, more concurrent connections can take place simultaneously (at the expense of decreased SNIR and hence higher decoding failure rates). A third approach is centered at the notion of *channel diversity*, and equips each node with one or more radios. Concurrent transmission is made possible by having neighboring nodes transmit (and their corresponding receiver nodes receive) at different (non-overlapping) channels.

In this paper, we consider the problem of mitigating interference and improving network capacity in wireless mesh networks from the angle of *spatial diversity*. In a nutshell, while the achievable throughput on a multihop wireless path is limited by intra-flow interference, the overall capacity of a multihop wireless network can be increased by exploiting spatial diversity that exists among a number of multihop paths. Connections that are routed along these paths can be scheduled to take place simultaneously if their transmissions do not interfere with each other (significantly). To make a case of exploiting spatial diversity, we focus on transporting *downstream* traffic at gateway nodes with Internet access. This is because (i) most of the Internet accesses are intended for downloading large video/audio/text files; and (ii) by virtue of the way how wireless mesh networks operate, all the downloaded traffic is handled by gateway nodes. As a result, the downstream throughput at gateway nodes affects most significantly the performance as perceived by users.

There are two major issues that must be addressed in order to realize the notion of spatial diversity: (i) how to reason about the level of inter-flow interference among different paths; and (ii) how to schedule frame transmission among connections that are believed to incur the least level of interference. To address the first issue, we propose to construct, based on measurements of received signal strengths, a *virtual coordinate system*. This is in contrast to most existing work which relies on geographic locations of mesh nodes. With the coordinate system derived with principal component analysis, we will be able to determine the sets of paths along which transmissions can take place with the least inter-flow interference. To address the second issue, we coordinate, based on the sets of non-interfering paths, the order with which a gateway node schedules frames of different connections to be transmitted. To allow a gateway node to send frames consecutively in an non-interruptible manner, we leverage the *transmission opportunity (TXOP)* option in the IEEE 802.11e specification [7]. That is, a gateway node that succeeds in grasping the medium is granted the right to use the medium for a period of time specified by TXOP. The gateway uses a TXOP to transmit multiple frames, with SIFS (instead of DIFS) as the inter-frame space between the sequence of DATA-ACK exchanges. If the DATA-ACK exchange has been completed, and there is still time remaining in the TXOP, the node may transmit another frame (after an idle time of SIFS), provided that the frame to be transmitted and its necessary acknowledgment can fit into the time remaining in the TXOP. Through extensive simulation (with real-life measurement traces on CUWiN [2], we show that (1) the virtual coordinate system constructed with the received signal strength infers interference much better than the geographical locations of nodes (as was used in most existing work); and (2) with the proposed approach, the downstream throughput of a gateway node in a wireless mesh network can be improved by 10-55% under a variety of network topologies and traffic distributions. This, coupled with the fact that the proposed approach needs only minor code change in the gateway nodes and does not require additional hardware, makes it a viable option to improving network capacity in existing wireless mesh networks.

The reminder of the paper is organized as follows. In Section 2, we give a succinct summary of IEEE 802.11 and its associated intra-/inter-flow interference problems. We also motivate our proposed work with an illustrative example. In Section 3, we elaborate on how gateway nodes collect necessary information of received signal strengths and construct the coordinate system. In Section 4, we discuss the procedures with which gateway nodes coordinate the order of frame transmission among different connections. We also present an implementation with IEEE 802.11e. This is then followed by a discussion of related work in Section 4 and the performance evaluation (with real-life measurement traces on CUWiN) in Section 6. Finally, we conclude the paper in Section 7 with a list of future research agendas.

2. BACKGROUND AND MOTIVATION

2.1 An Overview of IEEE 802.11 DCF

We consider a stationary, multihop wireless network that operates on IEEE 802.11 distributed coordination function (DCF). The basic access method of DCF is carrier sensing multiple access with collision avoidance (CSMA/CA). A node that intends to transmit senses the channel and defers its transmission while the channel is sensed busy. When the channel is sensed idle for a specific time interval, called *distributed inter-frame space* (DIFS), the sender again waits for a contention window, CW (determined by the binary exponential backoff algorithm) and then transmits a data frame. After the data frame is received without errors, the receiver sends an acknowledgment frame to the sender after a specified interval, called the *short inter-frame space* (SIFS), that is less than DIFS. If an acknowledgment frame is not received, the data frame is presumed to be lost, and a retransmission is scheduled. The value of CW is set to CW_{min} in the first transmission attempt, and is doubled at each retransmission up to a pre-determined value CW_{max}. Retransmissions for the same data frame can be made up to a pre-determined retry limit, L, times. Beyond that, the pending frame will be dropped.

In addition to the above basic method, the request-to-send/clear-to-send (RTS/CTS) mechanism is optionally used, in the hope to solve the hidden/exposed terminal problems [10]. A node that intends to transmit first transmits a short control frame, called *Request To Send* (RTS), after sensing the medium to be idle for a DIFS interval. The RTS frame includes the source, the destination, and the duration it takes to transmit the data and ACK frames. The receiver node then responds with a control frame, called *Clear to Send* (CTS), that includes the same duration information. All the other nodes receiving either the RTS and/or CTS frames set their virtual carrier sense indicator, called *Network Allocation Vector (NAV)*, to the duration given in the RTS and/or CTS frames, and use it together with the physical carrier sense in determining whether or not the channel is idle.

It has been shown in the literature that the RTS/CTS floor acquisition mechanism handles the hidden/exposed *sender* problem well, but fails to solve the hidden/exposed *receiver* problem [10, 15, 27]. As a result, the RTS/CTS mechanism is often turned off in practice for the sake of improving the throughput performance.

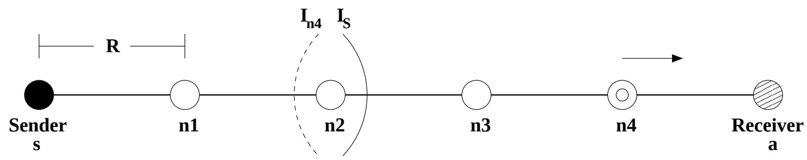

(a) A single multihop path.

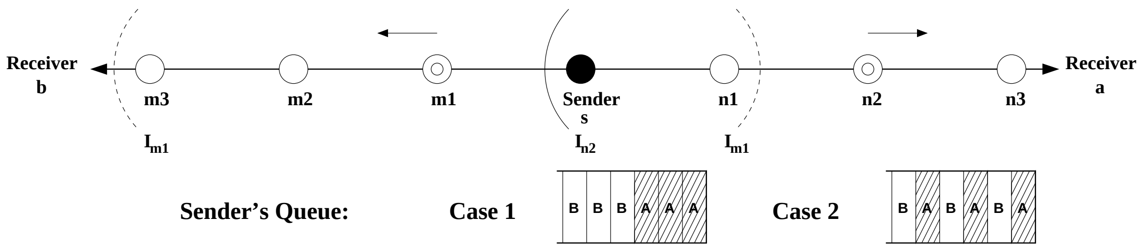

(b) Two multihop paths in the opposite directions.

Figure 1: Illustration of spatial diversity in multihop wireless networks.

2.2 An Illustrative Example

Figure 1 illustrates how the interference (caused by the hidden/exposed terminals) affects the network throughput in multiple-hop wireless network. In Figure 1, adjacent nodes are within the transmission range of each other, and the interference range is approximately twice of the transmission range. Let R and $I_{\mathbf{x}}$ denote, respectively, the transmission range and the set of nodes within the interference range of node $\mathbf{x}$. For example, $I_{\mathbf{s}} = \{\mathbf{s}, \mathbf{n1}, \mathbf{n2}\}$. As shown in Figure 1(a), sender $\mathbf{s}$ transmits a sequence of data frames to receiver $\mathbf{a}$. Due to intra-flow interference, the number of concurrent transmissions that can take place without interfering with one another is limited. For example, when $\mathbf{s}$ is transmitting frames to $\mathbf{n1}$ (denoted by $\mathbf{s} \longmapsto \mathbf{n1}$), only $\mathbf{n4}$ can transmit simultaneously without interfering with the transmission of $\mathbf{s} \longmapsto \mathbf{n1}$. This restriction on the number of concurrent transmissions makes it difficult for a connection to attain high throughput along a single multihop path.

In a large network consisting of more than one multiple-hop paths, it is, however, possible to increase the number of concurrent transmissions by exploiting spatial diversity. Figure 1(b) shows a simple wireless network with two connections ($\mathbf{s} \rightarrow \mathbf{a}$ and $\mathbf{s} \rightarrow \mathbf{b}$) along two multihop paths in opposite directions. Now we consider two extreme cases at $\mathbf{s}$ (Figure 1 (b): i) the frames destined for $\mathbf{a}$ are transmitted first, followed by those destined for $\mathbf{b}$ and ii) the frames destined for $\mathbf{a}$ interleave with those destined for $\mathbf{b}$. In case (i), spatial diversity that exists along the two different paths cannot be exploited, and the number of concurrent transmissions is almost the same as that in Figure 1(a). This is because after the first frame of connection $\mathbf{s} \rightarrow \mathbf{a}$ is transmitted by $\mathbf{s}$, it will complete with subsequent frames from $\mathbf{s}$ until it is at least three hops away from any subsequent frame (e.g., it reaches $\mathbf{n_4}$ while the second frame of $\mathbf{s} \rightarrow \mathbf{a}$ is still at $\mathbf{s}$).

On the other hand, in case (ii) of Figure 1(b), as the frames destined for $\mathbf{a}$ interleave with those destined for $\mathbf{b}$, it is possible for concurrent transmissions to take place without interfering with each other. For example, once the first frames destined for $\mathbf{a}$ and $\mathbf{b}$ reach, respectively, $\mathbf{n_2}$ and $\mathbf{m_1}$, the transmission $\mathbf{n_2} \longmapsto \mathbf{n_3}$ can take place concurrently with $\mathbf{m_1} \longmapsto \mathbf{m_2}$. As a matter of fact, if $\mathbf{m_1}$ and $\mathbf{n_1}$ initiates the transmission at approximately the same time (and the RTS/CTS mechanism is not exercised), the transmissions $\mathbf{m_1} \longmapsto \mathbf{m_2}$ and $\mathbf{n_1} \longmapsto \mathbf{n_2}$ may simultaneously take place.

As shown in the above example, intra-/inter-flow interference is dependent both upon how traffic is distributed and how traffic is routed along potentially interfering paths. As a result, to explore spatial diversity, we have to consider the following two issues: (i) how to reason about the level of inter-flow interference among different paths, and find sets of paths along which frames can be transmitted concurrently with the least inter-flow interference; and (ii) how to schedule frame transmission among connections with the least inter-flow transmission. To address the first issue, we will construct in Section 3, based on measurements of received signal strengths, a virtual, PCA-based coordinate system. Then we elaborate in Section 4 on how we coordinate the order in which frames of different connections are scheduled

3. DETERMINING VIRTUAL COORDINATES BASED ON RSSS

To make a case of exploiting spatial diversity to improve network capacity, we focus on transporting *downstream* traffic at gateway nodes with Internet access. This is because most of the gateway nodes are responsible for transporting a large amount of downstream traffic, and how they schedule transmission of frames to downstream mesh nodes will have a significant impact on the performance as perceived by users.

To infer the level of inter-flow interference among different paths, one may choose to use geographic locations of (next-hop) nodes as the references. Although this information can be readily obtained by GPS or as part of the static network configuration, it is sometimes quite misleading. For example, even though two next-hop nodes are geographically close to each other, the interference may not be significant if there is an obstacle between them. Due to this reason, we propose to exploit received signal strength (RSS) measurements among neighbors (to be defined below) as the references. This is because RSS measurements are more "representative" in determining the level of interferences between nodes. Moreover, they can be readily obtained through the sensory functions which are implemented in most of the modern IEEE 802.11 interfaces. Based on the RSS measurements among neighbors, we will construct a *virtual coordinate system* and use the "virtual distance" between mesh nodes to infer the level of interferences between them.

3.1 Measuring RSSs Between Mesh Nodes

To find the sets of paths with the least inter-flow interference, we will construct, for each gateway node (GN), a virtual coordinate system centered at the GN. The first step to constructing such a virtual coordinate system is to instrument nodes that can communicate with the GN *directly* or through a "relay" node in between to perform RSS measurements. Each such node periodically transmits hello packets, measures RSSs from the other nodes, and reports the measurements to its neighbors. In some sense, the RSS measurements are performed within two hops of the GN. If we increase the measurement area to within h hops from the GN ($h \geq 2$), it is possible to infer the level interference between nodes that are far from the GN and find the sets of paths with the least interferences more accurately. However, this is at the expense of higher control overhead and higher complexity of the algorithm that selects the set of interference-free paths (to be discussed in Section 3.3). For practicality and scalability, we restrict the measurement area to be within two hops from from the GN.

Through exchange of hello packets, a GN $\mathbf{n}$ gathers the RSS measurement between a node $\mathbf{m}$ that can directly communicate with itself and that between a neighbor node of $mathbfm$'s and $mathbfm$. Let $\mathcal{M}(n)$ denote the set of neighbor nodes that can directly communicate with $\mathbf{n}$ and $\mathbf{n}$ itself. Then, the RSS measurements between $\mathbf{n}$ and nodes in $\mathcal{M}(n)$ can be written in an $p \times p$ square matrix $\mathbf{S} = [s_{ij}]$ for $i, j \in \{1, \cdots, p\}$, where s_{ij} is the (– RSS) [1] measurement made in dBm by the ith node to jth node, $s_{ii} = 0$, and $p = |\mathcal{M}(n)|$ is the number of nodes. $\mathbf{S}$ is termed as the signal strength matrix, and will be used to construct a virtual coordinate system centered around $\mathbf{n}$.

One point that is worthy of mentioning is that all the components of $\mathbf{S}$ may not be available because the distances between some neighbor nodes may be occasionally larger than the wireless transmission range. Consider Figure 2 for example: both nodes ℓ and $\mathbf{m}$ are the neighbor nodes of the GN, but each node is located outside the transmission range of the other node. We will further discuss in Section 3.2 how we deal with the case that some of the pairwise RSS measurements are not available.

After the coordinates of nodes in $\mathcal{M}(n)$ are obtained, one can further derive the coordinate of a node $\mathbf{k}$ that cannot directly communicate with $\mathbf{n}$ but can via relay node(s) in $\mathcal{M}(n)$, with the use of RSS measurements made between $\mathbf{k}$ and relay nodes that can directly communicate with node $\mathbf{n}$. In some sense, construction of the virtual coordinate system proceeds in a ring-by-ring manner, starting from the innermost ring composed of $\mathbf{n}$'s neighbor nodes and proceeding outward.

3.2 Constructing Virtual Coordinate System

[1] As the sign of the RSS measurement is negated, a smaller value of s_{ij} implies stronger signal strength.

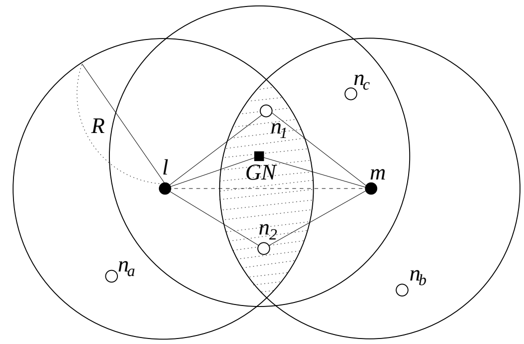

Figure 2: Extrapolation of RSS in the case that two nodes among the neighbor nodes are out of the other node's transmission range.

The RSS measurements between $\mathbf{n}$ and its neighbors are represented by the $p \times p$ square matrix $\mathbf{S}$, the columns of which can be considered as the coordinates of the corresponding nodes in a p-dimension space. Note that the ith column vector of $\mathbf{S}$ is the RSSs measured by ith node to all the nodes in $\mathcal{M}(GN)$. As these coordinates are correlated with each other, it is difficult to identify components that play an important role in determining the interferences. Hence we propose to construct, with principal component analysis (PCA), an Euclidean coordinate system with a smaller dimension . In a nutshell, PCA transforms a data set that consists of a large number of (possibly) correlated variables to a new set of uncorrelated variables, principal components. The principal components are ordered so that the first several components have the most important features of the original attribute variables. Moreover, the kth principal component can be interpreted as the direction of maximizing the variation of projections of measured attributes while orthogonal to the first $(k-1)$th principal components [28].

The most common approach to determining principal components is singular value decomposition (SVD). Specifically, the SVD of $\mathbf{S}$ is obtained by

$$\mathbf{S} = \mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^T, \quad (1)$$

$$\mathbf{W} = \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_p \end{bmatrix},$$

where $\mathbf{U}$ and $\mathbf{V}$ are column and row orthogonal matrices, and σ_i's are the singular values of $\mathbf{S}$ in the decreasing order (i.e., $\sigma_i \geq \sigma_j$ if $i < j$). Note that $\mathbf{S}^T\mathbf{S} = (\mathbf{U}\mathbf{W}\mathbf{V}^T)^T(\mathbf{U}\mathbf{W}\mathbf{V}^T) = \mathbf{V}(\mathbf{W}^T\mathbf{W})\mathbf{V}^T$. This means that the eigenvectors of $\mathbf{S}^T\mathbf{S}$ make up $\mathbf{V}$ with the associated (real nonnegative) eigenvalues of the diagonal of $\mathbf{W}^T\mathbf{W}$ [24]. Similarly, $\mathbf{S}\mathbf{S}^T = \mathbf{U}^T(\mathbf{W}\mathbf{W}^T)\mathbf{U}$. The columns of the $p \times p$ matrix $\mathbf{U} = [\ \mathbf{u}_1, \ldots, \mathbf{u}_p\ ]$ are the principal components and the orthogonal basis of the new subspace. By using the first q columns of $\mathbf{U}$ denoted by $\mathbf{U}_q$, we project the p-dimensional space into a new q-dimensional space:

$$\mathbf{c}_i = \mathbf{U}_q^T \cdot \mathbf{s}_i = [\ \mathbf{u}_1,\ \ldots,\ \mathbf{u}_q\ ]^T \cdot \mathbf{s}_i, \quad (2)$$

where $\mathbf{c}_i$ is the new coordinate of the ith node, and $\mathbf{s}_i$ is the ith column vector of $\mathbf{S}$.

After obtaining the coordinates, we have to perform two post-processing operations. The first operation is to translate the coordinate system so that the GN (i.e., $\mathbf{c}_{GN}$) becomes the origin. The second operation is to scale the coordinate system so that the distance between two nodes coincides with the corresponding, actual measured signal strength. The optimal scaling factor α^* that minimizes the discrepancy between the Euclidean distance and the measured signal strength can be determined by minimizing the following objective function

$$J_1(\alpha) = \sum_i^p \sum_j^p \left(L(\alpha \mathbf{c}_i, \alpha \mathbf{c}_j) - s_{ij} \right)^2, \qquad (3)$$

where the L is the Euclidean distance between two vectors (i.e., $L(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x}-\mathbf{y})^T(\mathbf{x}-\mathbf{y})}$). After a few algebraic operations, the positive solution, α^*, can be shown to be

$$\alpha^* = \frac{\sum_i^p \sum_j^p d_{ij} L(\mathbf{c}_i, \mathbf{c}_j)}{\sum_i^p \sum_j^p L(\mathbf{c}_i, \mathbf{c}_j)^2}. \qquad (4)$$

The new coordinate of a node is written by

$$\mathbf{x}_i = \alpha^* \left(\mathbf{c}_i - \mathbf{c}_{GN} \right). \qquad (5)$$

Determining Coordinates for Nodes That Are Two Hops Away

Recall that as shown in the illustrative example in Section 2.2, if two nodes are outside the transmission range, but within the interference range, of each other, their transmissions interfere each other, but cannot be directly measured by either node without the support of intermediate nodes. As a result, the GN has to be aware of the level of interference with two hops in order to exploit spatial diversity (Fig. 1 (b)). That is, the virtual coordinate system has to be "extended" to include the coordinates of nodes that are two hops away.

As discussed in Section 3.1, a neighbor node of the GN measure RSSs from its neighbor nodes. For a node $\mathbf{k}$ that is two hops from GN, if $l = |\mathcal{M}(\text{GN}) \cap \mathcal{M}(\mathbf{k})| \geq q+1$, we obtain the coordinate of $\mathbf{k}$, $\mathbf{x}_k$, by minimizing the following objective function:

$$J_2(\mathbf{x}_k) = \sum_{i \in \mathcal{M}(\text{GN}) \cap \mathcal{M}(\mathbf{k})} \left(L(\mathbf{x}_i, \mathbf{x}_k) - s_{ik} \right)^2. \qquad (6)$$

Because solving this non-linear optimization problem incurs high computational complexity, we use a alternative, non-iterative multilateration algorithm [17]: From the quadratic version of the Euclidean distance equations (i.e., $L(\mathbf{x}_k, \mathbf{x}_i)^2 = \sum_{j=1}^q (x_k^j - x_i^j)^2 = s_{ik}^2$ and x^j is the jth component of $\mathbf{x}$.), a linear system is derived by subtracting one of the equations from the other equations.

$$\mathbf{A}\mathbf{x_k} = \mathbf{b}$$

and

$$\mathbf{A} = 2 \begin{bmatrix} (\mathbf{x}_1 - \mathbf{x}_l)^T \\ \vdots \\ (\mathbf{x}_{l-1} - \mathbf{x}_l)^T \end{bmatrix}$$

$$\mathbf{b} = \begin{bmatrix} \mathbf{x}_1^T \mathbf{x}_1 - s_{1k}^2 \\ \vdots \\ \mathbf{x}_{l-1}^T \mathbf{x}_{l-1} - s_{(l-1)k}^2 \end{bmatrix} - (\mathbf{x}_l^T \mathbf{x}_l - s_{lk}^2)\, \mathbb{1}_{l-1},$$

where $\mathbb{1}_{M-1} \in \mathbb{R}^{M-1}$ is the 1's column vector. The coordinate is given by the least square solution of the linear system, i.e.,

$$\mathbf{x}_k = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}. \qquad (8)$$

Figure 3: How to determine the coordinate of a node that is two hops away from a GN when the number of RSS measurements is not sufficient.

It requires neither a judicious guess of the initial location nor a time-consuming iteration process to solve Eq. (6).

In the case that node $\mathbf{k}$ does not have sufficient neighbor nodes (i.e., $l < q+1$), it is not possible to uniquely determine the coordinate of $\mathbf{k}$. In this case, we obtain the coordinate by making the assumption that the signal strength is inversely proportional to d^β, where d is the geographical distance between the nodes and β is the path loss coefficient ($2 \leq \beta \leq 4$). Figure 3 illustrates an example in which the dimension of the coordinate system is two, and only two RSS measurements $s_{(GN)i}$ and s_{ik} are available. If the geographic distances between nodes, $d_{(GN)k}$ and $d_{(GN)i}$), are known, the RSS between GN and $\mathbf{k}$ is computed under the assumption as

$$s_{(GN)k} = s_{(GN)i} + \beta \log\left(\frac{d_{(GN)k}}{d_{(GN)i}}\right).$$

Then, we have two possible coordinates for $\mathbf{k}$, labeled as $\mathbf{x}_k$ in Figure 3, which are symmetric with respect to the corresponding one-hop neighbor, $\mathbf{x}_i$, of the GN. Since the symmetry holds in general, we simply take a vector whose direction is the same as $\mathbf{x}_i$ and whose magnitude is $s_{(GN)i} + \beta \log(\frac{d_{(GN)k}}{d_{(GN)i}})$ for $\mathbf{x}_k$. That is,

$$\mathbf{x}_k = \left(1 + \frac{\beta \log \frac{d_{(GN)k}}{d_{(GN)i}}}{|\mathbf{x}_i|} \right) \mathbf{x}_i. \qquad (10)$$

Note that this approach renders an approximate estimate of x_k when there exists obstacles between the GN and node $\mathbf{k}$ and hence the assumption does not hold.

3.3 Identifying Paths with Least Interference

We can infer the interferences between mesh nodes by computing the Euclidean distances in the coordinate system. If the distance between two nodes is large, it implies that they will not interfere with each other's transmission. In this subsection, we devise an algorithm that selects, based on the topological relationship between nodes in the coordinate system, paths along which packets can be transmitted concurrently with the least inter-flow interference.

Recall that as shown in Figure 1(b), it is not possible to avoid inter-flow interference near $\mathbf{s}$ because $\mathbf{s}$, $\mathbf{m1}$, $\mathbf{n1}$ are so close to each other. However, if the interference range is f times larger than the transmission range and $2 < f \leq 3$, then both $\mathbf{m1}$ and $\mathbf{n2}$ (or $\mathbf{m2}$ and $\mathbf{n1}$) can transmit at the same time. Motivated by this example, we take into account of the first two hops of paths in determining the set of connections whose frames can be transmitted concurrently. In some sense, we determine the set of first relay nodes that give the least interference when packets are being forwarded by them.

Figure 4 illustrates how to inter the interference between nodes that are two hops away from the GN with the use of their coordinates. In the figure, the coordinates of the 1-hop neighbor nodes of a GN are labeled as $\mathbf{x}_i$ and $\mathbf{x}_j$ and calculated in Eq. (5), while those of the corresponding two-hop neighbor nodes are labeled as $\mathbf{x}_{i'}$ and $\mathbf{x}_{j'}$ and calculated in Eqs. (8) or (10). The signal-to-interference ratio (SIR) at $\mathbf{x}_{i'}$ due to the interference from $\mathbf{x}_{j'}$, denoted by $\text{SIR}_{i'}(j')$,

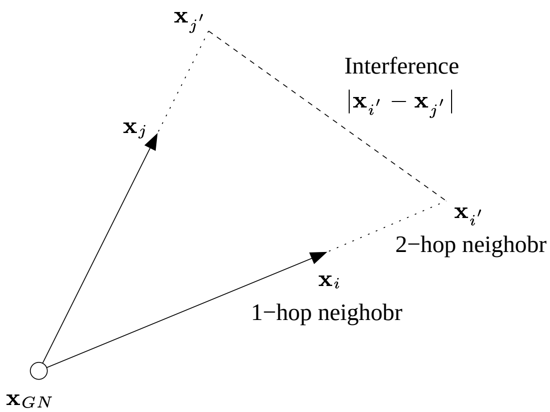

Figure 4: How to infer interference between nodes that are 2 hops away from a gateway node with the use of their coordinates.

is given by

$$\mathrm{SIR}_{i'}(j') = |\mathbf{x}_{i'} - \mathbf{x}_{j'}| - s_{i'i}. \quad (11)$$

For each pair of nodes, if SIR is larger than an SIR threshold γ (in dB), we consider the pair as one with negligible inter-flow interference. We construct an interference table T, of which the ijth entry is given by

$$t_{ij} = \begin{cases} 1 & \text{if } \mathrm{SIR}_i(j) > \gamma, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

We refer to the table T to determine whether the transmission of a node will interfere with another. It is important to note that the probing process is locally performed among neighbor nodes of a GN and that of constructing the interference table only at the GN.

One issue that is worthy of discussion is the availability of routing information. To infer the potential interference between transmissions at two nodes that are two hops away from the GN (Eq. (12)), it is required that the routing information up to two hops be available. While the next node $\mathbf{m}$ to which a frame is to be transmitted is usually available at GN, whether or not the next node of $\mathbf{m}$ is available to the GN depends upon the selection of the routing algorithm. For example, under the dynamic source routing (DSR) algorithm, the GN maintains the route cache for source routing and can obtain the two-hop node information without any overhead. In the case that the two-hop node information is not available at the GN, it can be piggy-backed as part of the hello packets.

4. COORDINATING ORDER OF TRANSMISSION

To enable the GN to coordinate the order of transmission and fully explore spatial diversity, the GN should be given the chance to transmit a sequence of frames, once it grasps the medium. For example, in Figure 1 the best way to increase the level of spatial diversity is to enable node $\mathbf{s}$ to transmit frames alternatively to each receiver. If after the GN transmits a frame to $\mathbf{n1}$ it cannot continue holding the medium and transmit a frame to $\mathbf{m1}$ until the first frame reaches $\mathbf{n3}$, then the number of concurrent transmissions essentially does not increase. On the other hand, if the GN alternatively transmits frames to $\mathbf{m1}$ and $\mathbf{n1}$, respectively, then after the two frames are more than two hops away from each other (e.g., at $\mathbf{n_2}$ and $\mathbf{m_1}$, respectively, or at $\mathbf{n_1}$ and $\mathbf{m_2}$, respectively), they can be transmitted concurrently. Figure 5 shows one possible scenario

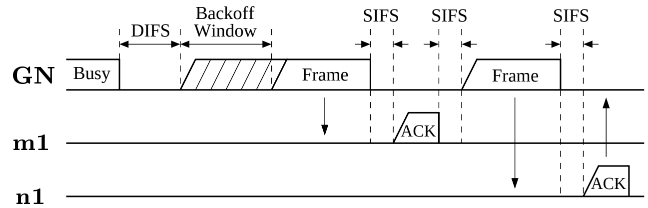

Figure 5: A scenario in which the GN coordinates transmission between itself and its neighbor nodes $\mathbf{m1}$ and $\mathbf{n1}$ in Fig. 1.

of how the GN coordinates the transmission between itself and its neighbor nodes $\mathbf{m1}$ and $\mathbf{n1}$. After the channel is sensed free for a DIFS time interval the GN waits for a random backoff interval and transmits the first frame to $\mathbf{m1}$. After that, the GN chooses a frame that does not interfere with transmission of the first frame (in the near future) and transmits it without backoff.

To select subsequent frames that give the least interference, we leverage the interference table T computed in Section 3.3. Specifically, after sending the first frame, the GN looks up a candidate frame from the head of the queue in the logical link layer (LLC). Let Λ denote the set of neighbor nodes to which frames were sent after the GN grasps the medium. For example, $\Lambda = \{\mathbf{m1}\}$ after the first frame is transmitted in Figure 5. The GN looks up to N frames in the LLC queue in order to locate a frame f that satisfies $t_{routing(f)i} = 1$ for $\forall\, i \in \Lambda$, where $routing(\cdot)$ is the function that returns the next hop of a frame.

When the GN finishes transmitting the last frame eligible for transmission based on T, it relinquishes the medium and the neighbor nodes will complete for the medium to relay frames that were sent to them by the GN. In order to give more opportunities to its neighbors, the GN is instrumented to set a contention window size that is larger than that originally specified in IEEE 802.11 DCF. Specifically, the contention window is set to the sum of the backoff intervals that were skipped when the GN is in possession of the medium. That is, the contention window is set to

$$CW = \min(|\Lambda| \times CW_{min},\ CW_{max}), \quad (13)$$

where CW, CW_{min}, and CW_{max} are the current, maximum, and the minimum contention window sizes, respectively. In the case that the every frame in the queue in the LLC layer interferes with those transmitted earlier, only one frame is transmitted (i.e., $|\Lambda| = 1$), and CW in Eq. (13) essentially falls back to that originally specified in IEEE 802.11.

4.1 Algorithms for Coordinated Transmission in LLC/MAC Layers

The pseudo-code of the proposed algorithm is given in Algorithms 1, 2, and 3. They should be, respectively, implemented in the LLC layer (Algorithm 1) and the MAC layer (Algorithm 2) at a GN.

Algorithm 1 outlines how the LLC layer operates to support the coordinated transmission in the MAC layer. The LLC layer selects, based on Λ and T (computed in Section 3.3), a frame whose transmission will not interfere with the frames sent earlier in this run of coordinated transmission. The LLC layer inspects in sequence N frames from the head of the queue. If the queue is not empty and an appropriate frame is identified, the LLC layer updates Λ and passes the frame onto the MAC layer.

Algorithm 1 LLC layer implementation for GN.

```
if the queue is not empty then
  is_the_first_frame ← true
  look up to N frames in the queue and select the set F
      of eligible frames based on Λ and T computed in Section
  3.3
  while F ≠ NULL do
    retrieve a frame p from F
    p.is_the_first_frame ← is_the_first_frame
    is_the_first_frame ← false
    insert the next-hop node of p into Λ
    invoke Algorithm 2 and wait for its return
  end while
end if
```

Algorithm 2 MAC layer implementation for GN.

```
Request a frame p from the LLC layer.
if p.is_the_first_frame then
  Λ ← ϕ // a set of neighbor nodes
  pre_bulk_count ← bulk_count
  bulk_count ← 0
  transmit_with_backoff(p, 0)
else
  // bulk transmission without backoff
  transmit frame p after an SIFS interval
  if an ACK frame is received then
    bulk_count++
  else
    // if corrupted, transmit again with backoff
    transmit_with_backoff(p, 1)
  end if
end if
```

Algorithms 2 and 3 show how the MAC layer implements the proposed algorithm. Algorithm 2 is called whenever a frame arrives from the LLC layer. If the frame is the first for this run of coordinated transmission, the transmit_with_backoff subroutine in Algorithm 3 is called to transmit the frame after a random backoff interval. The backoff value is randomly drawn from the range of $[0, CW - 1]$ and CW is computed by Eq. (13).

When a GN detects collision in the middle of coordinated transmission, it follows the binary exponential backoff algorithm defined in IEEE 802.11 DCF (Algorithm 3). Even in this case, the GN does not terminate its bulk transmission but continues with an additional random backoff interval to avoid potential collisions. If a frame fails to be transmitted for retransmission_limit times, it is dropped but the GN continues with its bulk transmission and attempts to transmit the next frame after an SIFS interval.

4.2 Implementation with IEEE 802.11e EDCA

The way the proposed algorithm sets the contention window size (Eq. 13) deviates from that specified in IEEE 802.11 DCF. Moreover, the proposed algorithm also requires that a GN be granted an extended interval (sufficient to transmit up to N frames) after it grasps the medium. Fortunately these functions are defined and available in IEEE 802.11e *Enhanced Distributed Channel Access (EDCA)* draft [16]. In what follows, we first give a succinct overview of EDCA and then discuss how we will leverage the functions provided in EDCA to implement the proposed algorithm.

Algorithm 3 Subroutine – transmit_with_backoff(**p**, retx).

```
// retx denotes the number of retransmissions tried
if retx = 0 then
  CW ← min(pre_bulk_count × CW_min, CW_max)
  bulk_count ++
else
  CW ← min(2^n × CW_min − 1, CW_max)
end if
backoff_counter ← rand(0, CW)

// contention period
while backoff_counter > 0 do
  for every idle T_aSlotTime do
    backoff_counter ← backoff_counter − 1
  end for
end while

transmit frame p
if an ACK frame is received then
  return
else if retx < retransmission_limit then
  transmit_with_backoff(p, retx+1)
else
  drop frame p
  return
end if
```

IEEE 802.11e EDCA has been proposed to support quality-of-service (QoS) in WLANs. In EDCA, several parameters control how and when a node gains access to the medium among different priority levels (called access categories (ACs)), so as to favor/disfavor data transmission from high-priority/low-priority flows. These parameters include the minimum idle delay before contention (AIFS), the minimum and maximum contention windows (CW_{min} and CW_{max}), and the transmission opportunity limit (TXOP). In particular, EDCA associates different ACs with different values of CW_{min} and CW_{max}, and allows traffic of different priorities to back off for different time intervals, so as to increase/decrease their probability of medium access. Also, a station that succeeds in grasping the medium is granted a transmission opportunity (TXOP) — the right to use the medium and transmit multiple frames without backoff. The TXOP value differs for different ACs. If the DATA-ACK exchange sequence has been completed, and there is still time remaining in the TXOP, the station may transmit another frame in the same access category, provided that the frame to be transmitted and its necessary acknowledgment can fit into the time remaining in the TXOP. All the parameters can be dynamically updated by the access point (AP) through the EDCA parameter set, and are sent from the AP as part of the beacon frames, and probe/re-association response frames. This adjustment allows stations in the WLAN to adapt to changing conditions, and gives the AP the ability to manage the overall QoS performance.

We can readily leverage the functions of setting CW_{min} and TXOP values in IEEE 802.11 EDCA to implement the proposed algorithm. The GN sets its CW_{min} value to the value specified in Eq. (13) before the contention period starts, and sets it back to the previous value after the contention period finishes. Moreover, the GN sets its TXOP to an interval sufficient to transmit N frames. In this manner, a GN can transmit up to N frames consecutively within a TXOP.

5. RELATED WORK

6. PERFORMANCE EVALUATION

7. CONCLUSION

8. REFERENCES

[1] Bay area wireless user group. http://www.bawug.org/.

[2] Champaign-urbana community wireless network. http://www.cuwireless.net/.

[3] MIT roofnet. http://pdos.csail.mit.edu/roofnet/doku.php.

[4] Seattle wireless. http://www.seattlewireless.net/.

[5] SFLan. http://www.sflan.org/.

[6] Southampton open wireless network. http://www.sown.org.uk/.

[7] I. 802.11e/D13.0. Draft supplement to part 11: Wireless medium access control (MAC) and physical layer (PHY) spe cifications: Medium access control (MAC) enhancements for quality of service (QoS), Jan. 2005.

[8] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *IEEE BroadNets*, July 2004.

[9] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proc. of ACM SIGCOMM*. ACM, September 2004.

[10] D. Allen. Hidden terminal problems in wireless lan's. 802.11 Working Group paper 802.11/93-xx.

[11] D. Berger, Z. Ye, P. Sinha, S. Krishnamurthy, M. Faloutsos, and S. K. Tripathi. TCP-friendly medium access control for ad-hoc wireless networks: alleviating self-contention. In *Proc. of IEEE MASS*, October 2004.

[12] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an 802.11b mesh network. In *Proc. of ACM Mobicom*. ACM, September 2005.

[13] D. S. J. De Couto, D. Aguayo, J. C. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *ACM Mobicom*, 2003.

[14] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *Proc. of ACM SIGCOMM*. ACM, September 2004.

[15] Z. Fu, H. Luo, P. Zerfos, S. Lu, L. Zhang, and M. Gerla. The impact of multihop wireless channel on TCP performance. *IEEE Trans. on Mobile Computing*, 4(2):209–221, March/April 2005.

[16] IEEE Computer Society. Part 11: Wireless LAN Medium Access (MAC) and Physical Layer (PHY) Specifications, Amendment 8: Medium Access Cntorl (MAC) Quality of Service Enhancement. *IEEE Standard 802.11e*, 2005.

[17] K. Langendoen and N. Reijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Elsevier Computer Networks*, 43:499–518, 2003.

[18] L. Li, J. Y. Halpern, P. Bahl, Y. M. Wang, and R. Wattenhofer. Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 264–273, August 2001.

[19] N. Li and J. C. Hou. FLSS: A fault-tolerant topology control algorithm for wireless networks. In *Proc. ACM International Conference on Mobile Computing and Networking (MOBICOM)*, September 2004.

[20] N. Li and J. C. Hou. Topology control in heterogeneous wireless networks: Problems and solutions. In *Proc. of IEEE INFOCOM*, March 2004.

[21] N. Li, J. C. Hou, and L. Sha. Design and analysis of an MST-based topology control algorithm. In *Proc. of IEEE INFOCOM*, pages 1702–1712, April 2003.

[22] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau. A game theoretic approach to provide incentive and service differentiation in P2P networks. In *Proc. ACM SIGMETRICS*, New York, NY, June 2004.

[23] Microsoft Networking Research Group. Self-organizing neighborhood wireless mesh networks. http://research.microsoft.com/mesh/.

[24] B. Noble and J. W. Daniel. *Applied Linear Algebra*. Prentice Hall, 1988.

[25] Nortel Networks. The business case for wireless mesh networks. http://www.nortelnetworks.com/corporate/events/2003d/wmn_eseminar/collateral/wmn_eseminar.pdf, December 2003.

[26] K. Sanzgiri, I. D. Chakeres, and E. M. Belding-Royer. Determining intra-flow contention along multihop paths in wireless networks. In *Proc. of Broadnets Wireless Networking Symposium*, October 2004.

[27] S. Xu and T. Saadawi. Does the IEEE 802.11 MAC protocol work well in multihop wireless adhoc networks? *IEEE Communications Magazine*, 39(6):130–137, 2001.

[28] K. Y. Yeung and W. L. Ruzzo. Principal component analysis for clustering gene expression data. *Bioinformatics*, 17(9):763–774, 2001.