AFRL-RI-RS-TR-2008-136
Final Technical Report
May 2008

RESEARCH IN WIRELESS NETWORKS AND COMMUNICATIONS

Harvard University

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-RI-RS-TR-2008-136 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/	/s/
DAVID L. HENCH Work Unit Manager	WARREN H. DEBANY, JR. Technical Advisor, Information Grid Division Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
MAY 2008	Final	Jan 05 – Nov 07

4. TITLE AND SUBTITLE

RESEARCH IN WIRELESS NETWORKS AND COMMUNICATIONS

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA8750-05-1-0035

5c. PROGRAM ELEMENT NUMBER
62702F

6. AUTHOR(S)

H. T. Kung

5d. PROJECT NUMBER
CITE

5e. TASK NUMBER
RW

5f. WORK UNIT NUMBER
NC

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

President and Fellow of Harvard University
Harvard University
1350 Massachusetts Ave, Room 550
Cambridge MA 02138-3846

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RIGC
525 Brooks Rd
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TR-2008-136

12. DISTRIBUTION AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA # WPAFB 08-3196

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This project has yielded a number of results in wireless networking. We have developed algorithms of steering directional antennas in wireless mesh networks and for controlling radio interference in multi-hop wireless meshes. We have measured performance of 802.11a wireless links from an unmanned aerial vehicle (UAV) to ground stations. We have devised protocols for parallel use of multiple channels in a multi-radio, multi-hop 802.11 wireless network, with the goal of maximizing the total multi-hop throughput. We have proposed transmit antenna selection based on receiver feedback of channel information obtained via link layer probing. To validate these results, we have built a low-altitude UAV testbed using commodity 802.11 hardware and performed field experiments. Under this project, we have also obtained some mathematical results, such as the "hub theory," which provides a sufficient and necessary characterization on the effectiveness of some important communications and network systems.

15. SUBJECT TERMS

Antenna Selection, Radio Interference, UAV, Wireless Mesh, Hub Theory

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			David L. Hench
U	U	U	UU	50	19b. TELEPHONE NUMBER (*Include area code*) N/A

Table of Contents

Summary of Results and Conclusions	1
Publications	4
Constructing Multiple Wireless Meshes in the Same Region with Beam-Crossing Grids	5
Adjacent Channel Interference in Dual-Radio 802.11a Nodes and Its Impact On Multi-hop Networking	12
Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes With Various Antenna Orientations	18
Parallel Use of Multiple Channels in Multi-hop 802.11 Wireless Networks	24
Transmit Antenna Selection Based on Link-layer Channel Probing	33
A Hub Matrix Theory and Applications to Wireless Communications	39

Summary of Results and Conclusions

Work of this project has led to a number of findings and conclusions. Most of these results have been published. A list of publications is given in the next section. Reprints of published papers are attached to this report.

In the following paragraphs, we summarize our results.

Given a set of wireless access points (APs) with directional antennas, we have studied the problem of steering their antennas to form multiple non-interference wireless meshes in the same region. The resulting multiple wireless meshes provide not only a simultaneously usable, but also a fail-safe, wireless network infrastructure. We have developed a mesh construction algorithm, called "beam-crossing grids", that allows the constructed multiple wireless meshes to cover a large common area, while not interfere with each other. This means that, from locations in a large area, wireless terminals can have access to more than one of these meshes for improved network bandwidth and redundancy. By using analysis and simulation, we show that when the AP density is relatively high, the algorithm converges rapidly to a high-performance construction of wireless meshes

We have evaluated the performance impact of adjacent channel interference (ACI) in multi-hop wireless networks based on dual-radio 802.11a nodes. Although these nodes use chipsets that satisfy the transmit-mask requirements set by the IEEE 802.11 standard, the multi-hop performance is still significantly affected by ACI. That is, a node's transmitter can interfere with its own receiver on a different channel; as a result, multi-hop throughput is severely degraded. This degradation is especially pronounced for 802.11a. We use a spectrum analyzer with a signal combiner to quantify ACI under various conditions and propose solutions to mitigate the effect of such interference on multi-hop forwarding. Field experiments with multihop relay have validated these findings as well as the effectiveness of our solutions.

We have measured performance of 802.11a wireless links from an unmanned aerial vehicle (UAV) to ground stations. In a set of field experiments, we record the received signal strength indicator (RSSI) and measure the raw link-layer throughput for various antenna orientations, communication distances and ground-station elevations. By comparing the performance of 32 simultaneous pairs of UAV and ground station configurations, we are able to conclude that, in order to achieve the highest throughput under a typical flyover UAV flight path, both the UAV and the ground station should use omni-directional dipole (as opposed to high-gain, narrowbeam) antennas positioned horizontally, with their respective

antenna null pointing to a direction perpendicular to the UAV's flight path. In addition, a moderate amount of elevation of the ground stations can also improve performance significantly.

We have considered parallel use of multiple channels in a multi-radio, multi-hop 802.11 wireless network, with the goal of maximizing the total multi-hop throughput. We first quantify several fundamental forms of radio interference that cause performance degradation when the number of hops increases and that prevent total throughput from scaling up with number of radio interfaces at each node. We then evaluate three different methods of parallel channel use: Ad-Hoc, Frequency- Division Multiplexing (FDM), and Time-Division Multiplexing (TDM). We measure their performance on a linearly connected multi-hop network of dual-radio nodes. Although theoretically these three methods should have comparable performance, their actual measured performances are quite different. We find that TDM has the best performance, followed by Ad-Hoc and then FDM. The performance differences are due to these methods' capabilities of combating interference. We conclude that interference, especially adjacent channel interference, has significant effect on the achievable performance of a multi-radio, multi-hop network and hence should be carefully taken into account in the design and deployment of such a network.

We have proposed transmit antenna selection based on receiver feedback of channel information obtained via link layer probing. Furthermore, we report the performance gain of the proposed antenna selection scheme in an experimental multi-antenna 802.11 network. We built a low-altitude Unmanned Aerial Vehicle (UAV) testbed using commodity dual antenna 802.11 hardware and performed field experiments to collect traces of link performance using antennas of various types and orientations. Based on the collected traces, we demonstrate that transmit antenna selection can achieve a significant amount of gain using a link-layer channel probing protocol at a relatively low probing rate. The largest improvement we observed with joint transmit/receive antenna selection in 2×2 systems was 32%, about twice as much as that of receive-only antenna selection in 1×2 systems, which achieved 17%. Moreover, a similar improvement is obtained with probing intervals up to about 200 milliseconds, which is infrequently enough to consume only a small fraction of the available 802.11 channel capacity. Since these results require only a low implementation and operational cost, we conclude that transmit antenna selection is a worthwhile technique to use with the kind of multi-antenna mobile ad-hoc networks we examined.

In this effort, we have developed some mathematical theory. There are many communications and network systems whose properties are characterized by the gaps of the leading eigenvalues of $A^H A$ for a matrix A. We have shown that a sufficient and necessary

condition for a large eigen-gap is that A is a “hub” matrix in the sense that it has dominant columns. We have illustrated the applications of this hub theory in multiple-input and multiple-output (MIMO) wireless systems.

Some of these results have been summarized in a Harvard Ph.D. thesis by Dario Vlah.

Publications

"A Hub Matrix Theory and Applications to Wireless Communications," EURASIP Journal on Advances in Signal Processing, 2007

"Transmit Antenna Selection Based on Link-layer Channel Probing," Proceedings of the 2nd IEEE Workshop on advanced EXPerimental activities ON WIRELESS networks & Systems (EXPONWIRELESS 2007), June 2007

"Adjacent Channel Interference in Dual-radio 802.11a Nodes and Its Impact on Multi-hop Networking," Proceedings of IEEE Globecom, November 2006

"Parallel Use of Multiple Channels in Multi-hop 802.11 Wireless Networks," IEEE MILCOM 2006, October 2006

"Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations," Proceedings of the 15th International Conference on Computer Communications and Networks (ICCCN'06), October 2006

"Constructing Multiple Wireless Meshes in the Same Region with Beam-Crossing Grids," IEEE Wireless Communications and Networking Conference (WCNC), March 2005

"Antenna Selection Performance in 802.11 Networks," Harvard Ph.D. Thesis by Dario Vlah, October 2007

Constructing Multiple Wireless Meshes in the Same Region with Beam-Crossing Grids

Pai-Hsiang Hsiao and H. T. Kung
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
{*shawn, htk*}*@eecs.harvard.edu*

***Abstract*— Given a set of wireless access points (APs) with directional antennas, we consider the problem of steering their antennas to form multiple non-interferencing wireless meshes in the same region. The resulting multiple wireless meshes provide not only a simultaneously usable, but also a fail-safe, wireless network infrastructure. We describe a mesh construction algorithm, called "beam-crossing grids", that allows the constructed multiple wireless meshes to cover a large common area, while not interfere with each other. This means that, from locations in a large area, wireless terminals can have access to more than one of these meshes for improved network bandwidth and redundancy. By using analysis and simulation, we show that when the AP density is relatively high, the algorithm converges rapidly to a high-performance construction of wireless meshes.**

I. Introduction

A wireless mesh, as we mean in this paper, is a set of wireless access points (APs), where some pairs of them are directly connected with radio links. There has been an increasing interest in such wireless mesh networks due to their relatively ease of deployment. For example, Beyer [1] discusses the advantages of using such wireless mesh networks to provide residential broadband Internet access. Nortel Networks recently started trials of their wireless mesh networks with British Telecom and other organizations. They use APs equipped with smart antennas and integrated routers. [2]

Consider a set of APs placed in a 2-dimensional space. We assume that these APs use directional antennas for inter-AP wireless connections. Typically a wireless IP networking infrastructure can be formed using these APs in two steps. First, select pairs of APs to connect directly with radio links and then steer/configure their antennas to make the radio connections. Second, by using a routing protocol, compute a routing path (if it exists) between any two APs. These computed paths will then be used for transporting IP packets on the network.

A main concern of these wireless IP infrastructures is their reliability. Unlike wireline links, the performance of a wireless link can be affected by many environment-related factors, including path loss, shadowing, fading, interference from other users, and the non-linearity of underlying devices. A wireless infrastructure must tolerate not only permanent link failures, but also transient failures. Conventional routing protocols can recover from link failures, but their routing message overheads and end-to-end control delays often make them not suited to wireless infrastructures where bandwidths are more limited and transient link failures are frequent. To combat these problems, alternative routing methods have been proposed, such as those where multiple routes are maintained [3]–[5]. But these methods tend to be more complicated and difficult to manage.

To support a fail-safe wireless infrastructure capable of fast recovery from failure, in this paper we take the position that additional redundancy at the wireless link level ought to be provided. That is, the infrastructure should include multiple wireless meshes, so that any node or terminal can immediately switch over to another wireless mesh when detecting failure in the mesh that it currently uses.

We assume that we can configure and steer the directional antennas of the APs under software control (e.g., see [6] for such antennas). By steering the beams of these antennas, we will form multiple wireless meshes, which will cover as large a common area as possible, while not interfering with each other. We will use these multiple wireless meshes to provide redundancy at the link level. When there are link failures, the multiple wireless meshes can back up each other. When there is no failure, these multiple wireless meshes can operate in parallel to increase the total bandwidth of the wireless infrastructure.

We describe an algorithm, called *beam-crossing grids*, to construct multiple wireless meshes. The main insight of the algorithm, as its name suggests, is that it uses a grid of square regions, where two or triple crossing beams have been configured, as building blocks to form multiple wireless meshes. Given square regions with these crossing beams, it becomes relatively easy to connect them to form multiple meshes. In contrast, if square regions with parallel beams were used, then it tends to be the case that only one mesh can be constructed.

The paper is organized as follows. We first state the simplifying assumptions we make about the radios and antennas in Section II. We then describe the beam-crossing grids algorithm for constructing two wireless meshes in Section III. An upper bound on the running time of the algorithm is given in Section IV. We use simulations to evaluate the performance of this algorithm. The simulation results are reported in Section V. Section VI discusses the applications of our method, and its generalization in constructing more than two collocated

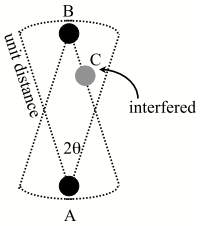

Fig. 1. Two ideally sectorized transmit beam patterns for inter-AP communication, where A and B are two communicating APs and C is an AP whose reception is interfered with by both A and B's transmission

networks in the region. We summarize this paper and provide conclusions in Section VII.

II. Simplifying Radio and Antenna Model

We make several simplifying assumptions on the radio and antenna models in this paper. We assume that inter-AP radio communication uses a dedicated frequency band which is different from those used by APs in communicating with wireless terminals. For inter-AP communication, the transmitting AP uses a directional antenna while the receiving AP uses an omnidirectional antenna. All APs produce the same transmit beam pattern, which is an ideally sectorized beam pattern, as illustrated in Fig. II. The length of the beam sector is of unit distance, and the angular width of the beam is 2θ. The gain inside the beam pattern is constant and is negligible both outside the pattern and on the boundary. To avoid interference, unintended AP receivers, e.g., C as shown in Fig II, cannot lie in the area of a transmit beam pattern of an AP. These assumptions are similar to those used by some typical models in the literature (see, e.g., [7]–[9]).

Given the assumptions about radio and antenna and a set of APs, we can specify a *connectivity graph*. In the connectivity graph, each vertex represents an AP, and two vertices are connected with an edge if the two corresponding APs can transmit to each other directly. In addition, we can also define an *edge conflict graph*. In the edge conflict graph, a vertex represents an edge of the connectivity graph. An edge connecting two vertices in the edge conflict graph indicates the two underlying radio links interfere with each other. For the case when two edges of the connectivity graph share a common vertex, we declare that there is no conflict between the two edges. This is because, in practice, the two transmitting APs in this case will content with the shared AP using a medium access control protocol to avoid collision (see, e.g., [7]). In the example illustrated in Fig. II, AP C lies in the union area of the two radiated patterns will be interfered with by the communication, and the edge AB conflicts with all edges CD where $D \neq A$ and $D \neq B$.

Although we make several simplifying assumptions on the radio and antenna models, the algorithm to be described below only considers the connectivity graph and the edge conflict graph. It is possible to derive these two graphs from more sophisticated radio and antenna models, or from site-measurement database.

For all the simulations performed in Section III and Section V, θ is set to be $\frac{\pi}{12}$.

III. Constructing Two Wireless Meshes in the Same Region

In this section we describe an algorithm, called beam-crossing grids, for constructing two wireless meshes, from a given a set of APs randomly placed in a 2-dimensional space. The algorithm can be extended to constructing three or more collocated wireless networks as discussed in Section VI.

The goal of the algorithm is to produce two wireless meshes, as illustrated by Fig. 7, which cover as large a common area as possible, using links which are sufficiently apart to prevent possible radio interference.

We divide the entire area into unit squares, and in each square region we will find link pairs that can be used simultaneously without causing interference. We prefer beam-crossing link pairs, as illustrated in Fig. 2. Beam-crossing links will reduce the chance one wireless mesh may block the connectivity of another wireless mesh.

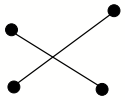

Fig. 2. Example of a beam-crossing link pair

The algorithm takes two inputs, a *connectivity graph* and an *edge conflict graph* (see Section II for the definitions of the two graphs). The output of the algorithm are two *subgraphs* of the connectivity graph, where any two edges in any subgraphs are interference-free according to the conflict graph. These two subgraphs represent the two wireless meshes that we are want to construct.

The algorithm assumes the location of each AP is known. The location information is used for two purposes: first, it is used to compute which square region an AP belongs to, and second, it is used to find proper link pairs in each square region, preferably beam-crossing ones. (The location information will not be necessary if the above two purposes can be achieved by other means.)

The algorithm consists of the following six steps:

1) Divide the area into unit squares, where the unit distance is the radio range of the APs use for inter-AP communication.
2) For each square we find a pair of links, if possible, that has no conflict. Both APs of a link must lie in the same square and the link pair can not share an AP. When there are multiple link pairs satisfying these conditions, the algorithm chooses the one with the largest weight, where weight is the sum of the length of the two links minus the sum of the minimal distances from APs to corners as shown in Fig. 3. Note that this selection criterion favors link pairs whose beams cross.
3) For each square for which a link pair has been found, recompute or eliminate link pairs if necessary so that

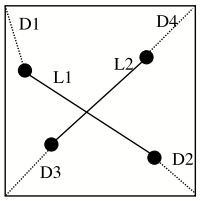

Fig. 3. An illustration of the weight computation in step 2. The weight of a link pair is $L1 + L2 - (D1 + D2 + D3 + D4)$. The use of each corner is exclusive, meaning that a corner can not be used twice. Note that this weight metric favors link pairs whose beams cross.

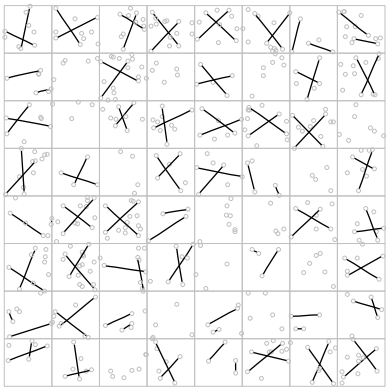

Fig. 4. The resulting link pairs after step 3. The link pairs are shown in black lines. Note that most of them are the beam-crossing type. It is possible that in some squares, no link pair can be found.

there will be no conflicts between links in adjacent squares. The algorithm achieves this by eliminating the link pair that has the most conflicts, and replaces it with a new link pair. Specifically, for each link pair, the algorithm computes its *penalty number* which is the number of other link pairs that it has conflicts with. The algorithm then eliminates the link pair with the largest penalty number and tries to find a replacing link pair in the same square. This time, the link pairs of adjacent squares are used as constraints. That is, the new link pair cannot have conflicts with the link pairs of the adjacent squares. No link pair will be chosen for the square if the no link pair satisfies the constraints. This process repeats until all link pairs' penalty numbers are zero. Fig. 4 is an example after this step.

4) For each square for which a link pair has been found, new links are added to connect them to link pairs in adjacent squares. (For computational simplicity, we include diagonally adjacent squares in our simulation). The two APs of an added link must be APs of previously chosen link pairs, and must belong to different squares. The algorithm looks at two adjacent squares at a time. Among the several ways of connecting two link pairs, we use the one that gives the largest number of links that are conflict-free. (This maximizes the number of in- and out-degrees of a square, thus helping find multiple networks.) The added links must also satisfy the constraint that there be no conflicts with nearby link pairs or other added links. This means that all the links, including the initial link pairs and the added links, are all conflict-free and can be used simultaneously. An example of the output of this step is shown in Fig. 5.

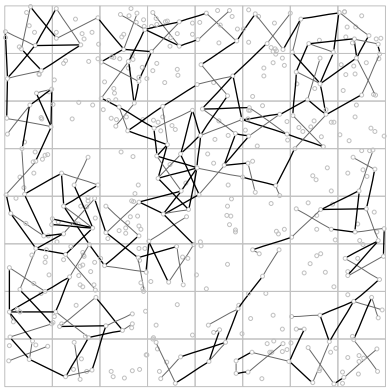

Fig. 5. The resulting link pairs after step 4. Link pairs from step 3 are shown in grey lines, and the added links of step 4 are shown in black lines.

5) Find connected components or subgraphs among links that are selected in step 3 and step 4. A beam-crossing link (any grey edge shown in Fig. 5) is randomly selected as the seed. The subgraph now contains the single seed edge. Starting from this seed edge, other connected edges are added to this subgraph through breadth-first search. A subgraph can only contain at most one beam-crossing link for each square. If the subgraph can not grow any more, the process starts with another randomly selected link and forms another subgraph. This process is repeated until every grey edge is included in a subgraph.
   This entire step is executed multiple times (we can bound the number of execution times for this step in the next section). Among these multiple executions, we will use the result of the execution that produces the largest combined size between the largest and second-largest subgraph.
   Fig. 6 shows an example of the resulting computed subgraphs. Only the largest four subgraphs are shown in this figure; the smaller subgraphs are not shown.
6) The algorithm now merges these subgraphs into two subgraphs. For each subgraph that is not the largest two subgraphs, we will perform a breadth-first search for a path that will connect the subgraph to one of the two subgraphs. The subgraph will select the connection path that will maximize the common area occupied by the resulting two subgraphs. An example of the two subgraphs resulting from step 6 is shown in Fig. 7.

IV. An Upper Bound on the Running Time of the Beam-Crossing Grids Algorithm

In this section we provide an upper bound on the running time of the beam-crossing grids algorithm presented in Section III. We will see that the upper bound can be expressed in terms of the size of the largest connected component, and

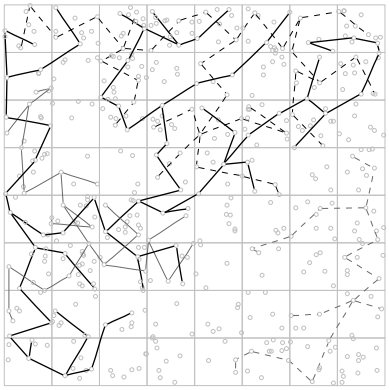

Fig. 6. The resulting subgraphs after step 5. There are four subgraphs shown in the figure. The largest one is depicted by the black lines, the second largest one by the dashed black lines, the third largest one by the grey lines, and the smallest one by the dashed grey lines.

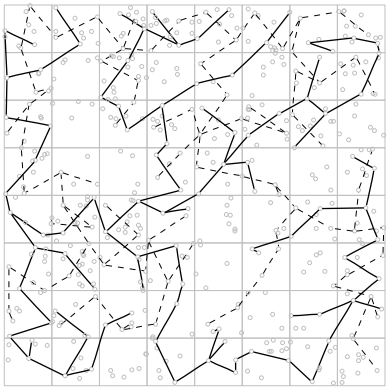

Fig. 7. The resulting two subgraphs after step 6. The first one is represented in black lines and has 99 nodes, while the second one is represented in dashed black lines and has 89 nodes.

then, in terms of the density of APs. We focus only on steps 5 and 6, since they are the dominating steps in running time.

Consider the graph resulting from steps 3 and 4. Let α be the ratio of the number of nodes in the largest connected component over that in the entire graph, and E the number of edges in the entire graph.

For any execution of step 5 of the algorithm, if the randomly selected seed happens to belong to the largest connected component, then through breadth-first search the execution will find this largest connected component. Thus the number of execution times for step 5 is bounded above by E minus the number of edges in the largest connected component.

Similarly, we can provide an upper bound on the number of times that step 6 of the algorithm will need to perform its breadth-first search operations. The number is bounded above by E minus the number of edges in the largest two connected components, which in turn, is number of edges in the largest connected component.

We can give tighter upper bounds for steps 5 and 6 in the following manner.

Theorem 1 1: Let T be the expected running time of step 5 or 6 of the algorithm of Section III. Then $T = O(\frac{\alpha^3-\alpha+1}{\alpha^2} \times E)$.

Proof: Let C be the expected running time of step 5 or 6 when it completes in 1 run and, for i = 1, 2, ..., let $A(i)$ be the expected running time of the algorithm when it completes in i runs and $B(i)$ be the expected running time of the algorithm when it completes in more than i runs. Then

$$\begin{aligned} C &= O(E) \\ A(i) &\le iC \\ T &= \frac{\alpha E}{E} A(1) + (1 - \frac{\alpha E}{E}) B(1) \\ &= \alpha A(1) + (1-\alpha) B(1) \end{aligned}$$

Note that

$$\begin{aligned} B(1) &= \frac{\alpha E}{E-1} A(2) + (1 - \frac{\alpha E}{E-1}) B(2) \\ &\le A(2) + (1 - \frac{\alpha E}{E}) B(2) \\ &\le 2C + (1-\alpha) B(2) \end{aligned}$$

Similarly, for i = 2, 3, ..., we have

$$\begin{aligned} B(i) &= \frac{\alpha E}{E-i} A(i+1) + (1 - \frac{\alpha E}{E-i}) B(i+1) \\ &\le A(i+1) + (1 - \frac{\alpha E}{E}) B(i+1) \\ &\le (i+1)C + (1-\alpha) B(i+1) \end{aligned}$$

Thus

$$\begin{aligned} T &\le \alpha C + (1-\alpha) \frac{\alpha E}{E-1} A(2) + (1 - \frac{\alpha E}{E-1}) B(2) \\ &\le \alpha C + (1-\alpha) 2C + (1-\alpha)^2 3C + (1-\alpha)^3 4C + \cdots \end{aligned}$$

By computing a closed-form expression for the right-hand side of the above inequality, we have

$$T \le \frac{\alpha^3 - \alpha + 1}{\alpha^2} C$$

■

Fig. 8 and Fig. 9 are simulation results which relate the size of the largest connected component to the average density of APs in a square. The former shows the largest component size as a function of probability that adjacent squares can be connected, and the latter shows the probability that adjacent squares can be connected as a function of the average density of APs in a square. Combing these results with the earlier result in this section, we see that the expected running time of the beam-crossing grids algorithm of Section III is a function of the average density of APs in a square. In particular, we see that the former will decrease as the latter increases.

V. Evaluation and Simulation Results

We evaluate the performance of the proposed algorithm using simulations, and report the simulation results in this section. We first introduce a coverage ratio that measures how well the area is covered by the meshes. Finally, we describe the simulation setup and report the results in terms of coverage ratio and number of APs used.

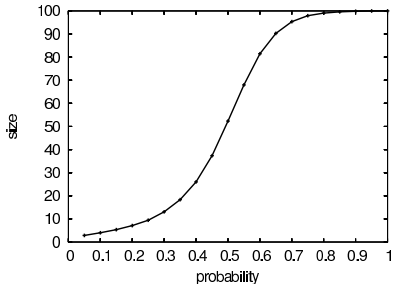

Fig. 8. Largest component size as a function of probability that adjacent squares can be connected

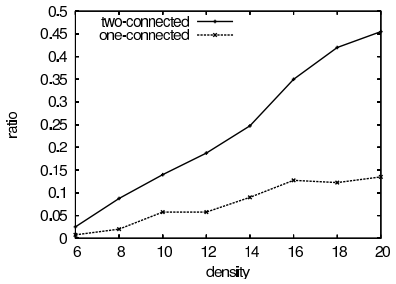

Fig. 9. Probability that two adjacent squares are connected as a function of density. Two-connected probability represents the probability that the two links can be found to connect the pairs of beam-crossing links; one-connected probability represents the probability that only one link can be found to connect the pairs of beam-cross links.

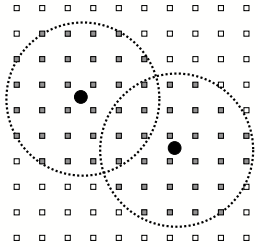

Fig. 10. An example of coverage ratio computation. There are a total of 100 uniformly placed sampling locations in the network and there are two APs, each covering a part of the network. The AP on the left covers 32 sampling locations, while the AP on the right covers 29 sampling locations. They jointly cover 55 sampling locations and their coverage intersects at 6 sampling locations. The union coverage ratio of the two APs combined is then 0.55 and the intersection coverage ratio is 0.06.

density	no. of APs	size of network 1	size of network 2
4	256	31.39	39.78
6	384	69.96	70.84
8	512	95.36	91.96
10	640	108.64	104.76
12	768	117.51	111.83
14	896	122.61	116.71
16	1024	124.72	117.83

TABLE I
SIZE OF EACH NETWORK

We assume each AP covers a circular area centered at the AP and the radius is its *access radio range*. The area a network covers is the coverage area of all its APs. The coverage ratio measures approximately the fraction of area covered by a network, by sampling at uniformly placed sampling locations. More specifically, all sampling locations lying in the circular area of any AP belonging to a particular network are said to be covered by the network. The coverage ratio is the number of covered sampling locations to the total number of sampling locations. An example of the coverage ratio computation is shown in Fig. 10. The union coverage ratio and the intersection coverage ratio of two networks can also be computed. For example, if each AP in Fig. 10 is a network, then the union coverage ratio is 0.55 and the intersection coverage ratio is 0.06.

For all simulations, the APs are randomly placed in an area of size 8 by 8. A *density* is used to control the total number of APs to be deployed in the networking area. We consider the following densities: 4, 6, 8, 10, 12, 14, 16. This translates into a total of 256, 384, 512, 640, 768, 896 and 1024 APs, respectively. For computing coverage ratio, we place sampling locations 0.05 apart, so there are a total of 25600 sampling locations in the networking area. Because this access radio range affects coverage ratio, we consider the following access radio ranges: 1.0, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4. For step 5 of the algorithm, 10 random trials are performed and the best results are chosen.

Fig. 11(a) to Fig.11(g) show the coverage ratios for various configurations. Each figure has a fixed access radio range, and shows the coverage ratios of all APs, the union and intersection coverage ratio of two networks, and the coverage ratios of individual networks. The coverage ratios reported are the averages of 100 simulations, each with randomly generated AP locations. Table I shows the average size, in number of nodes, of each network.

The results of coverage ratios show that, when the density is 8 or higher than and the access radio range is 0.6 or larger, each network on average covers more than 80% of the area. The union coverage of both networks in this case covers more than 95% of the area. This demonstrates that the proposed algorithm is effective in finding two networks without losing coverage.

Because the algorithm connects single links found in small regions, the resulting individual network uses no more than two APs per square. On the other hand, to achieve good performance the density must be greater than or equal to 8. This means that in order to achieve good performance, twice as many APs must be supplied to the algorithm than those would be used.

VI. DISCUSSION

The described method and algorithm can be used to organize APs in order to increase the throughput of broadband Internet services [1]. For example, in a rural area where wireline broadband has not yet be deployed, household APs can be connected using the method to increase the throughput of the network.

The method and algorithm can also be used for planning the deployment of APs. Instead of choosing from the set

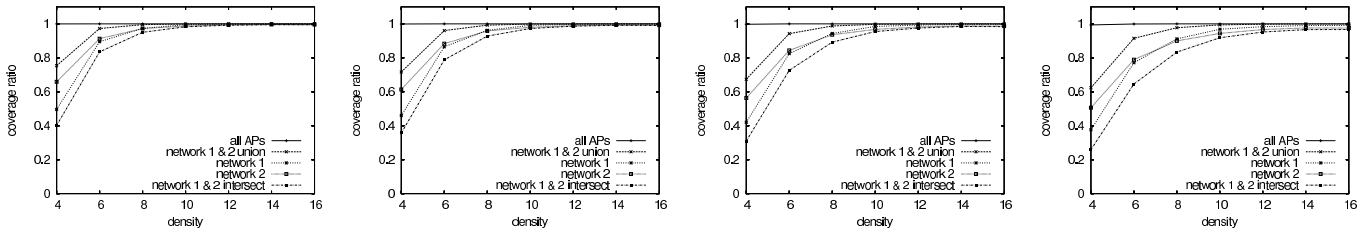

(a) Access Radio Range = 1.0 (b) Access Radio Range = 0.9 (c) Access Radio Range = 0.8 (d) Access Radio Range = 0.7

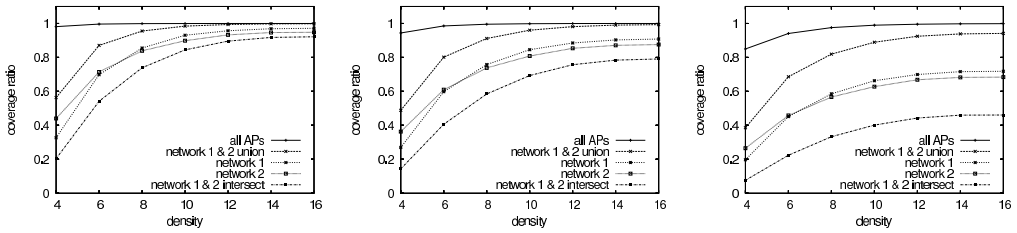

(e) Access Radio Range = 0.6 (f) Access Radio Range = 0.5 (g) Access Radio Range = 0.4

Fig. 11. Density vs. coverage ratio performance with different access radio ranges. The coverage ratios reported are average coverage ratios of 100 randomly generated networks.

of already-deployed APs, we consider all potential AP deployment locations. The output of the algorithm is a list of locations where APs should be placed and how they form networks for higher throughput. In other words, the method and algorithm can be used to select AP deployment locations. If APs can be placed at arbitrary locations, then there are other strategies that can be used to deploy these APs. For example, in our prior work [10], we consider the problem of collocating identical square mesh wireless networks.

The algorithm generates two wireless networks that each has the topology of a tree. Tree topology has the property that it uses minimal number of edges to connect all vertices. Because each additional link used in the network imposes additional interference constraint to itself and the other network, it is nature to design the algorithm to construct the network with tree topology. We could add shortcut links that connect branches of the network to reduce path length between any pair of APs, but we don't consider this optimization technique in this paper.

The major disadvantage of tree topology is its lack of redundancy— any link breaks can result in the partitioning of the network. It may not be possible to reconnect the network by using other APs around the faulty area, so new networks should be recomputed using non-faulty links. When links break due to faulty APs, the algorithm excludes the faulty APs as well as all links associated with them, and reconstructs the networks. Fig. 12 demonstrates, by simulations, that the algorithm can still find large-coverage collocated networks when some APs used by the networks are removed. For those simulations, the networking area is a 8 by 8 area and initially there are a total of 640 APs randomly placed in the network

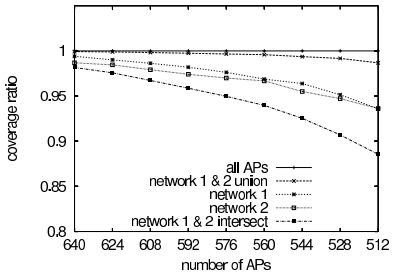

Fig. 12. Changes in coverage ratios when more APs used by the computed networks are removed. Initially, there are a total of 640 APs in the area (density of 10). At each iteration, 8 randomly selected APs are removed from each network and two new networks are computed from the remaining APs. The coverage ratios reported are the average of 100 randomly generated networks.

(density of 10). At each iteration, 8 randomly selected APs are removed from each computed network and the algorithm computes two new networks using the remaining APs. The process repeats until a total of 128 APs (64 from each network) are removed from the network. The coverage ratios reported are average coverage ratios from 100 randomly generated networks and are computed with access radio range set to be 1.0. The density becomes 8 after all of the 128 APs are removed, and the coverage ratios are in general lower than the coverage ratios reported in Fig. 11(a) with the same density and access radio range. This is because only the APs that were used in the networks are removed and these APs are better positioned to form networks than the others. When the density of remaining APs are high, the algorithm can still perform well even though a large number of better positioned APs are removed.

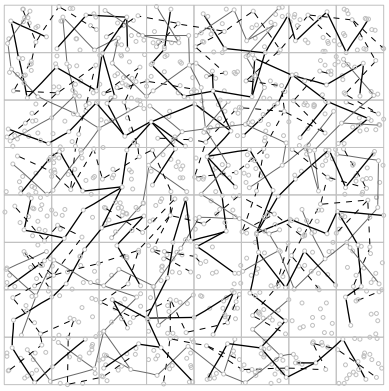

Fig. 13. An example of three collocated networks. The largest one is shown in black solid lines, the second largest one is shown in black dashed lines and the smallest one is shown in grey solid lines.

Although the algorithm described finds two collocated networks, it can be generalized to any number. For example, to find three networks, instead of finding a link pair for each square, the algorithm should seek link triplets. The largest three subgraphs are then used to merge with other smaller subgraphs. Fig. 13 shows an example of 3 networks constructed by this extended algorithm. The θ is reduced to $\frac{\pi}{18}$ when performing the simulation.

We do not attempt to minimize the number of APs used. Minimizing the number of APs used while covering at least a certain portion of the area requires the knowledge of AP coverage patterns. Our algorithm does not assume specific coverage patterns, and it only assumes (1) the union coverage does not decrease as more APs are used, and (2) the closer the two APs are the smaller their union coverage is. On the other hand, due to its design, each network never uses more than two APs in a square. This provides an upper bound on the total number of APs for each network.

VII. Conclusion

In this paper, we describe an algorithm, called *beam-crossing grids*, for constructing two collocated wireless mesh networks from a collection of APs in a region. The two wireless meshes can operate independently of each other and wireless terminals can use the networks for fault-tolerance or load-balancing access.

The algorithm is evaluated using simulations. The coverage area of the individual wireless networks increases as the density of APs increases. When the density is 8 or higher and access radio range is 0.6 or higher, our simulation results show that the algorithm effectively constructs two collocated networks that each covers more than 80% of the area. The coverage ratio increases as density or access radio range increases. Because the algorithm divides the area into small regions and for each network it uses at most two APs in a region, the total number of APs used by each network is bounded above by two times the number of regions.

We also discuss how to extend the algorithm to construct more than two networks, and show an example of constructing three networks.

We plan to extend this work in several directions. First, we will take the AP coverage patterns into consideration to allow each computed wireless network to use fewer APs without sacrificing their coverage areas. Second, we will extend the method to 3-dimensional space when the coverage patterns are well-defined. Finally, we plan to investigate the application of this method for cases where APs can adjust their transmitting power levels. This can minimize interference caused by a transmission and allow constructing an expanded number of collocated networks.

Acknowledgments

This research was supported in part by NSF grant #ACI-0330244, and in part by a research grant from Microsoft Research.

References

[1] D. Beyer, "Wireless mesh networks for residential broadband," in *National Wireless Engineering Conference*, San Diego, 2002.

[2] Nortel Networks, "Nortel Networks announces plans to trial breakthrough wireless LAN architecture with MIT and BT," 2003, press release.

[3] V. D. Park and M. S. Corson, "A highly adaptive distributed routing algorithm for mobile wireless networks," in *Proc. of IEEE INFOCOM*, 1997, pp. 1405–1413.

[4] A. Nasipuri, R. Castaneda, and S. R. Das, "Performance of multipath routing for on-demand protocols in ad hoc networks," *ACM/Kluwer Mobile Networks and Applications (MONET) Journal*, vol. 6, no. 4, pp. 339–349, 2001.

[5] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly-resilient, energy-efficient multipath routing in wirelese sensor networks," *Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 11–25, 2001.

[6] J. Kajiya, "Commodity software steerable antennas for mesh networks," in *Mesh Networking Summit*, 2004.

[7] Y.-B. Ko, V. Shankarkumar, and N. H. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *Proc. of IEEE INFOCOM*, March 2000.

[8] M. Takai, J. Martin, and R. Bagrodia, "Directional virtual carrier sensing for directional antennas in mobile ad hoc networks," in *Proc. ACM Intern. Sym. on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, June 2002.

[9] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "Energy limited wireless networking with directional antennas: the case of session-based multicasting," in *Proc. of IEEE INFOCOM*, 2002.

[10] P.-H. Hsiao and H. T. Kung, "Layout design for multiple collocated wireless mesh networks," in *Proc. of IEEE VTC 2004 Fall*, September 2004.

Adjacent Channel Interference in Dual-radio 802.11a Nodes and Its Impact on Multi-hop Networking

Chen-Mou Cheng Pai-Hsiang Hsiao H. T. Kung Dario Vlah
{doug, shawn, htk, dario}@eecs.harvard.edu
Department of Electrical Engineering and Computer Science
Harvard University
Cambridge, MA 02138

ABSTRACT

We evaluate the performance impact of adjacent channel interference (ACI) in multi-hop wireless networks based on dual-radio 802.11a nodes. Although these nodes use chipsets that satisfy the transmit-mask requirements set by the IEEE 802.11 standard, the multi-hop performance is still significantly affected by ACI. That is, a node's transmitter can interfere with its own receiver on a different channel; as a result, multi-hop throughput is severely degraded. This degradation is especially pronounced for 802.11a. We use a spectrum analyzer with a signal combiner to quantify ACI under various conditions and propose solutions to mitigate the effect of such interference on multi-hop forwarding. Field experiments with multi-hop relay have validated these findings as well as the effectiveness of our solutions.

I. INTRODUCTION

Commercial off-the-shelf (COTS) wireless equipment, such as the IEEE 802.11 wireless LAN ("WiFi"), has attracted a good deal of attention from both academia and industry due to its relatively low cost and high performance. Substantial research and development efforts have been spent on designing and deploying WiFi-based wireless mesh networks as well as community networks [1], [3], [11]. Over the years, many researchers have independently shown that the performance of such multi-hop wireless networks degrades rapidly as the number of hops increases; they attribute this to many reasons including inefficient medium access control, radio interference, wireless link errors resulting from changing channel conditions and multipath effects, frequent route changes, and improper TCP's interaction with lower-layer protocols [1], [9], [11]. New generation of technologies such as WiMAX [16] may be able to alleviate some of these problems, e.g., by using the time-division multiplexing (TDM) technique. However, at the time of writing, it still remains unknown how wireless multi-hop networks in the real world can effectively employ these new technologies to achieve high performance that scales with number of hops.

In the literature, we have found two approaches in enhancing the performance of multi-hop wireless networks: (1) use of directional antennas [12], [17], and (2) use of multiple omni-directional radios at each node [2]. The former enables spatial reuse, allowing a node to communicate with more than one neighbor at the same time using separate beams over the same frequency band, while the latter allows use of separate channels over multiple radios. This paper concerns an instance of the latter approach, where each node employs two 802.11a radios with omni-directional dipole antennas. This approach is attractive because it is relatively inexpensive these days to incorporate two COTS radios in each node, and the deployment of these nodes requires little antenna-related engineering. We will focus on the newer, OFDM-based 802.11a systems since they are capable of delivering higher performance with approximately the same bandwidth.

In this paper, we quantify the effect of *adjacent channel interference (ACI)* [14] on the performance of dual-radio, multi-hop 802.11s networks. Specifically, based on lab measurements and field experiments, we report that a node's tranmission can significantly interfere with its own reception, even though the transmit and receive radios use two separate channels. We demonstrate that this interference can lead to two-fold or worse performance degradation in data transfer (Section II). The performance degradation is particularly significant for an important scenario where the signal-to-noise ratio (SNR) of the receiving channel is at the lower end of the low-loss region (Section III). We show that the ACI problem is present with chipsets from various vendors, even though according to our spectrum-analyzer measurements (Section IV), they

This material is based on research sponsored by Air Force Research Laboratory under agreement numbers FA8750-05-1-0035 and FA8750-06-2-0154, and by the National Science Foundation under grant number #ACI-0330244. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, the National Science Foundation, or the U.S. Government.

all conform to the IEEE 802.11a spectrum-mask requirements [8].

Toward the end of this paper, we list several possible solutions to mitigate the effects of ACI. These include increasing the distance between the two dipole antennas on a node, collinear placement of the two antennas, and larger transmit and receive channel separation (Section V). Finally, we provide analysis on the expected throughput of multi-hop data transfer in the presence of ACI and validate it with field experiments.

II. Testbed Setup and Initial Multi-Hop Experience

As a proof of concept, we assembled a testbed platform of nodes based on 400MHz AMD Geode single-board computers made by Thecus Inc. On each node we installed two Wistron Neweb CM9 mini-PCI network adapters (based on the Atheros AR5213A 802.11a/b/g chipset). We used the `madwifi` Linux driver for the network adapters. We used two different types of antennas: 7dBi omni-directional dual-band antennas and 2dBi home-brew dipole antennas.

We conduct several preliminary experiments using a linear topology consisting of three nodes on a line (A—B—C) in a yard outside our office building in Cambridge, MA. We used 802.11a radios in our experiments; luckily, there were no other 802.11a signals on that part of the campus. We measured throughput of three different configurations. The first configuration, in which A transmits to B, serves as our baseline case. The second configuration, with A transmitting to B and B to C, is the traditional two-hop, single-channel configuration, in which all nodes use the same channel under the CSMA arbitration. The last configuration is a two-hop, two-channel configuration, in which A transmits to B on channel X, while B relays to C on channel Y.

The measurement result reveals an inherent problem with multi-radio nodes. When the link quality is good, all three configurations work fine: the throughput from A to C of the two-hop, single-channel configuration is about one half of that from A to B, whereas that of the two-hop, two-channel configuration is almost identical to that from A to B. However, when the link quality becomes marginal, the throughput achieved by the two-hop, two-channel configuration drops drastically. In the worst cases, it can drop below that of the two-hop, one-channel configuration when the link quality becomes really poor.

One may solve this problem by always operating the wireless links at high enough SNR. However, such a solution will also significantly reduce the transmission range of each hop, not to mention that marginal links are the norm in many real-world networks [1], [3]. Therefore, we investigate other solutions that address the problem directly; we begin with more detailed measurement experiments in order to determine why the performance became worse than expected under marginal link quality.

III. Initial Investigation of Causes Leading to In-Lab Experiments

Our initial explanation for the performance degradation was due to transmission power on one channel leaking into adjacent channels on the relay node. Although the IEEE 802.11 standards specify a minimum transmit mask to protect the adjacent channels from being interfered with, it may be the case that the particular hardware we use does not strictly abide by the standards.

To see if this is the case, we hook up one of our testbed nodes to a spectrum analyzer. Figure 1 shows the power spectral density (PSD) obtained when we have the testbed continuously broadcasting packets on channel 52 (centered at 5260MHz) with various `txpower` settings in the `madwifi` driver.

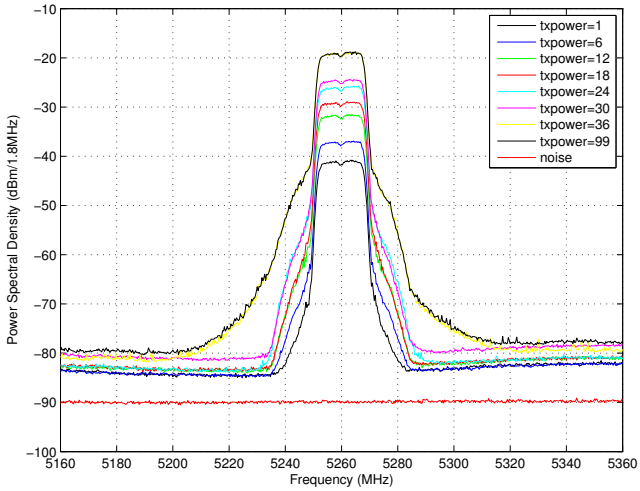

Fig. 1. **The power spectral density of an 802.11a broadcast signal on channel 52 (5260MHz).**

From the PSD, we can obtain the transmission power leakage for each channel, as summarized in Table I.

Ch.	36	40	44	48	52	56	60	64
dBm	-59	-59	-52	-27	0	-26	-53	-57

TABLE I

The measured channel power of adjacent channles when an 802.11a radio continuously broadcasts on channel 52 (5260MHz).

The IEEE 802.11 standard specifies that the peak power at 11MHz apart should not exceed -20dB of the peak power (similarly, -30dB at 22MHz). We see clearly that the data of Figure 1 satisfies this requirement. We also

measured a few cards from other manufacturers (such as the Intel PRO/Wireless 2915ABG network adapters). While satisfying IEEE's transmit-mask requirement, they all exhibit similar ACI problems when used for the two-hop, two-channel configuration. This leads us to believe that this transmit-mask requirement is not sufficient for our application scenario, where the two co-located antennas are very close to each other. We decided to pursue this avenue of investigation further via a set of controlled experiments, described in the following section.

IV. In-Lab Signal-Combiner Measurement Results

We conduct a set of controlled, in-laboratory experiments that quantify the impact of interference from adjacent channels on UDP throughput for 802.11a radios. In these experiments, we connect the transmit and receive radio interfaces via low-loss cables through signal combiners. Besides permitting better control over desired signal-to-interference ratio (SIR), this allows us to mitigate randomness in radio propagation. The result captures the effect of the interference mechanism intrinsic to the particular 802.11a implementation at hand and hence can be used to explain phenomena observed in various scenarios using the same radio configurations.

We use the Atheros-based testbed described previously. Specifically, there are three physically separate nodes: Tx, Rx, and Int. The output ports of Tx and Int are connected to the input port of a signal combiner, with the combined signal fed into the input port of Rx. To emulate a broad range of SIRs which may be observed at Rx and to protect receiver circuitry, we insert appropriate attenuation in the cabling. We then measure the UDP throughput from Tx to Rx under various SIRs, which are achieved by adjusting the transmission power at Tx and Int, as well as by using different combinations of attenuators.

We use `netperf` to measure UDP throughput. Unfortunately, we could not directly measure SIR in this particular testbed. Instead, we measure the signal power and the interference power separately, from which we can obtain the SIR for a particular configuration.

The signal power is simple to obtain once we have received legitimate 802.11 packets at Rx: upon successfully reception of a packet, the driver reports the RSSI for this packet. This RSSI reading reflects the received signal strength in that packet and can be converted to signal power. Additionally, we perform a set of calibrations using a spectrum analyzer, and the result shows that the received signal power derived from this method is accurate across a wide range of configurations. Other researchers have also reported similar findings on COTS 802.11 radios [10].

We use a similar method to measure the interference power. During the experiments, we measure the received signal power at Rx due to Int alone by temporarily putting Rx on the same channel as Int and having Int sending out broadcast packets. Since these packets are also legitimate 802.11 packets, Rx will report the RSSI readings, which can be converted to received signal power. We then use the PSD measurements as described in the previous section to extrapolate the perceived interference power on the target channel on which Tx and Rx are having the `netperf` sessions. We have also verified with a spectrum analyzer that the SIR so obtained is accurate.

Figure 2 shows some sample results. Here, the Tx-Rx pair always uses channel 52 (5260MHz), whereas Int may use channel 44 (5220MHz) or channel 48 (5240MHz). Tx-Rx can also use any appropriate transmission power, but Int can only use the lowest (0dBm) or the highest (18dBm) power levels.

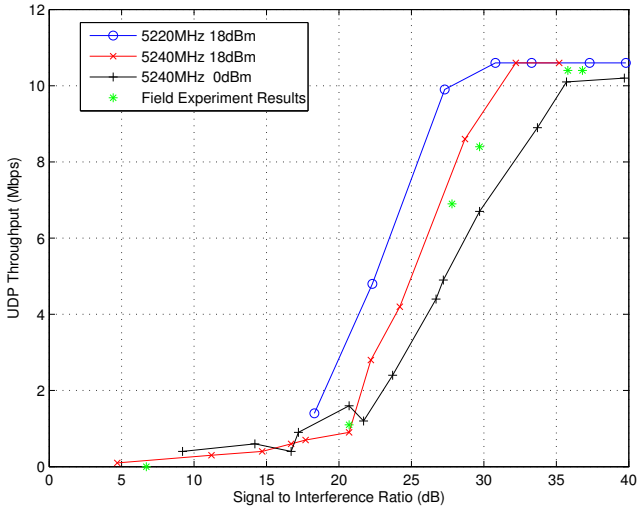

Fig. 2. **Effects of SIR on UDP throughput for a 12Mbps 802.11a link.**

In general, the achieved UDP throughput depends not only on the perceived SIR at Rx, but also on the specific configuration (including such factors as nominal Tx-Rx link speed and absolute received signal strength). In some configurations, UDP throughput can also depend on the particular transmission power levels used at Int. Nevertheless, the typical variation due to this different configurations is not very signifcant, so the SIR still serves as a good indicator for predicting the resulting UDP throughput in many cases. As a rule of thumb, a 20dB SIR is needed for a barely functioning Tx-Rx link, whereas a 30dB SIR will be sufficient for a link to be close to fully functioning. We also plot a few data points obtained from outdoor field experiments; we see clearly that these points well fall within the predicted range.

V. Proposed Solutions to Adjacent Channel Interference and Field Experiment Results

When considering possible ways to combat adjacent-channel interference, we focus on the case where forwarding nodes transmit and receive simultaneously; otherwise, the problem is trivially solved by avoiding it. A scheme like time-division multiplexing (TDM) might avoid ACI, but its own unique issues may not be acceptable in a system where ACI occurs. For example, nodes with inaccurate clocks might need a large latency in order to maintain efficient channel utilization—likely inappropriate for a delay-sensitive application. The interested readers are referred to another paper of the authors' for a more detailed discussion and comparison for other possibilities [6].

Any solution for ACI must reduce the amount of unwanted signal at a co-located receiver. We suggest the following three ways to achieve this:

1) **Antenna Engineering.** Antennas can be adjusted to reduce the amount of mutually radiated signal by exploiting different ways that signals attenuate.
   - **Distance.** We can maximize the distance between the antennas to increase the interference signal's path loss.
   - **Orientation.** Typically, antennas have at least one orientation with strong negative gain, which is called the "null" region of the antenna. We can orient the antennas so that they lie in each other's null regions.
   - **Shielding.** Shielding between antennas, such as a thin metal plate, can reflect some of the interfering signal. In our tests, simple foil can provide 10dB of attenuation. However, the downside is that the antenna radiation pattern becomes strongly directional.
   - **Cabling.** Ideally, all transmission lines connecting the antennas to the wireless card would be perfectly matched and lossless. However, it turns out that cable imperfections can not only attenuate the signal, but also help radiate it; in our tests, moving a thin one-meter coaxial cable to several different positions resulted in as much as 10 dB difference in emitted signal, depending on whether the signal radiated from the cable constructively or destructively combines with the antenna signal. Unfortunately, solving this problem may require complicated RF engineering, which is out of the scope of this paper. However, in our experiment setups, we address the problem by securely attaching cables in a consistent manner in order to reduce uncertainty due to signal variation.
2) **Filtering.** There exist commercially available band-pass filters which attenuate the out-of-channel signals by more than those built into the wireless cards. For example, the relatively inexpensive, indoor 2.4Ghz filter from Hyperlink Technologies [7] can attenuate unwanted signal by roughly 60dB, compared to the 50dB attenuation we measured in Section IV. The problems with this approach are first, the additional filter may increase the cost of the system, thus defeating the purpose of using COTS equipment. Second, the size of the actual filter may be unacceptable for some applications; for example, the Hyperlink Technologies filter is 8" long, and weighs 0.5 pounds, which is quite cumbersome for applications like unmanned aerial vehicle (UAV) wireless networking [4], [5].
3) **Power Control.** One could reduce the transmit power of the interfering radio, thus reducing the interfering signal by a proportional amount. This approach might work when the quality of the incident link is sufficiently above the sensitivity threshold. However, in multi-hop networks where ACI appears in the first place, there are likely to be few such links by design.

Of the above approaches, we focused on antenna engineering. In particular, we measured the effect of antenna separation on ACI by performing a set of field experiments, detailed in the following section.

A. Field Experiment

We performed a number of two-hop throughput measurements in an open field near the Harvard Stadium in Alston, MA. We used three nodes, labeled A, B, and C, with the intention of sending traffic from A to C via B. The length of each hop was 66 yards. A and C were placed on plastic stands at a height of 28", while node B was mounted on a wooden post 14' high (Figure 3 shows node B in position.).

We examined two link configurations, where link AB

Fig. 3. **Node B mounted on a wood tower. The antennas are mounted on horizontal, 3ft long masts**

was either *good* or *bad*, achieving around 99% or 50% of the possible throughput, respectively. The observed signal strenghts for these two link configurations were -65dBm and -72dBm. We created these conditions by varying the position of node A. The case where the first hop is bad is interesting because that is when node B's receiver suffers most from its transmitter's interference.

For each link configuration, we set the distance between antennas to *minimum* (11") and *maximum* (75"). For the minimum distance, we coiled the cables and taped them to the opposing sides of the node enclosure. For the maximum distance, we attached the cables to posts extending from the opposite sides of the enclosure, visible in Figure 3.

We ran the first hop on 802.11a channel 48. We used two channels for the second hop: channel 52 and 56, which were 20MHz and 40MHz away, respectively. This means there are 8 total combinations of first hop quality, antenna distance, and channel separation settings.

We measured the ACI power on node B as follows. We temporarily switched the second hop to channel 48—same as the first hop. We had node B broadcast a load via its intended transmitting radio, and collect RSSI figures from its receiving radio. Then, we calculated the ACI power by subtracting relative leakage values in Table I. The resulting numbers are shown in Table II.

Antenna distance	Measured RSSI on the same channel	Computed interference 20 MHz away	Computed interference 40 MHz away
Min (11")	-36 dBm	-64 dBm	-89 dBm
Max (75")	-49 dBm	-77 dBm	-102 dBm

TABLE II
Loopback RSSI measured at node B and the calculated adjacent-channel interference powers.

Main Measurement Results. We measured the two-hop throughput under eight combinations of the first hop quality, antenna distance and channel separation settings. We present the outcomes in Table III for 20MHz channel separation and Table IV for 40MHz. Each cell corresponds to one setting of antenna distance and first hop quality. Within each cell, the first value is the throughput of the first hop alone, and the second value is the two-hop throughput. The throughput of the second hop is not shown, because as we discussed earlier, its quality is fixed high enough to maintain a full throughput of 10.4 Mbps.

B. Discussion and Analysis

It is evident that increasing channel separation helps. As we can see from tables III and IV, the throughput increases from 56–81% to 73–94% of the bottleneck link. Furthermore, performance improves whenever the antenna distance increases, supporting the use of antenna engineering to mitigate ACI.

Antenna distance	Quality of hop AB			
	Good		Bad	
Min	9.0	5.0 (56%)	3.0	2.3 (77%)
Max	10.4	6.4 (62%)	4.4	3.55 (81%)

TABLE III
Throughput in Mbps measured with two hops running on channels 20MHz apart.

Antenna distance	Quality of hop AB			
	Good		Bad	
Min	10.2	7.9 (77%)	5.1	3.7 (73%)
Max	10.3	9.4 (91%)	5.5	5.15 (94%)

TABLE IV
Throughput in Mbps measured with two hops running on channels 40MHz apart.

We develop a simple two-hop throughput analysis that takes into account the adjacent-channel interference created by the relay node. Let us label the three nodes A, B and C. We assume the first hop alone, A-B, has infinite offered load and packet delivery probability p_1. We assume the second hop delivers all packets. To account for ACI, we define a second delivery probability, p_2, for link A-B *while link B-C is transmitting.*

Our goal is to find the overall two-hop throughput achieved by this network given p_1, p_2, and nominal link capacity C. We shall denote the overall throughput by x. Since hop B-C is perfect, there are no retransmissions and the total load from B to C is x. Therefore, the approximate amount of load on hop A-B affected by ACI is also x. From these facts we can compute the throughput of hop A-B as

$$r_{AB} = xp_2 + (C - x)p_1$$

This throughput flows over hop B-C with negligible loss, so we have the equality $r_{AB} = x$ and

$$x = \frac{Cp_1}{1 + p_1 - p_2} \tag{1}$$

Note that p_2 includes the effects of p_1, so $p_2 \leq p_1$ and $x \leq Cp_1$.

We apply the analysis to our field experiments by computing the p_1 and p_2 probabilities for each of the 8 combinations of parameter settings. The nominal link capacity C is 10.4 Mbps. We calculate p_1 as $MeasuredThroughput/C$. We calculate p_2 indirectly in two steps. First, we compute the SINR for link A-B by subtracting the appropriate ACI figure in Table II from the link A-B RSSI given earlier (-65dBm for good, -72dBm for bad links). Second, we map this SINR to a corresponding throughput in Figure 2, and then compute $p_2 = p_1 \cdot CorrespondingThroughput/C$. The obtained

values are shown in Table V, along with the predicted throughputs x, calculated using Equation 1.

(Antenna distance, First hop quality, Channel separation)	p_1	p_2	Predicted Throughput (Mbps)	Measured Throughput (Mbps)
Min,Bad,20MHz	0.29	0.00	2.3	2.3
Max,Bad,20Mhz	0.42	0.00	3.1	3.6
Min,Good,20Mhz	0.87	0.00	4.8	5.0
Max,Good,20Mhz	1.00	0.05	5.3	6.4
Min,Bad,40MHz	0.49	0.09	3.6	3.7
Max,Bad,40MHz	0.53	0.53	5.5	5.1
Min,Good,40MHz	0.98	0.52	7.0	7.9
Max,Good,40MHz	0.99	0.99	10.3	9.4

TABLE V
Values for p_1 and p_2 and analytical predictions

VI. Conclusions

In this paper, we have quantified the "adjacent-channel interference" problem for dual-radio 802.11a nodes. We would like to point out that this ACI problem has been reported previously in the literature, e.g., [13]. However, previous works are concerned with 802.11b radio systems, whereas we put more emphasis on the new OFDM-based 802.11a systems that are capable of delivering higher throughput performance. In a companion paper of the authors', we compare the severeness of the ACI problem for OFDM vs. non-OFDM 802.11a/g radio systems and conclude that ACI has greater impact on the multi-hop throughput performance of OFDM-based systems because the minimum SNR required is much higher in OFDM than non-OFDM [6] for a certain packet error rate. Although this paper did not give ACI measurement data related to OFDM-based 802.11g systems, we did see similar ACI problems with 802.11g in our field experiments. Despite the fact that the Atheros team reports a much lower required minimum SNR for baseband processing in their simulation [15], we find that many practical implementations require a higher minimum SNR at the antenna connector (greater than 20dB experienced for OFDM vs. typically less than 10 dB for non-OFDM). We have independently verified the minimum SNR discrepancy via in-lab spectrum-analyzer measurements, in-lab equipment calibration, as well as two-hop field experiments.

Due to the severe ACI we have observed, we suspect that the use of multi-radio nodes in multi-hop networking was not part of the considerations when the 802.11 spectrum-mask requirements were standardized. As interest in 802.11 multi-hop applications grows, it is perhaps justifiable that follow-on 802.11 standards call for more stringent requirements on transmit spectrum masks or receive sensitivity. Before the next-generation 802.11 chipsets satisfying these follow-on requirements become available, we could use the techniques described in this paper, such as increasing channel separation and antenna distance, to mitigate such adjacent-channel interference.

References

[1] Aguayo, D., Bicket, J., Biswas, S., Judd, G., and Morris, R., "Link-level Measurements from an 802.11b Mesh Network," *SIGCOMM 2004*, Aug 2004

[2] Alicherry, M., Bhatia, R., and Li, L., "Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks," *ACM MobiCom 2005*, August 2005

[3] Bicket, J., Aguayo, D., Biswas, S., and Morris, R., "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *ACM MobiCom 2005*, Aug 2005

[4] Brown, T. X., Argrow, B., Dixon, C., Doshi, S., "Ad Hoc UAV Ground Network (AUGNet)," *AIAA 3rd "Unmanned Unlimited" Technical Conference,* Chicago, IL, 2004.

[5] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations," *ICCCN 2006*, Arlington, VA, October, 2006.

[6] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Parallel Use of Multiple Channels in Multi-Hop 802.11 Wireless Networks," *IEEE MILCOM 2006*, October 2006

[7] Hyperlink Technologies, Inc., "2.4 GHz 802.11b and 802.11g Compatible 4-Pole Ultra High Q WiFi Bandpass Filters," *http://www.hyperlinktech.com/web/bpf24-4xx.php,* 2005

[8] IEEE 802.11 Working Group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications," *IEEE 802.11 standard, including 802.11a and 802.11b extensions*, Sep 1999

[9] Jain, K., Padhye, J., Padmanabhan, V. N., and Qiu, L., "Impact of Interference on Multi-hop Wireless Network Performance," *ACM MobiCom 2003*, September 2003

[10] Judd, G., and Steenkiste, P., "A simple mechanism for capturing and replaying wireless channels," *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005

[11] Li, J., Blake, C., De Couto, D. S. J., Lee, H. I., and Morris, R., "Capacity of Ad Hoc Wireless Networks," *ACM MobiCom*, July 2001

[12] Ko, Y. B., Shankarkumar, V., and Vaidya, N., "Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks," *IEEE INFOCOM*, March 2000

[13] Robinson, J., Papagiannaki, K., Diot, C., Guo, X. and Krishnamurthy, L., "Experimenting with a Multi-Radio Mesh Networking Testbed," *1st workshop on Wireless Network Measurements (WiNMee 2005)*, Trento, Italy, April 2005

[14] Sklar, B., "Digital Communications." Prentice Hall, 1988, p. 193.

[15] Thomson, J., Baas, B., Cooper, E. M., Gilbert, J. M., Hsieh, G., Husted, P., Lokanathan, A., Kuskin, J. S., McCracken, D., McFarland, B., Meng, T. H., Nakahira, D., Ng, S., Ratterhalli, M., Smith, J. L., Subramanian, R., Thon, L., Wang, Y.-H., Yu, R., and Zhang, X., "An Integrated 802.11a Baseband and MAC Processor," Proc. ISSCC 2002, San Francisco, CA, February 2002

[16] The WiMAX Forum, `http://www.wimaxforum.org/`

[17] Yi, S., Pei, Y., and Kalyanaraman, S.,"On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas," *ACM MobiHoc 2003*, June 2003

Proc. 15th International Conference on Computer Communications and Networks (ICCCN 2006), October 9-10, 2006

Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations

Chen-Mou Cheng Pai-Hsiang Hsiao H. T. Kung Dario Vlah
{doug, shawn, htk, dario}@eecs.harvard.edu
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138

Abstract

We report measured performance of 802.11a wireless links from an unmanned aerial vehicle (UAV) to ground stations. In a set of field experiments, we record the received signal strength indicator (RSSI) and measure the raw link-layer throughput for various antenna orientations, communication distances and ground-station elevations. By comparing the performance of 32 simultaneous pairs of UAV and ground station configurations, we are able to conclude that, in order to achieve the highest throughput under a typical flyover UAV flight path, both the UAV and the ground station should use omni-directional dipole (as opposed to high-gain, narrow-beam) antennas positioned horizontally, with their respective antenna null pointing to a direction perpendicular to the UAV's flight path. In addition, a moderate amount of elevation of the ground stations can also improve performance significantly.

I. Introduction

We envision that in the future, low-flying UAVs could provide a cost-effective wireless networking means for ground devices. Such UAV-based wireless networking could have a number of advantages, including (1) that UAVs can provide on-demand, high-quality communication due to line-of-sight signal propagation; (2) that UAVs can be sensing and data-fusion nodes dynamically deployable in the region of interest; (3) that UAVs can tailor their flight paths to enhance the quality of wireless networking and communication; and, finally, (4) that UAVs can themselves carry and forward huge amounts of data, e.g., gigabytes of terrain images or databases. With these capabilities, a UAV-based network can, for example, provide high-speed transport of multimedia data (e.g., videos and images) for ground nodes and overcome environmental shadowing effects caused by blocking structures such as mountains and tall buildings.

The UAV-based networking approach has become especially attractive in the recent years due to the availability of low-cost, Commercial Off-The-Shelf (COTS) wireless equipment, such as IEEE 802.11 wireless LAN ("WiFi") [4]. For example, by integrating compact, 802.11 wireless equipment into a small 94-inch wing-span UAV, we can readily create a powerful networking node in the air [8].

In this paper, we address the issue of configuring 802.11 antennas in UAV-based networking. It is well-known that antenna types (e.g., omni or directional) as well as their positions and orientations can greatly affect the performance of wireless links [2], [3]. In addition, when a UAV communicates with ground nodes, we need to consider ground effects (such as interference from reflected signals, modeled by two-ray propagation [6]). Given the large number of complicated issues involved, it is essential that we conduct field experiments in order to understand the performance impact of various antenna configurations at the application level.

We have instrumented a UAV and several ground nodes with two types of 802.11a antennas in various orientations. Using this equipment, we have conducted a set of flight tests to measure their raw link-layer throughput performance in the field. In this paper, we report and analyze our findings from these flight tests (Sections II, III). Furthermore, we report the measured received power as a function of communication distance and their correlation (Section IV). These results provide baseline performance information on 802.11a wireless links for UAV and ground node communication. They can be useful for future work in UAV networking, such as antenna selection strategies and multi-hop wireless networking.

II. Field Experiment Setup

Our networking testbed consists of a UAV node and several ground nodes, all equipped with 802.11a wireless devices. The ground nodes were placed on a line, with about 6 ft separation between the two end nodes. For these nodes, we used two types of single-board computers, made by Thecus and Soekris. Our UAV is based on the Senior Telemaster model [7]. We conducted our flight experiments at a private airfield in Woodstock, Connecticut. In these experiments, we used Atheros-chipset-based Wistron CM9 802.11a/b/g adapters, with 18dBm transmit power and channel 56 in the 802.11a band.

We encountered a serious problem where our single-board computer would interfere with the 72MHz R/C receiver on the airplane, possibly due to noise from the 66MHz system bus.

This material is based on research sponsored by Air Force Research Laboratory under agreeement numbers FA8750-05-1-0035 and FA8750-06-2-0154. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

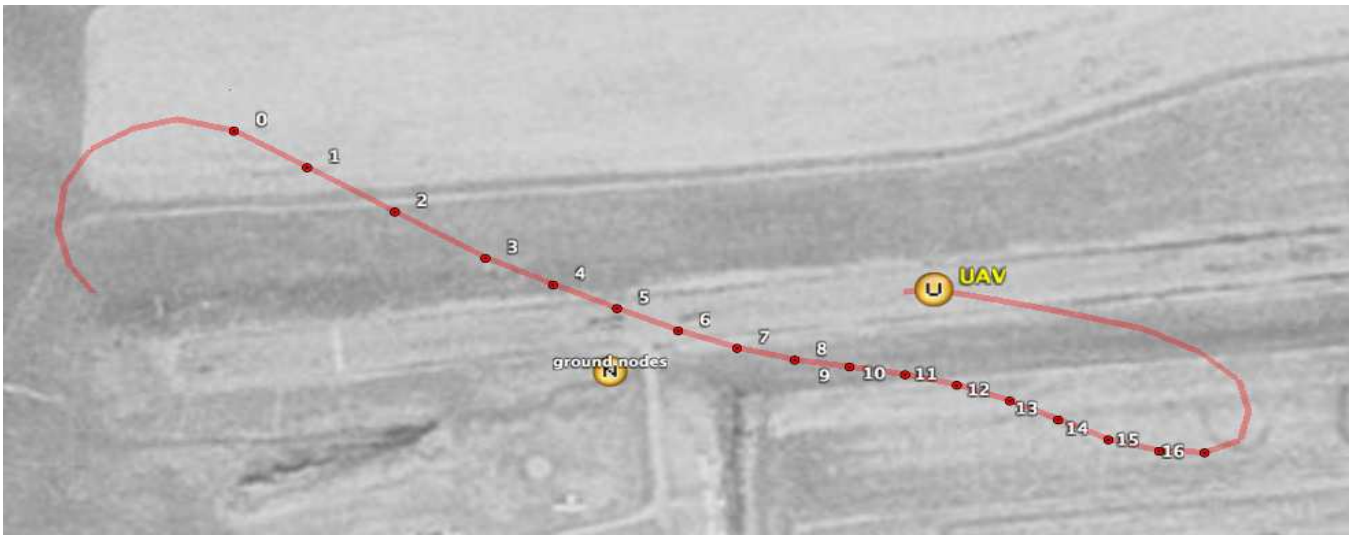

Fig. 1. **A U.S. Geological Survey (USGS) satellite map showing the location of an elevated ground node (N), node 3 in Figure 5, and a sample UAV path in a fly test. The horizontal, light-colored band under the UAV node (U) is the airport runway, which is approximately 25 yards wide. The dots show the positions reported by the GPS once a second.**

We solved the interference problem by doing three things: (1) moving the R/C receiver to the back of the airplane, (2) shielding the box hosting the single-board computer with metal screen wrap, and (3) moving the computer on/off switch and its wire—which was radiating the board noise—into the shielded enclosure.

We used two types of antennas on both the UAV and ground nodes. One was a 7-dBi, 2.4/5 GHz dual-band, omni-directional antenna purchased from a commercial vendor (Netgate), and the other was a custom 2-dBi dipole antenna. Samples of these antennas are shown in Figure 2. The key difference between these two antenna types is that the Netgate antenna produces an omni-directional beam that is much narrower in the vertical direction than the dipole, as can be seen from the manufacturer's radiation pattern plot in Figure 3.

The UAV was equipped with two wireless adapters, each with two antennas. The UAV would broadcast data packets using its four antennas in a round-robin manner. Each ground node was equipped with two wireless adapters, each with one antenna. Both adapters of a ground node could simultaneously receive UAV's packets. One of the four ground nodes was mounted on the top of a 14-ft wooden pole. Later in this section we will describe the antenna configurations and traffic patterns in detail.

The UAV flew approximately at 50-yard altitude and at 40 miles per hour over the ground nodes. The UAV had an on-board GlobalSat BU-353 GPS receiver, which provided position information at 1 second intervals. The GPS readings allowed us to visualize in real time, on a laptop, the UAV moving on a U.S. Geological Survey (USGS) satellite map shown in Figure 1. We have performed a coarse calibration of the GPS; our estimate of the position error was about 5 meters. The UAV GPS trace and the static ground node coordinates allowed us to analyze various performance parameters as functions of distance.

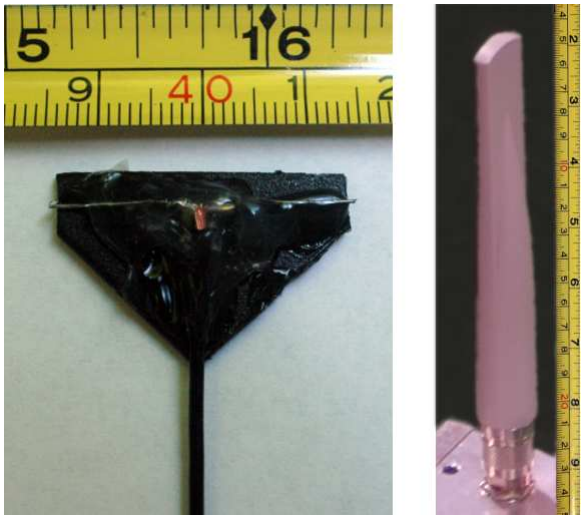

Fig. 2. **Samples of the two antenna types used in our testbed. Left: our hand-made dipole antenna, tuned to 5.28 GHz (channel 56). Right: the off-the-shelf Netgate antenna.**

A. Antenna Configurations

The antennas were configured as follows. First, let us define the following labels for referring to various antenna types and orientations:

H	horizontal dipole (i.e., dipole is parallel to the ground), orthogonal to flight direction
H_N	horizontal Netgate antenna, orthogonal to flight direction
H_p	horizontal dipole, parallel to flight direction
V	vertical dipole (i.e., dipole is perpendicular to the ground)
V_N	vertical Netgate antenna

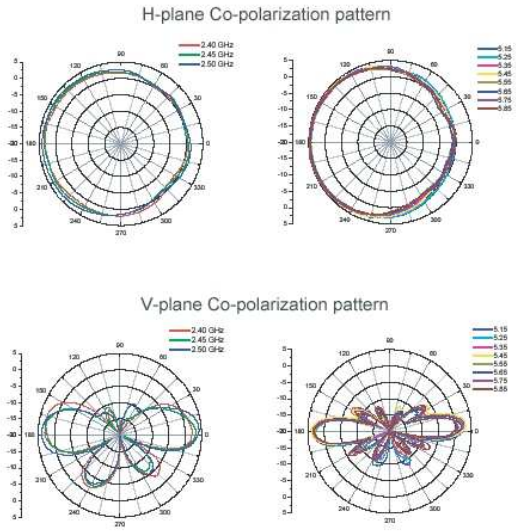

Fig. 3. **The radiation pattern of the Netgate antennas.**

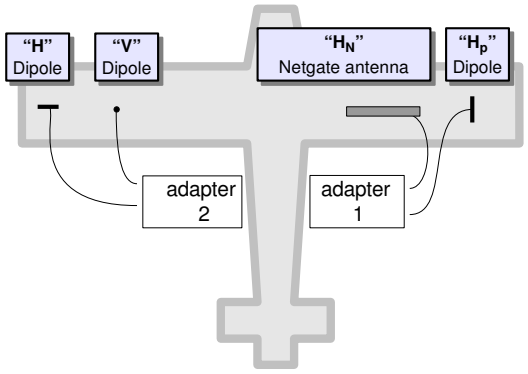

Fig. 4. **The UAV antenna setup.**

The UAV carried 4 antennas, H_p, H_N, H, and V, as depicted in Figure 4. The ground nodes each carried two antennas—one on each wireless adapter. The following table lists their antenna orientations, derived relative to a straight flight path along the direction of the runway:

	Node 1	Node 3	Node 5	Node 6
Antenna 0	V	H	H	H_p
Antenna 1	H	V_N	V_N	V_N

Figure 5 depicts the flight pattern of the UAV and the orientations of the ground antennas.

B. Description of Traffic Patterns

The UAV was the sole data transmitter during the experiment. It generated an endless stream of sequenced 320-byte UDP packets (which means roughly a 500 μs packet transmission time at the 6Mbps rate) and broadcast them over its 4 antennas in an approximately round-robin order. More specifically, a user-mode program alternated enqueueing pairs of packets into the first and second network adapter's socket queue; the kernel-level driver would add a timestamp and output each packet to the antenna identified by the least significant bit of the sequence number. For example, of the first 4 packets, numbered 0-3, packets 0 and 1 would go to the first adapter, while 2 and 3 would go to the second. Furthermore, packets 0 and 2 would be sent using each adapter's first antenna, while 1 and 3 would go out on the second antennas.

One reason for such a multiplexing scheme is to avoid interference between probe packets; in this scheme, the packets are interleaved in time so that at any moment, there is at most one packet in the air. It is for the same reason that we decided in the experiment to have the UAV node as the sole transmitter and evaluate only the performance of one-way communication from UAV to the ground. Full bidirectional measurements would otherwise require scheduled transmissions from ground nodes, which would lead to unacceptably large guard times and thus significantly decrease the temporal resolution of the measurement. Fortunately, under the typical symmetric-link assumption in free-space or nearly free-space propagation models, these unidirectional link measurements can still be useful in characterizing the bidirectional UAV-ground links.

Under this multiplexing scheme, ideally each antenna would send one packet every 2ms. Indeed, individual transmit queues always contain packets for alternating antennas; however, the combined output of the two adapters can not be perfectly interleaved since the sending pattern is subject to the random backoff in the 802.11 CSMA mechanism. We measured the resulting interleaving pattern in the lab, and found that runs of packets from the same adapter had at most 7 packets, while their mean length was 1.53 packets.

The ground nodes captured the broadcast packets using two wireless adapters and recorded the transmit timestamp, sequence number, size, and the RSSI figure. This way, from the data traces of just one ground node we can obtain the performance for the 8 different links created by the combination of 4 UAV and 2 ground node antennas. Therefore, the combined traces of all 4 ground nodes contain the performance for 32 different antenna combinations.

The reason for measuring so many link combinations nearly simultaneously is to eliminate the variations that would in-

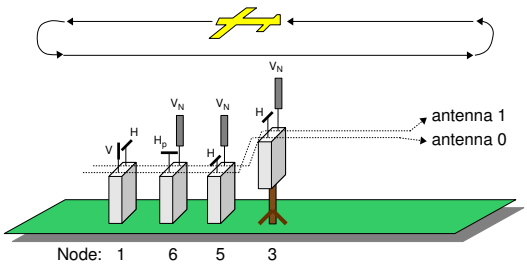

Fig. 5. **Orientations of the antennas on the ground nodes relative to the UAV flight pattern. The thick gray sticks represent Netgate antennas.**

evitably occur if we measured the different links using separate UAV flights. All of our 32 traces contain data points that lie at most several milliseconds apart, which is short enough to regard many physical parameters of the environment constant. For example, it may take several hundreds of milliseconds for the bank angle of the UAV to change enough to appreciably affect the receiver's position in the antenna pattern. This is also long enough that we can fairly compare many interesting properties for the 32 links, such as throughput, signal strength and packet loss.

III. Measured Throughput as a Function of Antenna Orientation

In this section, we report the throughput measurement results, based on which we compare the performance of several different antenna configurations and identify the best one.

As described in Section II, in the experiments, the UAV node was constantly sending out UDP packets at 6Mbps, or 1.5Mbps from each of the four transmit antennas. By counting the number of received packets at a receive antenna in a short period of time, we can measure the instantaneous performance of that particular antenna orientation configuration, which we will call "the UDP throughput." We use this throughput as the main performance metric in evaluating various antenna orientation configurations.

The total flight time in the two flights reported is approximately 24 minutes, during which the UAV node sent out more than 2.4 million packets. The total number of packets received at the eight antennas of the four ground nodes is about 1.8 million; however, most of the packets are received by more than one antenna, and therefore, the achieved end-to-end throughput, averaged over all antennas, is merely 120.8 kbps. Due to the relatively large flight area, the UAV node and the ground nodes are out of each other's communication range for a significant portion of time. For this reason, out of all possible antenna orientation configurations, even the best configuration (horizontal transmit antenna to elevated horizontal receive antenna) only receives about 33% of the packets. We plot the UDP throughput of the top four best-performing antenna configurations versus distance in Figure 6.

There are four curves in Figure 6. Following the order described in the legend, the topmost curve represents the throughput achieved from a horizontally oriented dipole antenna on the UAV node to another horizontally oriented dipole antenna on an elevated ground node; this combination is the best antenna orientation configuration we have seen in this experiment. The second topmost one differs from the previous one in that the ground node is not elevated. By comparing the two, we can see that an elevation of 14 feet helps achieve a significantly higher throughput. The third curve shows the throughput from a vertically oriented dipole antenna on the UAV node to another vertically oriented dipole antenna on a ground node. It is interesting to note that, although inferior to other three configurations when the distance is small, this configuration actually outperforms the two horizontal configurations in which the ground node is not elevated when the distance is large enough. This is because when the distance is small, the vertically oriented dipole antennas on the transmitter and the receiver are more likely to be in each other's null, resulting in worse performance, whereas it is less likely for the two antennas to be in each other's null when the distance is large. Lastly, the bottom curve corresponds to a pair of cross-polarized dipole antennas on the UAV node and a ground node. To our surprise, this configuration actually performs quite well, especially at the farthest distance of around 300 meters. We believe that this is because when the UAV is at that distance, it is probably banking at sharp angles such that it turns back towards the ground nodes; at this time, the antennas are no longer cross-polarized. We have looked more closely into the GPS trace and found that the furthest distance between the UAV and the ground nodes during the entire course of the flight is around 300 to 350 meters, further confirming our theory.

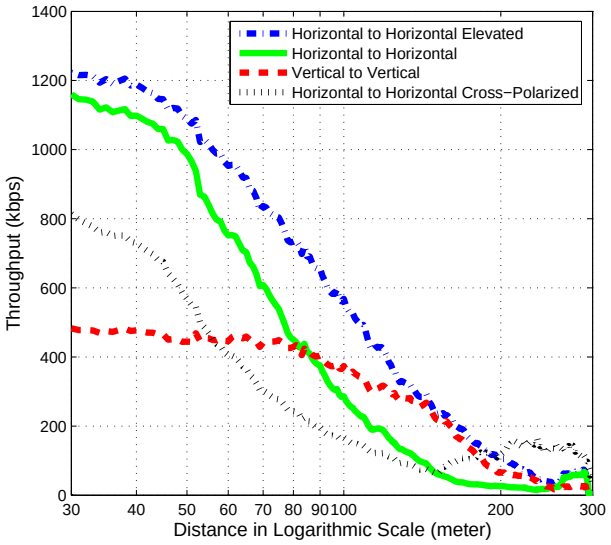

Fig. 6. **The UDP throughput of the top four best-performing antenna configurations.**

We summarize the best throughput results for various antenna orientation configurations in Table I. Overall, the horizontal/horizontal antenna orientation configuration has the best performance; furthermore, elevating the ground node can also help improve throughput. Cross-polarization in general has a negative impact on performance; this is evident from the vertical/horizontal combinations in Table I. Finally, we find that the off-the-shelf Netgate antennas perform poorly compared with dipole antennas, most likely due to their narrow beams along the vertical direction.

IV. Measured Received Power As A Function of Distance

In this section, we investigate the correlation between measured received power and distance. In particular, we perform linear regression on received power and distance following a

Transmit Antenna	Receive Antenna	Throughput
H	H Elevated	433 kbps
H	H	289 kbps
V	V	246 kbps
H_p	H Elevated	223 kbps
V	H Elevated	160 kbps
H_p	H	143 kbps
V	H_p	137 kbps
H_p	V	119 kbps
V	V_N	110 kbps

TABLE I

The throughput performance of various antenna orientation configurations.

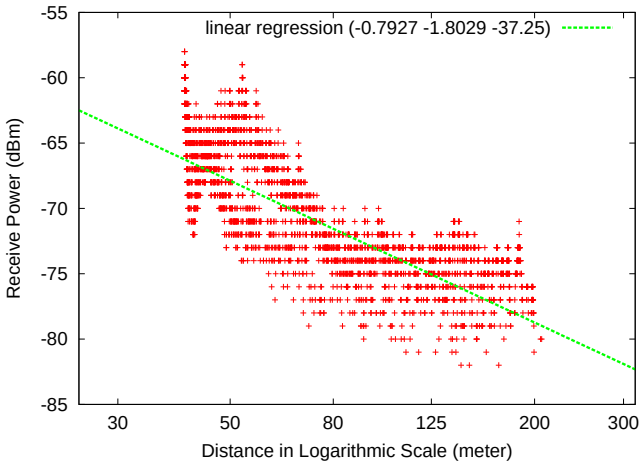

Fig. 7. **Received power vs. distance and linear regression result for the flyover shown in Figure 1. Each marker represents a packet, and the numbers in the parentheses are correlation coefficient, slope, and intercept.**

log-distance path loss model. The correlation coefficient of the linear regression can tell us whether there is a correlation between received power and distance. We are also interested in the slope of the linear regression, because it gives us the path loss exponent of the environment.

Under the log-distance path loss model, received power (measured in dBm) is expressed as a function of the logarithm of distance. More specifically, the received power $P_r(d)$ at distance d can be computed from received power $P_r(d_0)$ at distance d_0 with the following formula (α being the path loss exponent):

$$P_r(d) = P_r(d_0) - 10\alpha log_{10}(\frac{d}{d_0}) \quad (1)$$

For each packet received by the ground nodes, we logged its RSSI as reported by the Atheros cards. We then derived the received power for each packet from the reported RSSI. It has been shown that there is a constant difference of -95 between the RSSI and the actual received power when RSSI is greater than 6 [9]. Based on that finding, we derived received power (in dBm) by adding -95 to the reported RSSI.

We recorded GPS readings on the UAV node and on the elevated ground node (node 3 in Figure 5) to derive distance between the two nodes. Because the GPS only reports coordinate every second, we need to estimate the coordinate of the UAV when a packet is being transmitted. We estimated the coordinate by performing linear interpolation between two enclosing GPS reports and use it to compute distance.

As input to the correlation computation we use data from the 17 second flyover segment depicted in Figure 1. We chose this particular segment because its path is close to a straight line, so the variations in relative antenna orientations due to UAV turning and banking are expected to be small.

Figure 7 shows the correlation of the measured received power and distance during this flyover; each red marker represents one packet. The packets plotted are those transmitted from the UAV H antenna and received by the H antenna of Node 3 (the elevated node). We chose this particular antenna pair because the distance—not randomness induced by, e.g., ground reflection—should be the dominating effect on the receive power due to elevation. The straight line in the figure represents the output of linear regression. Specifically, the correlation coefficient is -0.7927, slope -1.8029, and intercept -37.25.

The correlation coefficient indicates there is a good correlation between received power and distance. However, the absolute value of the slope, which can be interpreted as the path loss exponent, is lower than expected: for this experiment, we expected α to be greater than 2, but the slope of -1.8029 suggests a path loss exponent less than 2.

We believe that such low path loss exponent can be explained by the limitation of our 802.11 equipment and the dynamic range of the received power in this experiment. More specifically, when the received power of a packet is too low, the equipment can not decode it; as a result, no RSSI is reported for that particular packet. As the distance increases, there will be more packets that can not be decoded due to low RSSI. This causes the distribution of received power vs. distance to be skewed when the distance is large. That is, at long distance only packets of sufficiently high received power are recorded, while packets of low received power are dropped. As a result, absence of packets of low received power results in a skewed regression slope.

To better measure the correlation of received power and distance, we would need to revise our experimental setup. We would need to raise the operating received power range of our experiments so the distribution of receive power are not distorted at long distance. We would also want to repeat the experiments using 802.11b/g in order to obtain measurements of wider range of distance.

V. Discussion

In Section III we reported throughput measurements based on data taken during 24 minutes of UAV flights. We now take a closer look at the flyover segment depicted in Figure 1, in order to study transmission performance apart from the various

other segments where the UAV motion is not as regular. Furthermore, we narrow down the data set to that received by antenna H of the elevated ground node (node 3), since it would be the least affected by ground effects.

Table II lists the throughputs achieved during the flyover from the four UAV antennas. The performance of antenna pairs orthogonal to the flight path (H-H and H_N-H) is noticeably better than the other two; even the second-best pair performs more than twice as well as the third-best. On the other hand, the worst performer is the V-H pair, possibly due to cross-polarization and the fact that the receiver is placed closer to the antenna null region of the transmitter.

H	0.63101 Mbps	(42.1%)
H_N	0.59425 Mbps	(39.6%)
V	0.23607 Mbps	(15.7%)
H_p	0.16682 Mbps	(11.1%)

TABLE II

The throughput performance of four UAV antennas to the H antenna on the elevated ground node, expressed in megabits per second and as fractions of the maximum possible throughput.

Figure 8 shows the raw RSSI data from the flyover. This data agrees with the throughput measurements in that the strongest signal comes from the H and H_N antennas. Furthermore, the plots uncover two additional observations.

First, we note that even though its performance is second best, the H_N antenna exhibits a significant peak around 10s into the flyover. This is not an isolated incident, since we observed similar peaks on other flyover traces. We believe that this is caused by the H_N antenna's narrow beam pattern—as the UAV flies and banks at varying angles toward the ground nodes, it occasionally "hits" them with the main lobe of the H_N antenna. The narrow beam seems to be a disadvantage here, since it increases the variation in link quality without actually beating the wide-beam antenna in throughput performance.

The second observation is that the signal strength of the best antenna pair does not vary smoothly as the UAV flies over the ground nodes; instead, we can see at least three major peaks in the curve for antenna H. Since the antennas in this pair have axially symmetric beam patterns, the peaks cannot be explained by the UAV motion alone. Instead, we believe that the cause is interference from a reflected ray as modeled by two-ray propagation.

Some receiving antennas, other than those in Figure 8, performed very poorly. For example, the V_N antenna of Node 5 did not receive any packets at all. It turns out that most poor performers were V_N antennas; we believe that their narrow horizontal beam patterns were largely underneath the UAV. This further illustrates the difficulty with use of directional antennas.

VI. Conclusions

Our measurement data have shown that, for UAV's communication with a ground node, horizontal dipole antennas with

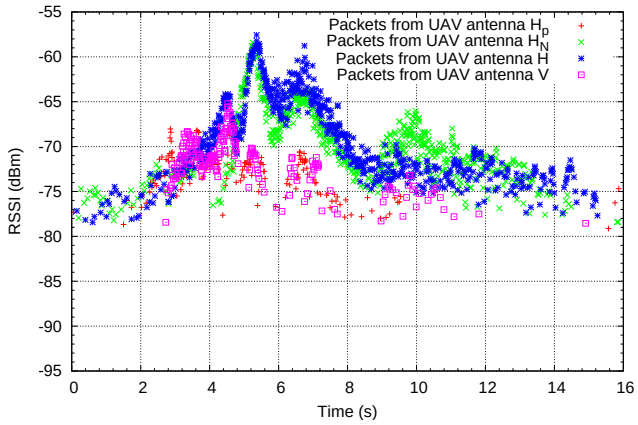

Fig. 8. **There are four sets of points shown, corresponding to packets coming from each of the UAV's four antennas.**

their nulls pointing to a direction perpendicular to the UAV flight path yields the highest throughput among 32 antenna pair configurations. In addition, the measurement data suggest that the path loss in an airfield environment is roughly proportional the square of the communicating distance. These results appear to be among the first antenna measurement results for 802.11 based UAV networking. We chose to start our measurement work with 802.11a because there is relatively less interference from the environment in the 5GHz band. We plan to conduct similar measurements for 802.11 b/g in the future.

References

[1] Li, J., Blake, C., De Couto, D. S. J., Lee, H. I., and Morris, R., "Capacity of Ad Hoc Wireless Networks," *ACM MobiCom*, July 2001

[2] Aguayo, D., Bicket, J., Biswas, S., Judd, G., and Morris, R., "Link-level Measurements from an 802.11b Mesh Network," *SIGCOMM 2004*, Aug 2004

[3] Bicket, J., Aguayo, D., Biswas, S., and Morris, R., "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *ACM MobiCom 2005*, Aug 2005

[4] IEEE 802.11 Working Group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications," *IEEE 802.11 standard, including 802.11a and 802.11b extensions*, Sep 1999

[5] Haykin, S., "Communication Systems," *John Wiley & Sons, Inc.*, 2001.

[6] Lee, W. C. Y., "Mobile Communications Engineering: Theory And Applications," *McGraw Hill*, 1997.

[7] Hobby Lobby International, Inc., "Senior Telemaster R/C Airplane," *http://www.hobby-lobby.com/srtele.htm,* 2006.

[8] Brown, T. X., Argrow, B., Dixon, C., et al., "Ad Hoc UAV Ground Network (AUGNet)," *AIAA 3rd "Unmanned Unlimited" Technical Conference,* Chicago, IL, 2004

[9] Judd, G., and Steenkiste, P., "A Simple Mechanism for Capturing and Replaying Wireless Channels," *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005

Parallel Use of Multiple Channels in Multi-hop 802.11 Wireless Networks

Chen-Mou Cheng Pai-Hsiang Hsiao H. T. Kung Dario Vlah
{doug, shawn, htk, dario}@eecs.harvard.edu
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138

ABSTRACT

We consider parallel use of multiple channels in a multi-radio, multi-hop 802.11 wireless network, with the goal of maximizing the total multi-hop throughput. We first quantify several fundamental forms of radio interference that cause performance degradation when the number of hops increases and that prevent total throughput from scaling up with number of radio interfaces at each node. We then evaluate three different methods of parallel channel use: Ad-Hoc, Frequency-Division Multiplexing (FDM), and Time-Division Multiplexing (TDM). We measure their performance on a linearly connected multi-hop network of dual-radio nodes. Although theoretically these three methods should have comparable performance, their actual measured performances are quite different. We find that TDM has the best performance, followed by Ad-Hoc and then FDM. The performance differences are due to these methods' capabilities of combating interference. We conclude that interference, especially adjacent channel interference, has significant effect on the achievable performance of a multi-radio, multi-hop network and hence should be carefully taken into account in the design and deployment of such a network.

I. INTRODUCTION

There has been a growing interest from the research community and industry in using low-cost Commercial Off-The-Shelf (COTS) wireless equipment, e.g., IEEE 802.11 wireless LAN ("WiFi"), for a variety of wireless networking applications. In particular, there has been substantial research and deployment activity in WiFi-based wireless multi-hop networks [2], [4], [15].

However, several studies in recent years have independently shown that the performance of such networks deteriorates rapidly as the number of hops increases, due to reasons such as inefficient medium access control, radio interference, wireless link errors resulting from changing channel conditions and multipath effects, frequent route changes, and improper TCP reaction to packet loss caused by poor link quality [12], [15] . In the future, newer and longer-range technologies such as WiMAX [22] may alleviate some of the problems by using, e.g., time-division multiplexing (TDM) protocols, but it is still too early to tell how wireless multi-hop networks based on these new technologies will perform in practice.

Some of the typical approaches for enhancing the performance of multi-hop wireless networks call for use of directional antennas [14], [24], or parallel use of multiple omni-directional radios at each node [3]. These would allow a node to transmit and receive at the same time using two separate beams over the same frequency band, or, respectively, using two separate channels over two radios. Our solutions are instances of the latter approach, where each node employs two COTS radios with omni-directional dipole antennas. The approach is attractive because COTS equipment has advanced enough that it is relatively inexpensive to incorporate two radios per node, and the deployment of nodes with omni-directional antennas is relatively easy in the sense that it does not require sophisticated antenna-specific engineering. Among 802.11 COTS equipment, we are mostly interested in the newer OFDM-based 802.11a/g systems since they can deliver higher performance in terms of throughput.

The rest of this paper is organized as follows. In Section II, we identify several radio interference issues related to 802.11 multi-hop networks. Specifically, we

This material is based on research sponsored by Air Force Research Laboratory under agreement numbers FA8750-05-1-0035 and FA8750-06-2-0154, and by the National Science Foundation under grant number #ACI-0330244. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, the National Science Foundation, or the U.S. Government.

observe that the commercially available 802.11 chips can reliably decode packets only when the Signal to Interference plus Noise Ratio (SINR) is fairly high; for example, the OFDM-based 802.11a/g radios require at least 20dB SINR for decoding 1000-byte packets at 8% packet error rate. This decoding capability is generally worse than what people would expect. Furthermore, in environments with low path loss, there is only a slow increase of SINR with increasing interference distance. This explains why the achievable throughput in a multi-hop system decreases with the number of hops.

In Section III, we consider the case when each of the nodes in a multi-hop system is equipped with multiple radio interfaces and discuss the adjacent channel interference (ACI) issue that prevents the total throughput from scaling up with the number of radio interfaces. In Section IV, we present three methods of parallel channel use: Ad-Hoc, FDM, and TDM. In Section V, we report our performance measurement results for the three methods applied to a linearly connected multi-hop network of dual-radio nodes.

We note that our target application scenario is multi-hop networking for low-attitude unmanned aerial vehicles (UAVs), in which multi-hop forwarding is used to extend network range and provide communication beyond line of sight [5], [7]. We choose this application for two reasons. First, UAV-based networking offers many inherent application advantages, including the UAVs' capability of carrying a huge amount of on-board data relevant to the current mission and dynamically positioning themselves over regions of interest. Second, free-space inter-UAV communication suffers less from multipath complications resulting from reflections off the ground and structures on the ground such as buildings and hills. This provides us with a clean environment where inherent interference issues in a multi-hop system can be observed and analyzed more clearly. It is for the same reason of keeping the scenario as simple and clear as possible that we focus our discussion and analysis on multi-hop linear arrays.

II. Inherent System Performance Limitations

Many researchers (e.g., Li et al. [15]) have examined the causes for the poor performance of multi-hop packet forwarding commonly observed in wireless ad-hoc networks. Among the many causes, the interference created by the packets of a connection when they travel along a chain of forwarding nodes imposes a stringent limit on the degree of possible spatial reuse and hence on the realizable utilization of available bandwidth. Namely, if any two hops in a chain of k hops interfere with each other when operated simultaneously, then the realizable utilization of this chain can not be more than $\frac{1}{k}$ because there can be at most 1 active hop at a time, and each packet has to traverse at least k hops. Li et al. also verified this reasoning with both simulation and experimentation and found a close match between the two, suggesting that this $\frac{1}{k}$-limit is probably the single most important factor contributing to the poor performance of multi-hop forwarding.

A popular model used in the ns-2 network simulator dictates that the interference distance is typically about twice that of the maximal effective communication distance; in this case, k equals to 4. However, in an outdoor experiment using ground nodes at about 20-inch elevation that we have conducted in a yard near our office building in Cambridge, MA, we observe that the interference can cover as much as four times the maximal effective communication distance for our commercially available 802.11a equipment. This makes k at least 6, rendering efficient multi-hop forwarding almost impractical for large numbers of hops. We note that our finding agrees with that of Padhye et al., who also reported the inaccuracy of the ns-2 model when experimenting with commercially available hardware [18].

We look closely into the reasons why the interference can travel so far by setting up a laboratory experiment. The experiment is designed to determine the minimum SINR required for our testbed hardware to successfully decode packets. We set up two testbed nodes, each equipped with a Wistron Neweb CM9 mini-PCI network adapter (based on the Atheros AR5213A 802.11a/b/g chipset) and use one of them as the transmitter (Tx) and the other as the receiver (Rx). We create a noise source using an Agilent 33250A 80MHz function/arbitrary waveform generator and a Maxim MAX2820 2.4GHz zero-IF transceiver: Maxim MAX2820 up-converts the baseband noise generated by Agilent 33250A to the appropriate 802.11g passband. We then connect Tx, Rx, and our noise source with a signal combiner that combines the Tx signal and the generated noise into Rx, as depicted in Figure 1.

We modify the `madwifi` driver of the Atheros chipset to let Tx broadcast 1000-byte packets continuously at the link speed of 6Mbps on channel 6 (2.427–2.447GHz). We then measure the number of packets received at Rx. We adjust the output gain of the Maxim MAX2820 transceiver until we have a packet error rate of approximately 8%. At this point, we unplug the cable going into Rx and plug it into a spectrum analyzer. The measured

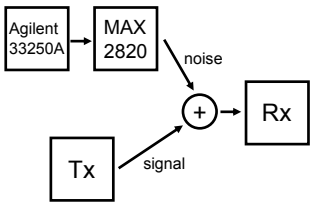

Fig. 1. **The laboratory experiment setup for measuring the required SINR in order to decode 6Mbps 802.11g packets with a packet error rate of about 8%**

Tx channel power is -42.5dBm, whereas the channel power of the noise source is -62.0dBm; in other words, our testbed hardware requires about 19.5dB SINR in order to achieve a packet error rate of about 8%.

Contrary to the general belief (e.g., [20]) and to the specification of the Wistron Neweb CM9 network adapter, the SINR required for 802.11g to achieve a packet error rate of 8% at 6Mbps is not merely a few dBs—but as large as 20dB[1]. This means that the interference distance can be as large as 10 times of carrier sensing range (k=12), assuming a free-space propagation model, or 3 times (k=5) assuming a lossy-terrain propagation model. Clearly, the multi-hop throughput will be poor if k is so large.

Because of the infeasibility of sustaining reasonable throughput with a long chain of forwarding nodes, we argue that one should focus on shorter chains when aiming to achieve a high throughput in a wireless ad-hoc network. If one needs to transfer a large amount of data over a long range, one should probably traverse a small number of hops, terminate the connections at some intermediate buffering nodes, and restart the connections when appropriate. This approach is similar in spirit to the Delay Tolerant Networks framework [6]. Following this reasoning, we shall perform our experiments in a linear topology consisting of no more than four or five nodes when we investigate various techniques of exploiting parallel channels with commercially available 802.11 radios.

III. Issues with Multi-Radio Methods

We examine a number of issues we have encountered in experimenting with multi-radio, multi-hop networks. We loosely divide them into issues unique to multi-radio nodes and those that apply to any multi-hop system.

[1]We note that this 20dB decoding margin still conforms to the IEEE 802.11g standard [11], c.f. 17.3.10.1, Table 91—assuming a -102dBm noise floor at the antenna connector in room temperature.

A. Issues Specific to Multi-Radio Nodes

Rx-Tx Adjacent Channel Interference. The *Rx-Tx ACI* occurs when a multi-radio node transmits on one channel and, due to imperfect transmit filters in radio hardware, outputs part of the RF power into a second (adjacent) channel that another radio happens to be receiving on at the same time.

Other researchers have already encountered the ACI problem and found that combating the interference generally requires isolating the affected radios [1], [19]. For example, Robinson et al. achieve an improvement after increasing the antenna separation and shielding the radio cards [19].

In a separate paper, we reported measurements of the ACI interference using a spectrum analyzer [8]. We include some of those measurement results in Table I: the values represent relative interference power in 802.11a channels adjacent to channel 52. From this data we can see, for example, that a transmitter on channel 52 would introduce interference power to the adjacent channel 48 about 27dB lower than its main signal. Unfortunately, there is very little path loss to a co-located radio; even reduced by 27dB, the interference power is still significant, especially when competing with a faint sender.

TABLE I
Channel 52's relative interference power in the adjacent channels of 802.11a

Ch.	36	40	44	48	52	56	60	64
dB	-59	-59	-52	-27	0	-26	-53	-57

Tx-Tx Adjacent Channel Interference. The *Tx-Tx ACI* occurs when the out-of-band power from one transmitting radio is mistakenly recognized as an active carrier at another radio attempting to transmit on a different channel, causing the latter radio to back off. This prevents a node from transmitting simultaneously on both radios, even though they supposedly operate on orthogonal channels. In contrast to Rx-Tx ACI, the severity of the Tx-Tx ACI only depends on the isolation between radios on the host node; it is independent of the distance to the destination node. Lastly, the effect of the Tx-Tx ACI could possibly be avoided altogether by disabling the carrier sensing mechanism.

We measured the effect of Tx-Tx ACI by emitting a broadcast stream from two radios on a single node. The radios were set to adjacent 802.11a channels (20MHz between center frequencies) and operated at 24Mbps. We varied the offered load from 0–21Mbps, for several

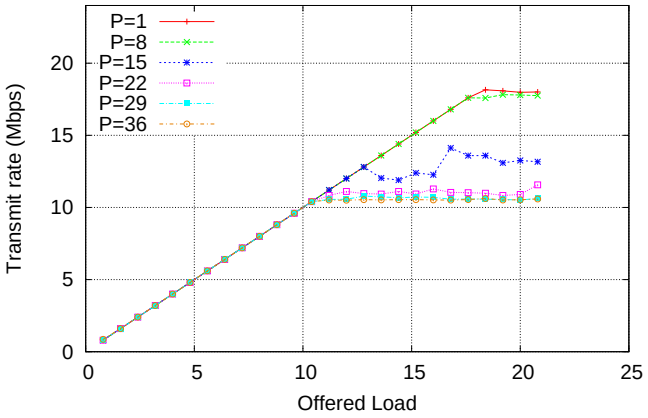

Fig. 2. **Effect of Tx-Tx ACI on carrier sensing behavior (24-Mbps modulation). Transmit power levels (P) are given in Atheros driver units**

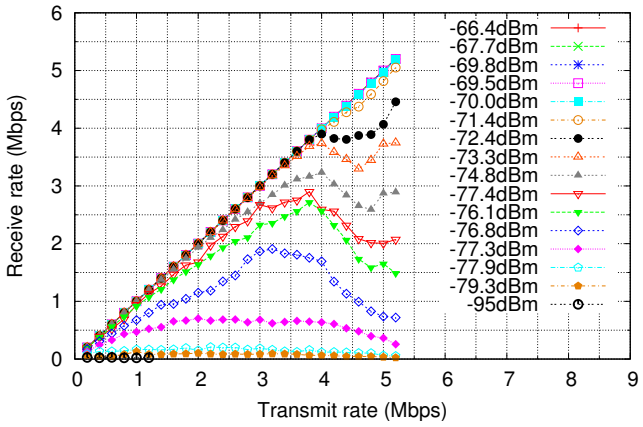

Fig. 3. **UDP broadcast performance of a marginal link for a range of signal strengths and offered loads**

different power settings. We report the observed output rates of one of the radios in Figure 2. We can see that with any power setting at or above 22, the achieved output rate reaches a peak about half of that achieved with power of 8 or lower. This is because at the higher powers, the carrier sensing mechanism gets triggered by the other, out-of-band transmission. For power settings between 8 and 22, it seems that the carrier sensing behavior is erratic; we ascribe it to the way in which the Atheros radios select the carrier sensing thresholds.

B. Issues in Single- or Multi-radio Multi-hop Systems

Co-Channel Interference from Hidden Terminals beyond Carrier Sensing Range. As we have detailed in Section II, it is possible for radios on the same channel to interfere each other even well outside their carrier-sensing range. Such a phenomenon is relatively well known in the literature; for example, Padhye et al. [18] examined this type of co-channel interference in an ad-hoc network by having all combinations of two nodes transmit at full rate and measuring the received rates everywhere else. This type of interference is the source of the well-known hidden-terminal problem [21]. Although the virtual carrier sensing mechanism (RTS/CTS) is supposed to alleviate it, we found it not quite effective in our preliminary tests; therefore, we decided to disable it in the rest of the experiments. Other researchers have also reported the ineffectiveness of RTS/CTS [15], [23].

Load Collapse on Marginal Links. We observed that when two nodes have a marginal link, increasing the transmitter's offered load resulted in a higher packet loss rate at the receiver. We measured this effect by having one node broadcast UDP packets at loads 0–6Mbps and 16 different power levels. We placed a receiver far enough to achieve a marginal link rate and recorded the incoming data rates as well as signal strengths. The results are shown in Figure 3; the 16 different power levels are identified by the average RSSI at the receiver. To rule out any problems with the sender, we placed a second receiver nearby and confirmed that it always received the full data rate. Note the sharp degradation around 4Mbps at some of the higher power configurations. We hypothesize that this effect might be a result of the specific way the Atheros radios set their receive sensitivity threshold. However, we have not been able to confirm this with Atheros engineers at this time.

These measurements show that the packet loss rate does not only depend on received signal strength, but also other transient factors; therefore, the assumptions of independent or identically distributed packet errors often used in wireless channel models may not be valid. For example, suppose we assumed that bit errors were independent. We might observe the 0.4 Mbps received for 1 Mbps sent at the -77.3 dBm power level, and conclude that the loss rate is 60%. This would lead us to predict that we should receive 2 Mbps when we send 5 Mbps across the same wireless channel. However, as we can see, the load response curve is distinctly non-linear, so the independence assumption is invalid.

To make sure that this phenomenon did not have a significant effect on our experiments, we first verified that each hop was able individually to achieve the full UDP throughput; thus, it operated above the marginal signal region which exhibits the problem.

IV. Three Methods of Parallel Channel Use

We consider the following three methods of parallel channel use: (1) Ad-Hoc, (2) Frequency-Division Multiplexing (FDM), and (3) Time-Division Multiplexing (TDM). In general, any good parallel-channel forwarding scheme should have the following properties:

1) **Infrequent collisions and low interference**: Reducing collisions and interference can be done at protocol level via the use of CSMA/CA and RTS/CTS. Alternatively, we can employ TDM by explicitly scheduling nodes' transmission so that no two nodes in the proximity transmit at the same time. Lastly, with FDM, we can allocate non-overlapping channels for each communication link. It is the main theme of this paper to determine experimentally how well these three methods perform.
2) **Aggressive spatial reuse**: Although the standard CSMA/CA and RTS/CTS mechanisms in 802.11 help reduce collisions and interference, they may lead to lower throughput by being overly conservative. Successful packet forwarding may still happen in the presence of a moderate amount of interference if the SINR is high enough (so called "capture effect"). To maximize throughput, a radio channel should be reused when the interference does not cause major performance degradation. In other words, we should seek spatial reuse as long as the link performance is acceptable in the presence of interference, rather than avoid it completely. For example, suppose we have a linear network of five nodes, A, B, C, D, and E, and we want to use the TDM method. Suppose each link can achieve 4 Mbps of throughput individually. If we use four time slots to completely avoid interference between links, we will have 1 Mbps end-to-end throughput. If, say, by operating A-B and D-E links simultaneously, we can obtain 0.7 Mbps on each link, then if we use three time slots, we will be able to achieve a higher end-to-end throughput by scheduling these two links in a time slot that is 1/0.7 times longer than the other two time slots; the resulting throughput is 4/(1/0.7+2)=1.17 Mbps.

In this section we describe the three methods and discuss interference issues that each method is likely to experience. A summary of each method's issues is presented in Table II.

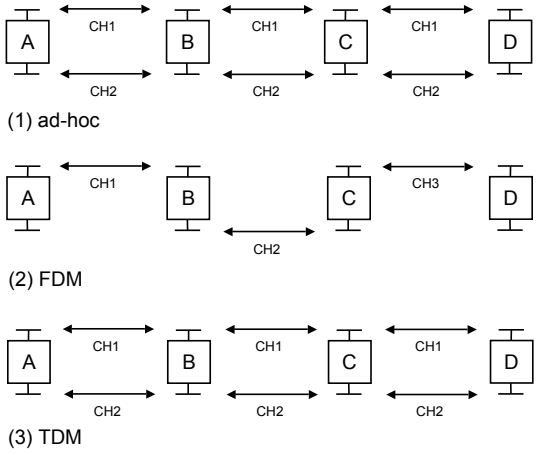

Fig. 4. **802.11 channel assignment examples for the three methods under comparison**

Let us introduce an exemplary network that consists of four nodes, as shown in Figure 4. Each of the four nodes (labeled A, B, C, D) is equipped with two radios (R_1 and R_2). The four nodes form a linear topology and are deployed so that any given node only has good links to its immediate neighbors, while its links to other non-adjacent nodes are intermittent at best. In our measurement tests, the network routes packets originating from node A to node D via B and C.

A. Ad-Hoc

Under the Ad-Hoc method, R_i on each node is assigned to 802.11 channel CH_i, for $i = 1, 2$, and multiple-access contention is resolved via the 802.11 CSMA/CA mechanism. In other words, we create two parallel ad-hoc networks, each enlisting one of the two radio interfaces from each node.

The two radios on each node can operate simultaneously; however, they share the channels with neighboring nodes using the CSMA/CA mechanism. Parallel use of the two radio interfaces this way is subject to the hidden-terminal problem and ACI (both Rx-Tx and Tx-Tx).

Hidden-terminal problem can occur when node C is outside of A's carrier-sensing range. More specifically, node A's packet transmission to B will not be received when C is also transmitting because C interferes with B's reception. We note that the RTS/CTS mechanism can only partially alleviate this problem; for example, it does not help when A interferes with D's reception from C because A and C are unlikely to be able to exchange RTS/CTS packets.

The Rx-Tx ACI occurs when one interface is receiving a packet while the other interface transmits—be it a

data packet, or a control packet like link-layer ACK. A possible way to reduce the Rx-Tx ACI while still maintaining the simplicity of ad-hoc channel access could be to "bond" the multiple radios by synchronizing all packet transmissions. In this way, any given node would likely either only transmit or only receive. We have not pursued this mechanism further, as it seemed that carrier sensing could disrupt the synchronized transmissions, and we did not have a means to disable it.

We note the severeness of both the hidden-terminal problem and Rx-Tx ACI increases with the load being transmitted at individual nodes. For the simple linear topology and single flow case, the interference may not be at its worst because the load is light at later hops. This is because the load at later hops depends on packets being delivered through the earlier hops. Thus, as some of the packets are lost due to interference, the load on later hops becomes lighter. We provide a more detailed discussion and a simple analysis of this point in Section VI.

B. Frequency-Division Multiplexing

Under the FDM method, each communication hop is assigned a unique channel. For example, node B uses channel CH1 on radio R_1 to communicate with node A and uses channel CH2 on radio R_2 to communicate with node C. With such a channel assignment, there is no need for multiple-access contention resolution such as CSMA/CA because there is a dedicated channel for each communication link. Ideally, there would be sufficiently many available channels, or enough path loss to allow reusing channels after a few hops.

In our example, the two radio interfaces on a node operate simultaneously—one for receiving packets from the upstream neighbor while the other for transmitting packets to the downstream neighbor. Clearly, use of multiple radio with this method will suffer from Rx-Tx ACI.

C. Time-Division Multiplexing

Under the TDM method, channels are assigned in the same way as in the Ad-Hoc method. The difference is that multiple-access contention now gets resolved through explicit scheduling. More specifically, we divide time into fixed-length time slots and let nodes transmit using both radios in their designated time slots. For the four-node network example we need three time slots to avoid contention. The number of time slots can be smaller than the number of hops if there are spatial reuse opportunities.

The two radio interfaces operate simultaneously during their designated time slot. Parallel channel use under TDM experiences Tx-Tx ACI and, for unicast traffic, Rx-Tx ACI due to the presence of ACKs. Because ACKs are not essential to our study of performance, we send broadcast packets instead of unicast packets so as to disable link-layer ACKs.

For TDM to work, nodes need to be synchronized in time, so all nodes can follow the schedule correctly. In addition, there is some queuing inside the operating system software that must be accounted for in order to properly schedule the transmissions.

Timing errors can reduce efficiency of parallel channel use due to overlapping time slots: during the overlap, the affected nodes either back off due to carrier sensing, or lose the data due to interference. With longer time slots, we can tolerate higher timing error. For example, 1 ms timing error per 100-ms time slot results in no more than 2% efficiency loss (1% at the beginning and 1% at the end of a time slot may overlap with adjacent time slots). But increasing the time slot length increases the end-to-end latency and the buffering requirements on each forwarder. This limits the size of the time slot to be a few hundred milliseconds in most practical applications.

There are several timing synchronization protocols that can achieve sub-millisecond accuracy, for example, Reference-Broadcast Synchronization (RBS) [9] and Continuous Clock Synchronization [17]. In our experiments, we manually synchronize our nodes using `ntpdate` before each experiment to achieve millisecond timing accuracy. Note that `ntpdate` is an NTP client that only synchronizes local clock's phase offset to a common server, and it does not continuously track timing error for local clock's frequency offset. Because our experiments only last a short period of time (several minutes at most), local clock's frequency offset has little effect on the timing accuracy.

The operating system employs packet queues to minimize overhead caused by context switching; often, these queues can buffer a large number of packets. For example, UDP sockets on Linux machines have a default queue size of 100KB, translating into more than 60 packets in our setting. Furthermore, the device driver uses a queue to amortize interrupt processing overhead. In the `madwifi` driver we use in our experiments, the default queue size is 200 packets. Given that the transmission timing control runs as a user process, packets stored in these two queues would continue to drain outside of the designated time slot even after the user process stops transmitting. To minimize the effect

TABLE II
Expected interference issues for each method

	Rx-Tx ACI	Tx-Tx ACI	Hidden Terminal
Ad-Hoc	√	√	√
FDM	√		
TDM	√ (ACK only)	√	

of queuing, we reduce these two queue sizes to 3 and 4 packets, respectively. In addition, we program the user process to stop sending data early to compensate for the transmission delay of queued packets (e.g., the application stops 14 ms before the end of the time slot if each packet takes 2 ms to transmit).

V. Measurement Results

We conducted a set of experiments to measure the throughput achieved by the three multi-radio, multi-hop methods described in Section IV. We ran the experiments on a testbed network consisting of 400MHz AMD Geode single-board computers made by Thecus Inc. We equipped each of these nodes with two Wistron Neweb CM9 mini-PCI network adapters, which are based on the Atheros AR5213A 802.11a/b/g chipset. The network adapters ran the `madwifi` Linux driver[2]. We used a mix of two types of antennas: (1) off-the-shelf, dual-band, 7dBi omni-directional antennas, and (2) 2dBi dipole antennas made in-house.

The experiments were conducted outdoors in a yard outside our office building as well as in a large athletic field across the campus. Our testbed consisted of five nodes, A, ..., E, which were arranged in a linear topology similar to that depicted in Figure 4. We set the transmit power equally on all nodes, and then placed the nodes on the ground so that each hop was as large as possible while still maintaining low-loss communication. We performed the 802.11a measurements in the yard outside our office building where we did not have any 802.11a networks and the 802.11b/g measurements in the athletic field where there were no other 802.11b/g networks.

All links ran at the lowest possible modulation rates—6Mbps for 802.11a/g and 1Mbps for 802.11b. Measured single-hop throughputs for these three modulations are 5.6Mbps, 4.4Mbps, and 0.9 Mbps, respectively, calculated from UDP payload based on 1472-byte packets.

[2]A newer version of the driver, called `madwifi-ng`, was available, but we didn't use it due to stability issues.

In all three multi-radio methods, a user-level source process at node A sent 1472-byte UDP packets toward a sink process at node E. In Ad-Hoc and FDM methods, we used a 100KByte kernel-level send buffer; when this buffer became full, the source process would be blocked from sending more packets. In contrast, in the TDM method, we buffered packets in a large user-level queue and used a minimal 7-packet kernel-level send buffer (3 in the socket buffer plus 4 in the hardware transmit queue) to avoid spillover outside the designated time slot. We have made sure that the buffer occupancies were reasonable in all three methods so that they would not have a significant effect on the throughput measurement result.

We divide the results into three subsections below: 1) comparing TDM, Ad-Hoc, and FDM on a three-hop chain, 2) comparing three- and four-slot TDM on a four-hop chain, and 3) comparing 802.11b and g. In all cases, the measurements are listed as three values: the measured end-to-end UDP throughput, the ideal rate achievable assuming orthogonal channels and zero packet-loss probability, and a ratio of measured to ideal rate for ease of comparison.

We calculate the ideal rates from single-hop capacities as follows. Let C be the single-hop capacity. For Ad-Hoc and TDM, we assume that the k hops (where k is either 3 or 4) are all in each other's interference range, and hence share the channel. For two-radio nodes, there are two independent k-hop forwarding chains where each hop gets C/k Mbps, and so the maximum achievable throughput is $2 \cdot C/k$ Mbps. With FDM on three hops, all hops can operate simultaneously at C Mbps, making that the end-to-end throughput. With FDM on four hops, we assume the first and last hop evenly share the same channel, leading to a bottleneck rate and ideal throughput of $C/2$.

A. Comparison of TDM, Ad-Hoc, and FDM over three hops

802.11a:

	Achieved	Ideal	Percentage
FDM	2.83 Mbps	5.6 Mbps	51%
Ad-Hoc	2.50 Mbps	3.7 Mbps	68%
TDM	3.30 Mbps	3.7 Mbps	89%

802.11g:

	Achieved	Ideal	Percentage
FDM	1.87 Mbps	4.4 Mbps	43%
Ad-Hoc	1.94 Mbps	2.9 Mbps	67%
TDM	2.78 Mbps	2.9 Mbps	96%

B. Three- vs. four-slot TDM over four hops

802.11g:

	Achieved	Ideal	Percentage
TDM-3	1.66 Mbps	2.9 Mbps	57%
TDM-4	2.00 Mbps	2.2 Mbps	91%

C. 802.11g vs. 802.11b over four hops

As stated above, since 802.11b and g have only 3 non-overlapping channels, at least one channel must be reused. In this case, we put the first and last hop on the same channel, leading to an ideal throughput of $C/2$. We note that Ad-Hoc performs quite well when 802.11b is used, possibly because 802.11b has a lower decoding margin, making it more resilient to interference.

802.11g:

	Achieved	Ideal	Percentage
FDM	1.40 Mbps	2.8 Mbps	50%
Ad-Hoc	1.49 Mbps	2.2 Mbps	68%
TDM	2.00 Mbps	2.2 Mbps	91%

802.11b:

	Achieved	Ideal	Percentage
FDM	0.34 Mbps	0.45 Mbps	76%
Ad-Hoc	0.37 Mbps	0.45 Mbps	82%
TDM	0.36 Mbps	0.45 Mbps	80%

VI. Discussion

In this section, we report the measurement result that demonstrates the impact of interference from hidden terminals beyond carrier sensing range, as described in Section III. Specifically, we compare the performance of a 3-slot TDM scheme for a 3-hop network (3HOP) and a 4-hop network (4HOP). These two cases are interesting because they are otherwise the same configuration, except that in 4HOP, the first hop and the last hop share the same time slots and hence will be transmitting simultaneously. Our initial 802.11g measurement result shows that the transmitters of these two hops can not sense each other's transmission; in other words, they both transmit at full rate as if there were no other transmitters in their neighborhood. Ideally, if there were no interference beyond carrier sensing range, then we would observe close, if not identical, performance results for 3HOP and 4HOP; however, our measurement result in Subsection A and B of the previous section shows that 4HOP's performance is only 60% of 3HOP's (1.66Mbps vs 2.78Mbps).

There is a simple calculation that we can do to quantitatively determine the impact of such interference. We first experimentally determine the effect of such interference by operating the two links simultaneously with full load. In this example, the interference is quite asymmetric: the first hop suffers much more than the last hop, most likely because the transmitter of the last hop is closer to the receiver of the first hop than the other way around. Specifically, when both links are busy, we get 30% throughput on the first hop and 96% on the last hop. Now in the TDM scheme, all packets originate from the first hop; if the transmission of packets on the last hop kills some of the packets on the first hop, then subsequently the last hop will get less packets, resulting in underutilization of the last hop and hence less interference for the first hop. In our experiment setup, all intermediate hops have reasonably good performance, so we are going to simplify the analysis by assuming that they are lossless; the same simplification can be applied to the last hop since its performance is fairly close to 100% even in the presence of interference from the first hop. Now call the probability of the transmission on the last hop killing the transmission on the first hop α (in this case, $\alpha = 0.7$). In equilibrium, suppose the percentage of packets getting through the first hop is x; this is going to be the utilization of the last hop since we assume the intermediate hops are lossless. These packets are going to kill αx of the packets on the first hop, making the throughput of the first hop $1 - \alpha x$, which must be equal to x by conservation of traffic flow. This means that $x = (1 + \alpha)^{-1}$, or 0.59 in our experiment, which is fairly close to the measurement result of 60%.

VII. Conclusion

In this paper, we considered the problem of efficient parallel use of 802.11 channels in multi-radio, multi-hop wireless networks. We experimentally evaluated three methods of parallel channel use, based on common MAC schemes (Ad-Hoc, FDM, TDM), on a linearly connected dual-radio network. The two main contributions of this paper are: 1) the finding that ACI is a significant obstacle for good performance in multi-radio networks; and 2) the proposed solution for mitigating the ACI effect using TDM channel access, supported by experimental evidence.

Among the three methods, TDM performs the best in all the experiments except for the four-hop case of 802.11b, in which Ad-Hoc performs slightly better than TDM. This is because TDM avoids most types of interference by design. By dividing channel uses into separate time slots, TDM schedules better than Ad-Hoc, but the two perform pretty closely when only a single

channel is used. Ad-Hoc, despite being the simplest method among the three, performs better than FDM because it manages to avoid some of the Rx-Tx ACI in FDM.

For nodes with more than two radios, we believe that among the three methods, only TDM's performance can scale up as the number of radios increases. The other two methods do not explicitly address the ACI issue, and so with more radios, the effect of ACI is likely to get worse.

In summary, the results of this paper show the importance of considering ACI in protocol design for multi-radio networks. Otherwise, as we have seen in the case of Ad-Hoc and FDM, real-world systems would not perform as well as expected. Thus, analysis taking into account ACI is required to design and understand the performance of multi-radio protocols properly.

REFERENCES

[1] Adya, A., Bahl, P., Padhye, J., Wolman, A. and Zhou, L., "A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks," First International Conference on Broadband Networks (BROADNETS'04), 2004

[2] Aguayo, D., Bicket, J., Biswas, S., Judd, G., and Morris, R., "Link-level Measurements from an 802.11b Mesh Network," *SIGCOMM 2004*, August 2004

[3] Alicherry, M., Bhatia, R., and Li, L., "Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks," *ACM MobiCom 2005*, August 2005

[4] Bicket, J., Aguayo, D., Biswas, S., and Morris, R., "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *ACM MobiCom 2005*, August 2005

[5] Brown, T. X., Argrow, B., Dixon, C., Doshi, S., "Ad Hoc UAV Ground Network (AUGNet)," *AIAA 3rd "Unmanned Unlimited" Technical Conference,* Chicago, IL, 2004.

[6] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and Weiss, H., "Delay-Tolerant Network Architecture," Internet Draft, March, 2006.

[7] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations," *ICCCN 2006*, Arlington, VA, October, 2006.

[8] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Adjacent Channel Interference in Dual-radio 802.11 Nodes and Its Impact on Multi-hop Networking," *IEEE GLOBECOM 2006*, San Francisco, CA, November 2006.

[9] Elson, J., Girod, L., and Estrin, D., "Fine-grained network time synchronization using reference broadcasts," *SIGOPS Operating System Review*, Volume 36, Issue SI, 2002

[10] Hyperlink Technologies, Inc., "2.4 GHz 802.11b and 802.11g Compatible 4-Pole Ultra High Q WiFi Bandpass Filters," *http://www.hyperlinktech.com/web/bpf24-4xx.php,* 2005

[11] IEEE 802.11 Working Group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications," *IEEE 802.11 standard, including 802.11a and 802.11b extensions*, September 1999

[12] Jain, K., Padhye, J., Padmanabhan, V. N., and Qiu, L., "Impact of Interference on Multi-hop Wireless Network Performance," *ACM MobiCom 2003*, September 2003

[13] Judd, G., and Steenkiste, P., "A simple mechanism for capturing and replaying wireless channels," *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005

[14] Ko, Y. B., Shankarkumar, V., and Vaidya, N., "Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks," *IEEE INFOCOM*, March 2000

[15] Li, J., Blake, C., De Couto, D. S. J., Lee, H. I., and Morris, R., "Capacity of Ad Hoc Wireless Networks," *ACM MobiCom*, July 2001

[16] Mikrotikls Ltd, "Routerboard 500," *http://www.routerboard.com/rb500.html* , 2006.

[17] Mock, M., Frings, R., Nett, E., and Trikaliotis, S., "Continuous Clock Synchronization in Wireless Real-Time Applications," *Symposium on Reliability in Distributed Software*, pp. 125–132, 2000

[18] Padhye, J., Agarwal, S., Padmanabhan, V., Qiu, L., Rao, A., and Zill, B., "Estimation of Link Interference in Static Multi-hop Wireless Networks," *Internet Measurement Conference (IMC)*, Berkeley, CA, October 2005

[19] Robinson, J., Papagiannaki, K., Diot, C., Guo, X. and Krishnamurthy, L., "Experimenting with a Multi-Radio Mesh Networking Testbed," *1st workshop on Wireless Network Measurements (WiNMee 2005)*, Trento, Italy, April 2005

[20] Thomson, J., Baas, B., Cooper, E. M., Gilbert, J. M., Hsieh, G., Husted, P., Lokanathan, A., Kuskin, J. S., McCracken, D., McFarland, B., Meng, T. H., Nakahira, D., Ng, S., Ratterhalli, M., Smith, J. L., Subramanian, R., Thon, L., Wang, Y.-H., Yu, R., and Zhang, X., "An Integrated 802.11a Baseband and MAC Processor," Proc. ISSCC 2002, San Francisco, CA, February 2002

[21] Tobagi, F. A., and Kleinrock, L., "Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple-Access Modes and the Busy-Tone Solution," *IEEE Transactions on Communications*, 23(12):1417-1433, 1975

[22] The WiMAX Forum, `http://www.wimaxforum.org`

[23] Xu, K., Gerla, M., and Bae, S., "How effective is the IEEE 802.11 RTS/CTS Handshake in Ad Hoc Network," Proc. IEEE GlobeCom'02, Taipei, Taiwan, Nov. 2002

[24] Yi, S., Pei, Y., and Kalyanaraman, S.,"On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas," *ACM MobiHoc 2003*, June 2003

Transmit Antenna Selection Based on Link-layer Channel Probing

Chen-Mou Cheng Pai-Hsiang Hsiao H. T. Kung Dario Vlah
{doug, shawn, htk, dario}@eecs.harvard.edu
Harvard University
Cambridge, MA 02138

Abstract

In this paper, we propose transmit antenna selection based on receiver feedback of channel information obtained via link-layer probing. Furthermore, we report the performance gain of the proposed antenna selection scheme in an experimental multi-antenna 802.11 network. We built a low-altitude Unmanned Aerial Vehicle (UAV) testbed using commodity dual-antenna 802.11 hardware and performed field experiments to collect traces of link performance using antennas of various types and orientations. Based on the collected traces, we demonstrate that transmit antenna selection can achieve a significant amount of gain using a link-layer channel probing protocol at a relatively low probing rate. The largest improvement we observed with joint transmit/receive antenna selection in 2×2 systems was 32%, about twice as much as that of receive-only antenna selection in 1×2 systems, which achieved 17%. Moreover, a similar improvement is obtained with probing intervals up to about 200 milliseconds, which is infrequently enough to consume only a small fraction of the available 802.11 channel capacity. Since these results require only a low implementation and operational cost, we conclude that transmit antenna selection is a worthwhile technique to use with the kind of multi-antenna mobile ad-hoc networks we examined.

1. Introduction

Antenna diversity is a well-known and commonly used technique for improving wireless communication performance. When multiple antennas are configured properly, they can take advantage of signals that traverse uncorrelated paths and thus compensate for the outages incurred in some of these paths. Several efforts have quantified the antenna diversity gain via the use of multiple antennas with proper polarizations or spatial separation [10]. Even in line-of-sight situations, the signal-to-noise ratio (SNR) gain can sometimes be as large as 12dB [3]. Furthermore, if the transmitter has knowledge of the channel state, then the channel capacity can be further increased by allocating more power to the transmit antennas with higher channel gain, a strategy known as the water-pouring or water-filling principle [7], [9]. To harness the diversity gain, complex and expensive radio-frequency (RF) circuitry is often required, e.g., for performing maximum ratio combining [8].

For the IEEE 802.11 wireless LAN ("Wi-Fi"), the use of multiple antennas has become easier in recent years due to the commercial availability of wireless LAN adapters equipped with dual antenna ports. These adapters use a switch to connect the antenna ports to only one set of RF transceiver circuitry. Moreover, such adapters implement receive antenna selection in the sense that they can detect the best receive antenna based on the signal strength measurements taken *within that packet's preamble* [6]. As a result, the probability that the data portion of the packet is received on the better of the two antennas is increased.

Dynamic antenna selection could be especially important in mobile networks, where the nodes' relative orientations change frequently. For example, in a scenario where a UAV with 802.11 communications capability is used for sensor data pickup or as a relay node, the maneuver of the aircraft due to turning or maintaining course could cause recurring changes in the relative positions in the radiation patterns of the communicating parties, as well as in the cross-polarization between the transmitting and receiving antennas [4], [5].

In this paper, we propose a transmit antenna selection scheme based on link-layer feedback from receivers on channel-probing packets. As a result, the sender can transmit packets on the antenna which, according to the feedback, yields better reception. We can use this transmit antenna selection together with receive antenna selection, which existing 802.11 hardware already supports.

Our evaluation of antenna selection consists of the following three units, which we present as the main contributions of this paper.

1. Trace collection. Via field experiments, we collected packet loss and signal strength data for constant bit rate (CBR) streams sent over 4 UAV antennas and received by 5 ground antennas. The resulting measurements describe 20 parallel channels in an outdoor environment under a mobility of about 40 miles per hour.

2. Evaluation of antenna selection using emulation. We built a multi-antenna channel emulator driven by pre-recorded packet traces, to do in-lab evaluation of antenna selection. Using the emulator, we evaluated the performance of a joint transmit/receive antenna selection implementation in the 60 2×2 systems constructible from the 4×5 channel traces obtained in step 1. We found that 1) within 2×2 systems, the joint antenna selection always performs better than the individual single-antenna channels, 2) the best overall 2×2 packet delivery rate (PDR) is 32% larger than the best overall single channel PDR, and 3) the best overall 1×2 PDR under receive-only antenna selection is 17% larger than the best overall single-channel PDR. These results show not only that the joint antenna selection provides significant gain, but also that half of that gain is contributed by feedback-driven transmit antenna selection based on channel probing.

3. Antenna selection field experiment. We tested the step 2 antenna selection implementation in a 2×2 UAV field experiment. We found that it was 6% better than with any single

This material is based on research sponsored by Air Force Research Laboratory under agreement numbers FA8750-05-1-0035 and FA8750-06-2-0154. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

antenna combination, though the sample variance was too high to draw definite conclusions. Even so, the measurements indicate that antenna selection works as could be expected based on the antennas' radiation patterns and their positions.

2. Channel probing protocol (CPP)

We now present the Channel Probing Protocol (CPP), a simple link-layer protocol that we use to inform the transmit antenna selection. Every probing interval of T_P seconds, the transmitter sends a probe packet over alternating transmit antennas. The probe is received on the best antenna using the receiver's hardware diversity circuit. The receiver feeds back to the sender the received signal strengths of the alternating probes, allowing the sender to choose the better transmit antenna to be used by subsequent packets. Figure 1 shows an example run of the protocol for three probing intervals between a transmit node A and a receive node B.

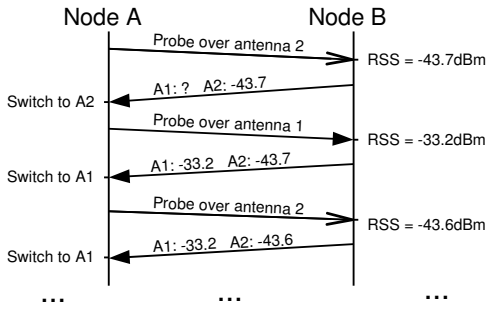

Fig. 1. An example CPP run. Labels $A1$ and $A2$ refer to the received signal strength of packets sent by A's antenna 1 or 2, respectively. Use of these antennas is indicated using solid and open arrows. Labels on the left side indicate the transmit antenna decision made using latest channel information fed back by node B. "RSS" stands for Received Signal Strength.

The sender's probes and receiver's reply packets can get lost over links with poor channel conditions. Channels under such conditions are not good candidates for communication anyway, so loss of their signal-strength information is not a serious problem. However, without up-to-date information about the bad channel conditions, the sender might use stale information obtained over a better channel. For this reason, the channel information expires after a certain timeout period, such as 5 probing intervals as we used in our implementation.

3. Trace collection

This section describes our trace gathering methodology and gives an overview of the collected data. We described a subset of these measurements and the methodology in a previous paper [4]. However, we did not address the issues of dynamic antenna selection in that previous paper. We begin by introducing our testbed, followed with a description of the antenna configurations. Lastly, we present the details of traffic patterns under observation and give a summary of the recorded data.

3.1. Testbed description

Our testbed consisted of a UAV node and 3 ground nodes. The ground nodes were placed on a line, with about 3 feet of separation between adjacent nodes. The nodes were powered by 400MHz AMD Geode single-board computers made by Thecus Inc. Our UAV was based on the Senior Telemaster [1] model, a training model known for its stable flight characteristics. The UAV flew in oval-shaped laps passing directly over the ground nodes, at speeds of around 40 miles per hour. In these experiments, we used Atheros-chipset-based Wistron CM9 802.11a/b/g radio cards, with dual antenna ports and automatic receive antenna selection as described in Section 1. Each testbed node had two of these dual-antenna radio cards. Like some previous works [2], in order to avoid interference and association delays, we turned off the 802.11 IBSS protocol by switching the radios into "ad-hoc demo" mode.

We used two types of antennas on both the UAV and ground nodes. One was a 7-dBi, 2.4/5 GHz dual-band, omnidirectional antenna purchased from a commercial vendor (Netgate). The other was a half-wavelength 2-dBi dipole antenna built in-house. The main difference between these two antenna types is that the Netgate antenna produces an omnidirectional beam that is much narrower in the vertical plane than that of the dipole antenna.

3.2. Antenna configurations

Our testbed's two antenna types were placed in a variety of orientations. In order to refer to these orientations in an unambiguous manner, we define the following of labels with respect to a level reference plane and a major axis that is either the direction of the runway for ground antennas, or that of the flight for UAV antennas:

H	horizontal dipole (i.e., dipole is parallel to the reference plane), orthogonal to the major axis
H_N	horizontal Netgate antenna, orthogonal to the major axis
H_p	horizontal dipole, parallel to the major axis
V	vertical dipole (i.e., dipole is orthogonal to the reference plane)
V_N	vertical Netgate antenna

We find it useful to specify the antenna orientations of multiple nodes at once. For this purpose, we introduce a notation for antenna configurations, which we define as pairs of antenna sets, and denote them as $\{a_1, a_2, ..., a_n\} \times \{b_1, b_2, ..., b_m\}$. The identifiers a_i and b_j are antenna labels. In each pair, the first set specifies the antennas on the UAV, while the second set specifies the antennas on the ground nodes. For example, the antenna configuration $\{V, H\} \times \{V, V_N\}$ refers to a vertical and a horizontal dipole antennas on the UAV, while a vertical dipole and a Netgate antenna on the ground nodes. This notation will be sufficiently descriptive for our purposes without mentioning specific nodes or radios to which the antennas are attached.

The antenna configuration that we measured in the flight experiment was $\{H_p, H_N, H, V\} \times \{V, H, H, V_N, H_p\}$. The UAV carried two antennas per radio card. The ground nodes had just one antenna per card, leaving one antenna port unused. There was a total of 6 radio cards on the ground nodes, but we only use the measurements from 5 due to hardware issues with one radio card. We put up two H receivers to obtain additional data for $\{H\} \times \{H\}$ systems, which were used in a co-located project regarding load-carry-and-deliver (LCAD) networking [5].

3.3. Description of traffic patterns

The UAV was the sole data transmitter during the experiments. It ran a user mode program that broadcast an endless stream of sequenced 320-byte UDP packets at the 6Mbps

modulation under 802.11a. At that rate, the transmission time of each packet was roughly 500 μs. All radios were tuned to the same channel. The duration of the experiment was 14.7 minutes.

The sender enqueued packets in round-robin order to both of its radio cards. On each radio card, alternating packets were sent via alternating transmit antennas by the kernel driver. Since both input queues on each radio card were always full, the channel was never idle. However, the two radios contended for the medium using the standard 802.11 method with random backoff, so the actual output sequence was not perfectly alternating. We measured the resulting interleaving pattern in the lab and found that runs of packets from the same radio card had at most 7 packets, while their mean length was 1.53 packets. This means that the probing of different antenna combinations can still happen nearly at the same times.

The ground nodes received the broadcast packets using two radio cards. Each radio card used only a single antenna to receive the packets and record the transmit timestamp, sequence number, size, and the received signal strength indication (RSSI) value. (In Section 5, we will use the collected traces from two radio cards at a time to emulate the reception over the two antennas in receive antenna selection.) Therefore, from the data traces of just one ground node, we can obtain performance of 8 different links formed by the combination of 4 UAV and 2 ground node antennas. In general, an experiment with n UAV antennas and m ground antennas would let us measure the performance of nm links at once.

The key benefit of this multiplexing scheme is that we are able to measure multiple antenna combinations in a single UAV flight. This provides a more controlled experimental environment by eliminating signal variations due to, e.g., difference in flight trajectories and aircraft attitudes in several separate UAV flights. Moreover, we can evaluate the performance of all antenna combinations in essentially the same channel conditions. The average delay between two consecutive measurements using the same transmit antenna is merely 2.4 ms, which is short enough to regard the most influential physical parameters of the environment as constant. For example, given the speed of our UAV, it would take several hundreds of milliseconds for its bank angle to change enough to appreciably affect the receiver's position in the antenna radiation pattern. Thus, we consider performance comparisons of antenna combinations, based on these collected traces, to be fair We will use these traces to evaluate performance with respect to various metrics such as signal strength and packet loss rate.

4. Trace-driven multi-antenna channel emulator

In order to evaluate antenna selection implementations quickly, we built a channel emulator driven by the pre-recorded traces. Recall that the traces are lists of entries containing the time, sequence number, size, and RSSI only of successfully received packets. For the purpose of emulation we inserted the missing entries corresponding to lost packets by interpolating the timestamps and sequence numbers of existing entries. The emulator receives packets from application processes, such as the CBR load and CPP implementation depicted in Figure 2, decides whether the packets should be dropped, and, if not, forwards them to receiving processes. Specifically, for a packet received at time T, the emulator decides using the first trace entry whose reception timestamp T_{near} is larger than T. If the trace entry denotes a successful reception, the packet is forwarded, and otherwise the entry is an interpolated entry for a missing packet and it is dropped. The times, T and T_{near}, are counted from the start of emulation and the trace, respectively. The emulator terminates once it has run for the duration of the traces.

Emulation of multi-antenna systems is achieved by using data from multiple traces. For example, to emulate the system $\{V\}\times\{V,H\}$ we would use data from single-channel traces of $\{V\}\times\{V\}$ and $\{V\}\times\{H\}$. Whenever the sender or receiver adjusts their active antenna, the emulator starts reading entries from the corresponding new trace. In the case of receive antenna selection, the emulator reads from the two traces corresponding to the active transmit antenna and the two receive antennas, and uses the entry with the highest received signal strength (RSS). Here we make the assumption that interpolated trace entries for dropped packets have a RSS of -95dBm, which is the noise floor of our radio cards.

Trace-driven emulation is deterministic in that it doesn't use a random process for drop decisions; thus, ideally multiple runs of the same emulation experiment would produce the same outcome. In practice, different runs could experience some variation due to timing of input packets, i.e., difference between T and T_{near}. This means that the input packet may get processed using different trace entries and could be subject to different decisions on whether it will be received or dropped. We evaluated this variation by comparing packet loss between measured and emulated packet streams, computed in 100ms windows. We found that 98% of these windows were less than 4% different. Furthermore, we computed the total number of packet losses of entire traces and found that the mean difference between measured and emulated was 0.05% with a 0.04 standard deviation. These differences are negligible for our purposes of evaluating packet delivery performance.

5. CPP evaluation using emulation

We used the traces from the flight experiments to evaluate transmit and receive antenna selection for 2×2 systems such as that formed by dual-antenna Atheros radio cards. From the 20 channels formed by 4 transmit and 5 receive antennas, it is possible to evaluate 60 distinct 2×2 systems ($\binom{4}{2}\cdot\binom{5}{2}$). In this section, we present the emulation results.

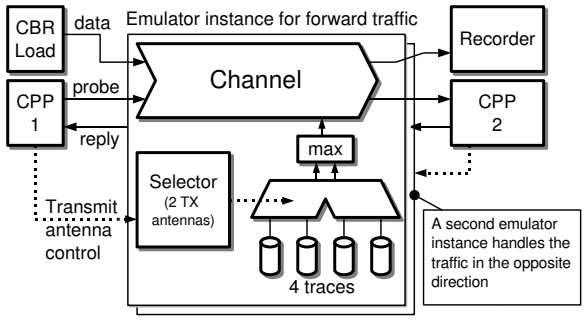

Fig. 2. Functional diagram of the emulator in a 2x2 experiment scenario. The emulator instance shown handles one direction of traffic, while an identical one handles the other direction. Applications control the transmit antenna, while the receive antenna is determined by the "max" component based on the recorded RSSI, in order to emulate hardware receive antenna selection. In field experiments, the emulator components shown here are replaced by two instances of the Linux OS using dual antenna radio cards, and the physical wireless environment as the propagation medium.

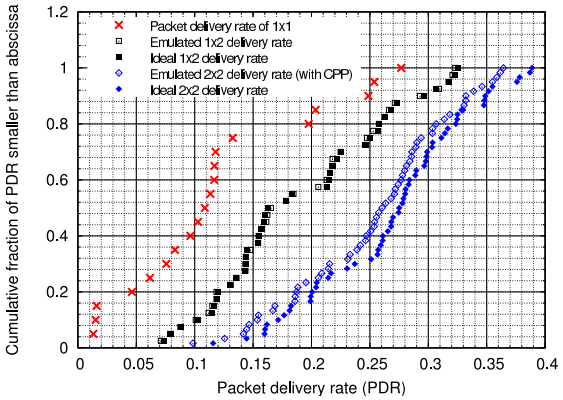

Fig. 3. Empirical cumulative distribution plots of packet delivery rate performance under emulation of 1×1, 1×2, and 2×2 systems, and ideal performance estimates for 1×2 and 2×2 systems.

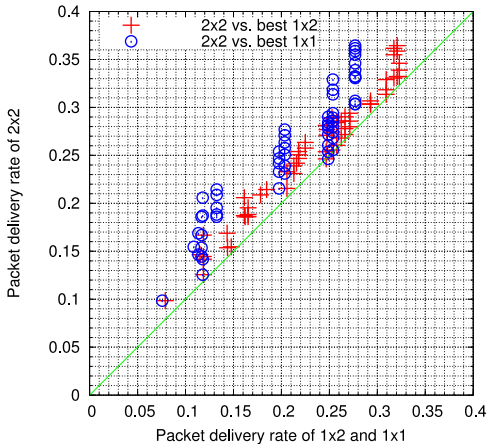

Fig. 4. This scatter-plot shows the performance of each 2×2 system plotted against that of its best 1×2 and 1×1 subsystems.

We ran CPP under emulation for each of the 60 2×2 systems, using a probing interval of 25ms. We will evaluate the effect of probing interval size in the next subsection. Along with an instance of CPP running on the sending and receiving sides, we used a CBR traffic source on the sender as the load, as depicted earlier in Figure 2. To study the performance of the 2×2 system, we analyzed the output collected in Recorder. The emulator would reselect the active trace whenever the CPP process at either endpoint toggled the active antenna.

For comparison purposes, we estimated the ideal performance achievable from 2×2 systems as follows. First, we computed the times at which the best antennas changed. Let us refer to the intervals between these times as *coherence intervals*. We then constructed a composite trace taking data from the best trace in each coherence interval. Even though this omniscient trace is not achievable in practice in that a real system would need instant channel information at the sender to achieve it, its packet delivery rate will serve as a performance upper bound for comparison purposes. We shall refer to the delivery rate as the "ideal performance estimate" hereon.

Figure 3 shows the performances of individual 1×1 channels, antenna selection in 1×2 and 2×2 systems, and the ideal performance estimates. The 2×2 system is the joint transmit/receive antenna selection system where the CPP-based transmit antenna selection augments the traditional receive antenna selection. The performances are expressed as packet delivery rates. This plot shows two main results: first, overall the 2×2 systems ("emulated" and "ideal") obtain performance higher than the 1×2 systems, indicating that channel conditions vary slowly enough for CPP feedback to be useful. Secondly, performance of joint transmit/receive antenna selection is close to the ideal performance estimate, indicating that the CPP feedback can capture much of the available gain beyond what the 1×2 systems achieve.

Figure 4 compares the performance of each 2×2 system to that of its best 1×2 and 1×1 subsystems. For example, the performance of system $\{V, H\} \times \{V, V_N\}$ would be plotted on the y-axis against the performance of the best of $V \times \{V, V_N\}$ and $H \times \{V, V_N\}$, and the best of $V \times V$, $V \times V_N$, $H \times V$ and $H \times V_N$. We see from this plot that 1) 2×2 antenna selection almost never performs worse than the sub-combinations; 2) the best 2×2 performance is 32% higher than best 1×1; 3) the best 2×2 performance is 17% higher than best 1×2. We note that in the last result, the additional gain due to transmit selection is similar to the 1×2 gain, in spite of transmit selection's inherent delay in receiving channel state information.

5.1. Effect of probing interval on transmit selection performance

In order to examine the effect of probing interval size T_P, we emulated a single 2×2 system while varying T_P from 5ms to 60s. We examined the 2×2 system with the highest packet delivery rate (0.36) and the second highest average rate of change in the best transmit antenna (1.5 per second), because it provides a high range of throughputs, and demands a high probing rate to capture the frequent best transmit antenna changes. The results are shown in Figure 5, along with the ideal performance estimate and the performance of the two 1×2 subsystems for comparison purposes.

As we can see, there is a region between 200ms and 1s where the performance of the emulated 2×2 system decreases from that near ideal to that of the 1×2, receive selection-only performance. This is a positive result in the sense that the probe interval can be as long as 200ms. Suppose for example we use a 8-byte UDP packet, which in 802.11a takes about 159μs to transmit at 6Mbps, including the DIFS interval and a minimal, 1-slot backoff. The amount of channel capacity taken up by two of these packets at 200ms intervals is negligibly low—less than 0.2% of channel capacity. Furthermore, the probing interval allowed by a larger, but still minimal channel overhead of 1% is 32ms, which is near the beginning of the high performance region.

6. CPP evaluation on a UAV testbed

In this section, we present the measurement results of CPP-based transmit antenna selection performance in the field, using a testbed that ran the unmodified software we used for the emulation in Section 5. We performed a series of flights where the UAV flew north and south along a 640-meter runway section, as depicted by the flight pattern in Figure 6 generated by an on-board GPS device. The UAV node ran the CBR packet source, transmitting to a stationary ground node. The two nodes utilized the same hardware as those in the trace collection testbed described in Section 3. The radios used channel 11 and transmitted at the 6 Mbit/s rate under 802.11g. The reason we used 802.11g instead of 802.11a as before was three-fold: 1) to obtain channel measurements at

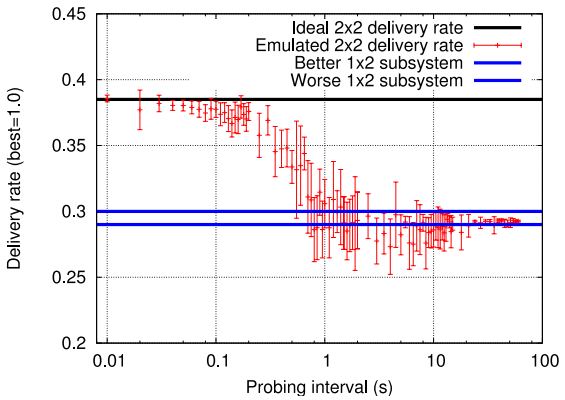

Fig. 5. Effect of probing interval on 2×2 selection performance with CPP. Each point is an average of 10 runs. The horizontal lines indicate the ideal performance estimate and the emulated performance of the individual 1×2 subsystems.

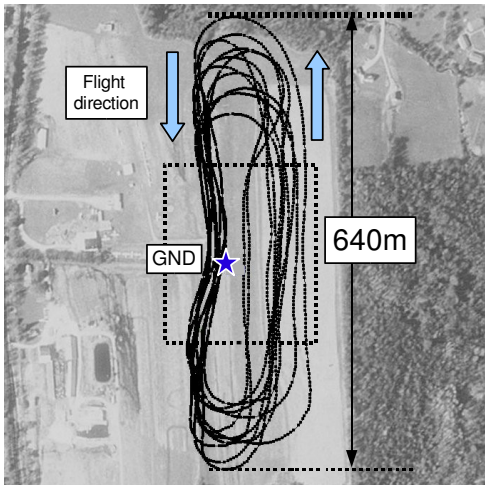

Fig. 6. The trajectory flown by the UAV during the experiment. The dotted rectangle indicates the flight data we isolated for performance measurements. The ground node is labeled "GND."

larger distances, 2) to measure antenna usage at far ends of the flight cycles, and 3) for the purposes of a co-located LCAD experiment.

Each node was equipped with two dipole antennas–one vertical (V), and the other horizontal, oriented perpendicular to the flight path (H). We chose this arrangement on the intuition that the dipole null regions would complement each other. There are possibly other antenna combinations with more diverse performance, which we might pursue in the future. However, in this work we were attracted to the simplicity of the plain dipoles.

During the flights, we tested five major system configurations: one 2×2system where the nodes ran CPP, and the other four 1×1systems where the nodes used the four possible antenna combinations, respectively. For each system configuration, we identified straight flight segments where the UAV flew either northbound or southbound, passing the ground node. These segments are more uniform than the U-turns near the ends of the flight path and therefore are better suited for a performance comparison. These segments are highlighted in Figure 6.

Table I shows performance results for the five system

	North	(stdev)	#segments
Diversity:	79.9%	(6.254)	11
V-V:	75.6%	(6.031)	4
V-H:	32.4%	(6.951)	4
H-V:	31.4%	(8.218)	9
H-H:	47.5%	(7.732)	4
	South	(stdev)	#segments
Diversity:	82.9%	(5.101)	10
V-V:	79.4%	(14.731)	3
V-H:	33.1%	(7.392)	5
H-V:	40.9%	(3.372)	7
H-H:	66.4%	(4.938)	4

TABLE I. Delivery rate performance for four single-antenna configurations and one antenna selection run.

configurations, computed from individual flight segments. The performance metric we focus on is the packet delivery rate, shown as a percentage of packets transmitted, and the sample deviation. The table also includes the number of flight segments used to compute each value.

According to these results, the performance with CPP-based transmit antenna selection, indicated by the diversity numbers in the table, is on average 5.7% better than the best static antenna pair. However, assuming normally distributed samples, the confidence of this conclusion is only 51%. One possible reason that the use of antenna diversity did not obtain a more significant gain in this particular field experiment is the relatively low altitude of the UAV. In particular, as the UAV approaches the ground node and the V-V antenna signal drops due to aligned null regions, but the distance also decreases and offsets this loss of signal. Another issue unique to these field results is that the CPP and the four individual antenna runs were measured on separate flights, introducing variations due to flight path discrepancies that were not present in the emulation data set. Thus, obtaining a more statistically significant result would require enough additional flight time to reduce the flight path variation. We plan to perform further such experiments in the future.

Figure 7 shows a more detailed comparison of the five system configurations using a plot of delivery rate versus distance. Here we can see that the performance with CPP is consistently near best of the five, whereas all of the static antenna settings have poor performing regions. We examine the CPP performance in more detail in Figure 8, which shows aggregate signal strength, delivery rate, and antenna usage data. As we can see from the antenna usage curve, the system tends to use the H-H pair at closer distances, while it spends an increasing fraction of the time on the V-V pair as the distance grows.

To further demonstrate the CPP operation, we compare its performance directly with that of V-V by plotting in Figure 9 the raw time-series performance of two individual southbound flight segments. We can clearly see the throughput drop in the middle of the V-V trace (Figure 9(a)), which occurs as the UAV flies over the ground node, bringing the V-V antennas into rough co-linear alignment. That throughput drop is absent in the CPP trace (Figure 9(b)), while there is a significant antenna usage shift toward the H-H pair.

7. Conclusion

In this paper we presented the first experimental results of joint transmit/receive antenna selection in outdoor mobile 802.11 networks using a UAV testbed. Our contributions are:

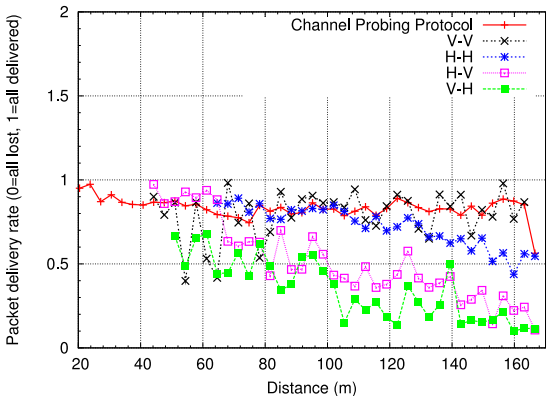

Fig. 7. A plot of delivery rate performance against distance between the UAV and ground node. We computed average delivery rates and distances over 50ms intervals, sorted them among 50 distance bins, and then plotted the averages within each distance bin.

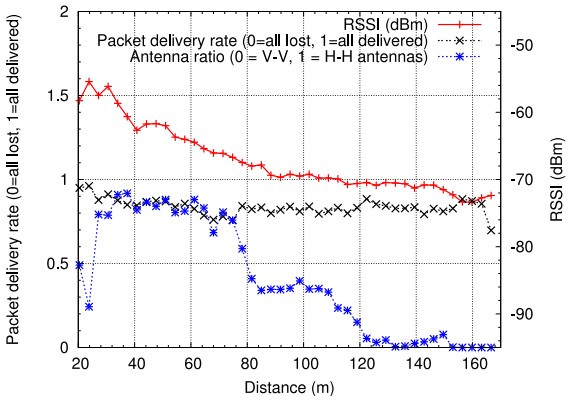

Fig. 8. Received signal strength, delivery rate, and antenna usage of CPP over distance (shown as a ratio between V-V and H-H only), computed from aggregate flight data. We computed averages for each metric over 50ms intervals, sorted them among 50 distance bins, and plotted averages from each distance bin.

- A set of high resolution, multi-channel packet loss and signal strength traces of UAV-to-ground communication.
- A trace-driven multi-antenna channel emulator.
- Emulation and experimental evaluation of antenna selection, which have shown that even with just two antennas per node, there are many antenna configurations for which transmit antenna selection achieves a performance improvement.
- A transmit antenna selection mechanism based on a link-layer feedback protocol.

The positive evaluation results prove three key points. 1) Different antennas can experience fading with enough statistical independence to let antenna selection obtain a performance gain. 2) For the UAV mobile network environment we examined, these fading conditions indeed change slowly enough for a relatively low-rate link-layer feedback protocol to adapt. 3) The link-layer feedback appears to work, i.e., the channel state information it measures translates into improved packet delivery performance, which is what the application really cares about. Based on these points, we conclude that transmit antenna selection is already a practical technique for achieving significant performance gain, even on commodity hardware and without changes to the 802.11 protocols.

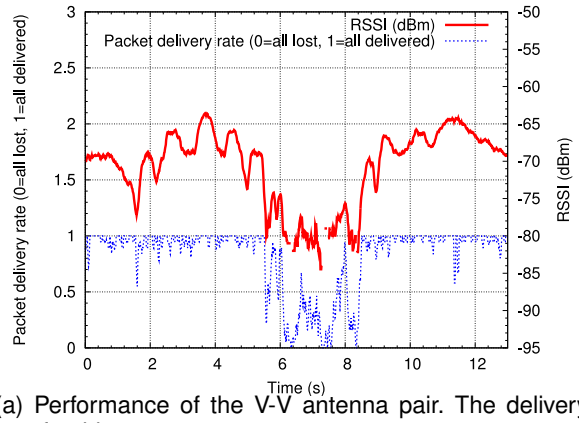

(a) Performance of the V-V antenna pair. The delivery rate for this segment was 82%.

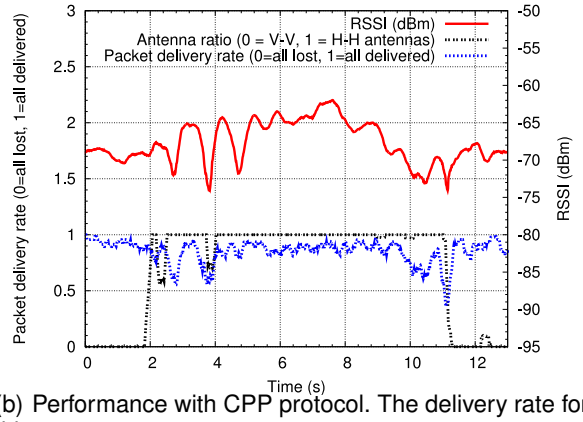

(b) Performance with CPP protocol. The delivery rate for this segment was 85%.

Fig. 9. Raw performance traces from individual southbound fly-by segments using CPP, and static V-V antennas. We selected the two fly-by segments with the closest average distance and altitude.

References

[1] "Senior Telemaster R/C Airplane by Hobby Lobby International, Inc." http://www.hobby-lobby.com/srtele.htm, 2006.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11b Mesh Network," in *SIGCOMM 2004*, Aug. 2004.

[3] Carl B. Dietrich, Jr., K. Dietze, J. R. Nealy, and W. L. Stutzman, "Spatial, Polarization, and Pattern Diversity for Wireless Handheld Terminals," *IEEE Transactions on Antennas and Propagation*, vol. 49, no. 9, pp. 1271–1281, September 2001.

[4] C.-M. Cheng, P.-H. Hsiao, H. T. Kung, and D. Vlah, "Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations," in *15th International Conference on Computer Communications and Networks (ICCCN 2006)*, Arlington, VA, USA, Oct. 2006.

[5] ——, "Maximizing Throughput of UAV-Relaying Networks with the Load-Carry-and-Deliver Paradigm," in *IEEE Wireless Communications and Networking Conference (WCNC 2007)*, Hong Kong, Mar. 2007.

[6] T. Fujisawa, J. Hasegawa, K. Tsuchie, T. Shiozawa, T. Fujita, T. Saito, and Y. Unekawa, "A Single-Chip 802.11a MAC/PHY With a 32-b RISC Processor," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 2001–2009, Nov 2003.

[7] A. Hottinen and R. Wichman, "Transmit Diversity by Antenna Selection in CDMA Downlink," in *IEEE 5th International Symposium on Spread Spectrum Techniques and Applications*, September 1998.

[8] J. G. Proakis, *Digital Communications*, 2nd ed. McGraw-Hill, 1989, pp. 721–725.

[9] S. Sanayei and A. Nosratinia, "Antenna Selection in MIMO Systems," *IEEE Communications Magazine*, vol. 42, no. 10, pp. 68–73, October 2004.

[10] A. M. D. Turkmani, A. A. Arowojolu, P. A. Jefford, and C. J. Kellet, "An Experimental Evaluation of the Performance of Two-Branch Space and Polarization Diversity Schemes at 1800 MHz," *IEEE Transactions on Vehicular Technology*, vol. 44, pp. 318–326, March 1995.

Hindawi Publishing Corporation
EURASIP Journal on Advances in Signal Processing
Volume 2007, Article ID 13659, 8 pages
doi:10.1155/2007/13659

Research Article

A Hub Matrix Theory and Applications to Wireless Communications

H. T. Kung[1] and B. W. Suter[1, 2]

[1] *Harvard School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA*
[2] *US Air Force Research Laboratory, Rome, NY 13440, USA*

Received 24 July 2006; Accepted 22 January 2007

Recommended by Sharon Gannot

This paper considers communications and network systems whose properties are characterized by the gaps of the leading eigenvalues of $A^H A$ for a matrix A. It is shown that a sufficient and necessary condition for a large eigen-gap is that A is a "hub" matrix in the sense that it has dominant columns. Some applications of this hub theory in multiple-input and multiple-output (MIMO) wireless systems are presented.

1. INTRODUCTION

There are many communications and network systems whose properties are characterized by the eigenstructure of a matrix of the form $A^H A$, also known as the Gram matrix of **A**, where **A** is a matrix with real or complex entries. For example, for a communications system, **A** could be a channel matrix, usually denoted **H**. The capacity of such system is related to the eigenvalues of $H^H H$ [1]. In the area of web page ranking, with entries of **A** representing hyperlinks, Kleinberg [2] shows that eigenvectors corresponding to the largest eigenvalues of $A^T A$ give the rankings of the most useful (authority) or popular (hub) web pages. Using a reputation system that parallels Kleinberg's work, Kung and Wu [3] developed an eigenvector-based peer-to-peer (P2P) network user reputation ranking in order to provide services to P2P users based on past contributions (reputation) to avoid "freeloaders." Furthermore, the rate of convergence in the iterative computation of reputations is determined by the gap of the leading two eigenvalues of $A^H A$.

The recognition that the eigenstructure of $A^H A$ determines the properties of these communications and network systems motivates the work of this paper. We will develop a theoretical framework, called a hub matrix theory, which allows us to predict the eigenstructure of $A^H A$ by examining **A** directly. We will prove sufficient and necessary conditions for the existence of a large gap between the largest and the second largest eigenvalues of $A^H A$. Finally, we apply the "hub" theory and our mathematical results to multiple-input and multiple-output (MIMO) wireless systems.

2. HUB MATRIX THEORY

It is instructive to conduct a thought experiment on a computation process before we introduce our hub matrix theory. The process iteratively computes the values for a set of variables, which for example could be beamforming weights in a beamforming communication system. Figure 1 depicts an example of this process: variable X uses and contributes to variables U_2 and U_4, variable Y uses and contributes to variables U_3 and U_5, and variable Z uses and contributes to all variables $U_1, \ldots, U_6$. We say variable Z is a "hub" in the sense that variables involved in Z's computation constitute a superset of those involved in the computation of any other variable. The dominance is illustrated graphically in Figure 1.

We can describe the computation process in matrix notation. Let

$$A = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1)$$

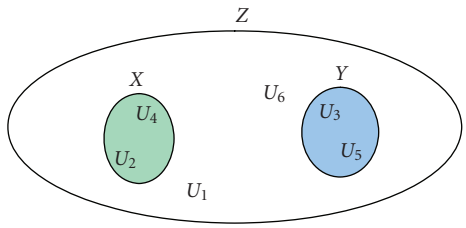

Figure 1: Graphical representation of hub concept.

This process performs two steps alternatively (cf. Figure 1).

(1) X, Y, and Z contribute to variables in their respective regions.
(2) X, Y, and Z compute their values using variables in their respective regions.

The first step (1) is $(U_1, U_2, \ldots, U_6)^T \leftarrow A^*(X, Y, Z)^T$ and next step (2) is $(X, Y, Z)^T \leftarrow A^{T*}(U_1, U_2, \ldots, U_6)^T$. Thus, the computational process performs the iteration $(X, Y, Z)^T \leftarrow S^*(X, Y, Z)^T$, where S is defined as follows:

$$S = A^T A = \begin{pmatrix} 2 & 0 & 2 \\ 0 & 2 & 2 \\ 2 & 2 & 6 \end{pmatrix}. \tag{2}$$

Note that an arrowhead matrix S, as defined below, has emerged. Furthermore, note that matrix A exhibits the hub property of Z in Figure 1 in view of the fact that the last column of A consists of all 1's, whereas other columns consist of only a few 1's.

Definition 1 (arrowhead matrix). Let $S \in \mathbf{C}^{m\times m}$ be a given Hermitian matrix. S is called an *arrowhead matrix* if

$$S = \begin{pmatrix} D & c \\ c^H & b \end{pmatrix}, \tag{3}$$

where $D = \operatorname{diag}(d^{(1)}, \ldots, d^{(m-1)}) \in \mathbf{R}^{(m-1)\times(m-1)}$ is a real diagonal matrix, $c = (c^{(1)}, \ldots, c^{(m-1)}) \in \mathbf{C}^{m-1}$ is a complex vector, and $b \in \mathbf{R}$ is a real number.

The eigenvalues of an arbitrary square matrix are invariant under similarity transformations. Therefore, we can with no loss of generality arrange the diagonal elements of D to be ordered so that $d^{(i)} \le d^{(i+1)}$ for $i = 1, \ldots, m-2$. For details concerning arrowhead matrices, see for example [4].

Definition 2 (hub matrix). A matrix $A \in \mathbf{C}^{n\times m}$ is called a *candidate-hub matrix*, if $m-1$ of its columns are orthogonal to each other with respect to the Euclidean inner product. If in addition the remaining column has its Euclidean norm greater than or equal to that of any other column, then the matrix A is called a *hub matrix* and this remaining column is called the *hub column*. We are normally interested in hub matrices where the hub column has much large magnitude than other columns. (As we show later in Theorems 4 and 10 that in this case the corresponding arrowhead matrices will have large eigengaps).

In this paper, we study the eigenvalues of $S = A^H A$, where A is a hub matrix. Since the eigenvalues of S are invariant under similarity transformations of S, we can permute the columns of the hub matrix A so that its last column is the hub column without loss of generality. For the rest of this paper, we will denote the columns of a hub matrix A by $a_1, \ldots, a_m$, and assume that columns $a_1, \ldots, a_{m-1}$ are orthogonal to each other, that is, $a_i^H a_j = 0$ for $i \ne j$ and $i, j = 1, \ldots, m-1$, and column a_m is the hub column. The matrix A introduced in the context of the graphical model from Figure 1 is such a hub matrix.

In Section 4, we will relax the orthogonality condition of a hub matrix, by introducing the notion of hub and arrowhead dominant matrices.

Theorem 1. *Let $\mathbf{A} \in \mathbf{C}^{n\times m}$ and let $S \in \mathbf{C}^{m\times m}$ be the Gram matrix of $\mathbf{A}$ that is, $S = A^H A$. $\mathbf{S}$ is an arrowhead matrix if and only if $\mathbf{A}$ is a candidate-hub matrix.*

Proof. Suppose A is a candidate-hub matrix. Since $S = A^H A$, the entries of S are $s^{(i,j)} = a_i^H a_j$ for $i, j = 1, \ldots, m$. By Definition 2 of a candidate-hub matrix, the nonhub columns of A are orthogonal, that is, $a_i^H a_j = 0$ for $i \ne j$ and $i, j = 1, \ldots, m-1$. Since S is Hermitian, the transpose of the last column is the complex conjugate of the last row and the diagonal elements of S are real numbers. Therefore, $S = A^H A$ is an arrowhead matrix by Definition 1.

Suppose $S = A^H A$ is an arrowhead matrix. Note that the components of the S matrix of Definition 1 can be represented in terms of the inner products of columns of A, that is, $b = a_m^H a_m$, $d^{(i)} = a_i^H a_i$, $c^{(i)} = a_i^H a_m$ for $i = 1, \ldots, m-1$. Since S is an arrowhead matrix, all other off-diagonal entries of S, $s^{(i,j)} = a_i^H a_j$ for $i \ne j$ and $i, j = 1, \ldots, m-1$, are zero. Thus, $a_i^H a_j = 0$ if $i \ne j$ and $i, j = 1, \ldots, m-1$. So, A is a candidate-hub matrix by Definition 2. □

Before proving our main result in Theorem 4, we first restate some well-known results which will be needed for the proof.

Theorem 2 (interlacing eigenvalues theorem for bordered matrices). *Let $U \in \mathbf{C}^{(m-1)\times(m-1)}$ be a given Hermitian matrix, let $y \in \mathbf{C}^{(m-1)}$ be a given vector, and let $a \in \mathbf{R}$ be a given real number. Let $V \in \mathbf{C}^{m\times m}$ be the Hermitian matrix obtained by bordering U with y and a as follows:*

$$V = \begin{pmatrix} U & y \\ y^H & a \end{pmatrix}. \tag{4}$$

Let the eigenvalues of V and U be denoted by $\{\lambda_i\}$ and $\{\mu_i\}$, respectively, and assume that they have been arranged in increasing order, that is, $\lambda_1 \le \cdots \le \lambda_m$ and $\mu_1 \le \cdots \le \mu_{m-1}$. Then

$$\lambda_1 \le \mu_1 \le \lambda_2 \le \cdots \le \lambda_{m-1} \le \mu_{m-1} \le \lambda_m. \tag{5}$$

Proof. See [5, page 189]. □

Definition 3 (majorizing vectors). Let $\boldsymbol{\alpha} \in \mathbf{R}^m$ and $\boldsymbol{\beta} \in \mathbf{R}^m$ be given vectors. If we arrange the entries of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ in

increasing order, that is, $\alpha^{(1)} \le \cdots \le \alpha^{(m)}$ and $\beta^{(1)} \le \cdots \le \beta^{(m)}$, then vector $\boldsymbol{\beta}$ is said to majorize vector $\boldsymbol{\alpha}$ if

$$\sum_{i=1}^{k} \beta^{(i)} \ge \sum_{i=1}^{k} \alpha^{(i)} \quad \text{for } k = 1, \ldots, m \tag{6}$$

with equality for $k = m$.

For details concerning majorizing vectors, see [5, pages 192–198]. The following theorem provides an important property expressed in terms of vector majorizing.

Theorem 3 (Schur-Horn theorem). *Let $V \in \mathbf{C}^{m\times m}$ be Hermitian. The vector of diagonal entries of V majorizes the vector of eigenvalues of V.*

Proof. See [5, page 193]. □

Definition 4 (hub-gap). Let $A \in \mathbf{C}^{n\times m}$ be a matrix with its columns denoted by $a_1, \ldots, a_m$ with $0 < \|a_1\|_2^2 \le \cdots \le \|a_m\|_2^2$. For $i = 1, \ldots, m-1$, the ith hub-gap of A is defined to be

$$\text{HubGap}_i(A) = \frac{\|a_{m-(i-1)}\|_2^2}{\|a_{m-i}\|_2^2}. \tag{7}$$

Definition 5 (eigengap). Let $S \in \mathbf{C}^{m\times m}$ be a Hermitian matrix with its real eigenvalues denoted by $\lambda_1, \ldots, \lambda_m$ with $\lambda_1 \le \cdots \le \lambda_m$. For $i = 1, \ldots, m-1$, the ith eigengap of S is defined to be

$$\text{EigenGap}_i(S) = \frac{\lambda_{m-(i-1)}}{\lambda_{m-i}}. \tag{8}$$

Theorem 4. *Let $A \in \mathbf{C}^{n\times m}$ be a hub matrix with its columns denoted by $a_1, \ldots, a_m$ and $0 < \|a_1\|_2^2 \le \cdots \le \|a_m\|_2^2$. Let $S = A^H A \in \mathbf{C}^{m\times m}$ be the corresponding arrowhead matrix with its eigenvalues denoted by $\lambda_1, \ldots, \lambda_m$ with $0 \le \lambda_1 \le \cdots \le \lambda_m$. Then*

$$\begin{aligned} \text{HubGap}_1(A) &\le \text{EigenGap}_1(S) \\ &\le \left(\text{HubGap}_1(A) + 1\right) \text{HubGap}_2(A). \end{aligned} \tag{9}$$

Proof. Let T be the matrix formed from S by deleting its last row and column. This means that T is a diagonal matrix with diagonal elements $\|a_i\|_2^2$ for $i = 1, \ldots, m-1$. By Theorem 2, the eigenvalues of S interlace those of T, that is, $\lambda_1 \le \|a_1\|_2^2 \le \cdots \le \lambda_{m-1} \le \|a_{m-1}\|_2^2 \le \lambda_m$. Thus, λ_{m-1} is a lower bound for $\|a_{m-1}\|_2^2$. By Theorem 3, the vector of diagonal values of S majorizes the vector of eigenvalues of S, that is, $\sum_{i=1}^{k} d^{(i)} \ge \sum_{i=1}^{k} \lambda_i$ for $k = 1, \ldots, m-1$ and $\sum_{i=1}^{m-1} d^{(i)} + b = \sum_{i=1}^{m} \lambda_m$. So, $b \le \lambda_m$. Since $b = \|a_m\|_2^2$, λ_m is an upper bound for $\|a_m\|_2^2$. Hence, $\|a_m\|_2^2/\|a_{m-1}\|_2^2 \le \lambda_m/\lambda_{m-1}$ or $\text{HubGap}_1(A) \le \text{EigenGap}_1(S)$.

Again, by using Theorems 2 and 3, we have $\sum_{i=1}^{m-1} d^{(i)} + b = \sum_{i=1}^{m} \lambda_m$ and $\lambda_1 \le d^{(1)} \le \lambda_2 \le d^{(2)} \le \lambda_3 \le \cdots \le d^{(m-2)} \le \lambda_{m-1} \le d^{(m-1)} \le \lambda_m$, and, as such,

$$\begin{aligned} &\left(d^{(1)} + \cdots + d^{(m-2)}\right) + d^{(m-1)} + b \\ &\quad = \lambda_1 + \left(\lambda_2 + \cdots + \lambda_{m-1}\right) + \lambda_m \\ &\quad \ge \lambda_1 + \left(d^{(1)} + \cdots + d^{(m-2)}\right) + \lambda_m. \end{aligned} \tag{10}$$

This result implies that $d^{(m-1)} + b \ge \lambda_1 + \lambda_m \ge \lambda_m$. By noting that $d^{(m-2)} \le \lambda_{m-1}$, we have

$$\begin{aligned} \text{EigenGap}_1(S) = \frac{\lambda_m}{\lambda_{m-1}} &\le \frac{d^{(m-1)} + b}{d^{(m-2)}} = \frac{\|a_{m-1}\|_2^2 + \|a_m\|_2^2}{\|a_{m-2}\|_2^2} \\ &= \frac{\|a_{m-1}\|_2^2}{\|a_{m-2}\|_2^2} + \frac{\|a_m\|_2^2}{\|a_{m-1}\|_2^2} \cdot \frac{\|a_{m-1}\|_2^2}{\|a_{m-2}\|_2^2} \\ &= \left(\text{HubGap}_1(A) + 1\right) \cdot \text{HubGap}_2(A). \end{aligned} \tag{11}$$

□

By Theorem 4, we have the following result, where notation "$\gg$" means "much larger than."

Corollary 1. *Let $A \in \mathbf{C}^{n\times m}$ be a matrix with its columns $a_1, \ldots, a_m$ satisfying $0 < \|a_1\|_2^2 \le \cdots \le \|a_{m-1}\|_2^2 \le \|a_m\|_2^2$. Let $S = A^H A \in \mathbf{C}^{m\times m}$ with its eigenvalues $\lambda_1, \cdots, \lambda_m$ satisfying $0 \le \lambda_1 \le \cdots \le \lambda_m$. The following holds*

(1) *if A is a hub matrix with $\|a_m\|_2 \gg \|a_{m-1}\|_2$, then S is an arrowhead matrix with $\lambda_m \gg \lambda_{m-1}$; and*
(2) *if S is an arrowhead matrix with $\lambda_m \gg \lambda_{m-1}$, then A is a hub matrix with $\|a_m\|_2 \gg \|a_{m-1}\|_2$ or $\|a_{m-1}\|_2 \gg \|a_{m-2}\|_2$ or both.*

3. MIMO COMMUNICATIONS APPLICATION

A multiple-input multiple-output (MIMO) system with M_t transmit antennas and M_r receive antennas is depicted in Figure 2 [6, 7]. Assume the MIMO channel is modeled by the $M_r \times M_t$ channel propagation matrix $H = (h_{ij})$. The input-output relationship, given a transmitted symbol s, for this system is given by

$$x = sz^H Hw + z^H n. \tag{12}$$

The vectors w and z in the equation are called the beamforming and combining vectors, respectively, which will be chosen to maximize the signal-to-noise ratio (SNR). We will model the noise vector n as having entries, which are independent and identically distributed (i.i.d.) random variables of complex Gaussian distribution $CN(0, 1)$. Without loss of generality, assume the average power of transmit signal equals one, that is, $E|s|^2 = 1$. For the beamforming system described here, the signal to noise ratio, γ, after combining at the receiver is given by

$$\gamma = \frac{|z^H Hw|^2}{\|z\|_2^2}. \tag{13}$$

Without loss of generality, assume $\|z\|_2 = 1$. With this assumption, the SNR becomes

$$\gamma = |z^H Hw|^2. \tag{14}$$

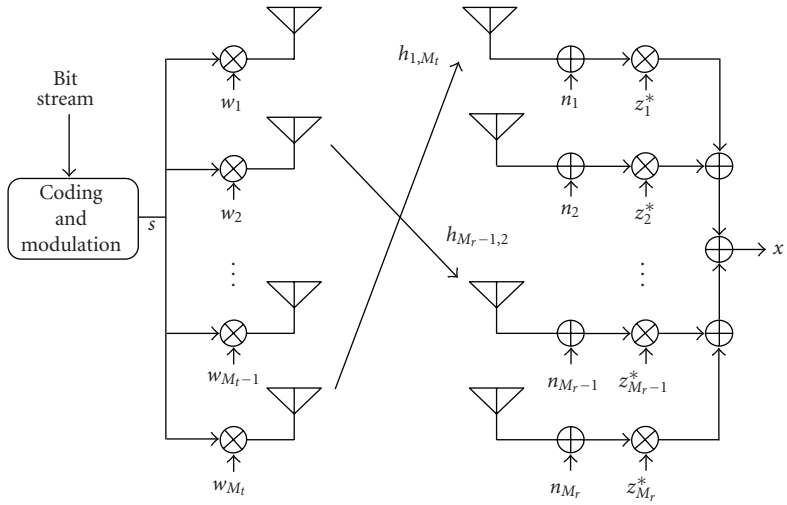

Figure 2: MIMO block diagram (see [6, datapath portion of Figure 1]).

3.1. Maximum ratio combining

A receiver where z maximizes γ for a given w is known as a maximum ratio combining (MRC) receiver in the literature. By the Cauchy-Bunyakovskii-Schwartz inequality (see, e.g., [8, page 272]), we have

$$\left|z^H Hw\right|^2 \leq \|z\|_2^2 \|Hw\|_2^2. \tag{15}$$

Since we already assume $\|z\|_2 = 1$,

$$\left|z^H Hw\right|^2 \leq \|Hw\|_2^2. \tag{16}$$

Moreover, since in MRC we desire to maximize the SNR, we must choose z to be

$$z^{\text{MRC}} = \frac{Hw}{\|Hw\|_2}, \tag{17}$$

which implies that the SNR for MRC is

$$\gamma^{\text{MRC}} = \|Hw\|_2^2. \tag{18}$$

3.2. Selection diversity transmission, generalized subset selection, and combined SDT/MRC and GSS/MRC

For a selection diversity transmission (SDT) [9] system, only the antenna that yields the largest SNR is selected for transmission at any instant of time. This means

$$w = \left[\delta_{1,f(1)}, \ldots, \delta_{M_t,f(1)}\right]^T, \tag{19}$$

where the Kronecker impulse $\delta_{i,j}$ is defined as $\delta_{i,j} = 1$ if $i = j$, and $\delta_{i,j} = 0$ if $i \neq j$, and $f(1)$ represents the value of the index x that maximizes $\sum_i |h_{i,x}|^2$. Thus, the SNR for the combined SDT/MRC communications system is

$$\gamma^{\text{SDT/MRC}} = \left|\left|h_{f(1)}\right|\right|_2^2. \tag{20}$$

By definition, a generalized subset selection (GSS) [10] system powers those k transmitters which yield the top k SNR values at the receiver for some $k > 1$. That is, if $f(1), f(2), \ldots, f(k)$ stand for the indices of these transmitters, then $w_{f(i)} = 1/\sqrt{k}$ for $i = 1, \ldots, k$, and all other entries of w are zero. It follows that, for the combined GSS/MRC communications system, the SNR gain is given by

$$\gamma^{\text{GSS/MRC}} = \frac{1}{k}\left|\left|\sum_{i=1}^{k} h_{f(i)}\right|\right|_2^2. \tag{21}$$

In the limiting case when $k = M_t$, GSS becomes equal gain transmission (EGT) [6, 7], which requires all M_t transmitters to be equally powered, that is, $w_{f(i)} = 1/\sqrt{M_t}$ for $i = 1, \ldots, M_t$. Then, for the combined EGT/MRC communications system, the SNR gain takes the expression

$$\gamma^{\text{EGT/MRC}} = \frac{1}{M_t}\left|\left|\sum_{i=1}^{M_t} h_{f(i)}\right|\right|_2^2. \tag{22}$$

3.3. Maximum ratio transmission and combined MRT/MRC

Suppose there are no constraints placed on the form of the vector w. Let us reexamine the expression of SNR gain γ^{MRC}. Note

$$\gamma^{\text{MRC}} = \|Hw\|_2^2 = (Hw)^H(Hw) = w^H\left(H^H Hw\right). \tag{23}$$

With the assumption that $\|w\|_2 = 1$, the above equation is maximized under maximum ratio transmission (MRT) [9] (see, e.g., [5, page 295]), that is, when

$$w = w_m, \tag{24}$$

where w_m is the normalized eigenvector corresponding to the largest eigenvalues λ_m of $H^H H$. Thus, for an MRT/MRC system, we have

$$\gamma^{\text{MRT/MRC}} = \lambda_m. \tag{25}$$

3.4. Performance comparison between SDT/MRC and MRT/MRC

Theorem 5. *Let $H \in \mathbf{C}^{n\times m}$ be a hub matrix with its columns denoted by $h_1,\ldots,h_m$ and $0 < \|h_1\|_2^2 \le \cdots \le \|h_{m-1}\|_2^2 \le \|h_m\|_2^2$. Let $\gamma^{\text{SDT/MRC}}$ and $\gamma^{\text{MRT/MRC}}$ be the SNR gains for SDT/MRC and MRT/MRC, respectively. Then*

$$\frac{\text{HubGap}_1(H)}{\text{HubGap}_1(H)+1} \le \frac{\gamma^{\text{SDT/MRC}}}{\gamma^{\text{MRT/MRC}}} \le 1. \tag{26}$$

Proof. We note that the A matrix in hub matrix theory of Section 2 corresponds to the H matrix here, and the a_i column of A corresponds to the h_i column of H for $i = 1,\ldots,m$. From the proof of Theorem 4, we note $b = \|a_m\|_2^2 \le \lambda_m$ or $\|h_m\|_2^2 \le \lambda_m$. It follows that

$$\frac{\gamma^{\text{SDT/MRC}}}{\gamma^{\text{MRT/MRC}}} \le 1. \tag{27}$$

To derive a lower bound for $\gamma^{\text{SDT/MRC}}/\gamma^{\text{MRT/MRC}}$, we note from the proof of Theorem 4 that $\lambda_m \le d^{(m-1)} + b$. This means that

$$\gamma^{\text{MRT/MRC}} \le ||a_{m-1}||_2^2 + ||a_m||_2^2 = ||h_{m-1}||_2^2 + ||h_m||_2^2. \tag{28}$$

Thus

$$\frac{\gamma^{\text{SDT/MRC}}}{\gamma^{\text{MRT/MRC}}} \ge \frac{||h_m||_2^2}{||h_{m-1}||_2^2 + ||h_m||_2^2} = \frac{\text{HubGap}_1(H)}{\text{HubGap}_1(H)+1}. \tag{29}$$

□

The inequality $\gamma^{\text{SDT/MRC}}/\gamma^{\text{MRT/MRC}} \le 1$ in Theorem 5 reflects the fact that in the SDT/MRC system, w is chosen to be a particular unit vector rather than an optimal choice. The other inequality of Theorem 5, $\text{HubGap}_1(H)/(\text{HubGap}_1(H)+1) \le \gamma^{\text{SDT/MRC}}/\gamma^{\text{MRT/MRC}}$, implies that the SNR for SDT/MRC approaches that for MRT/MRC when H is a hub matrix with a dominant hub column. More precisely, we have the following result.

Corollary 2. *Let $H \in \mathbf{C}^{n\times m}$ be a hub matrix with its columns denoted by $h_1,\ldots,h_m$ and $0 < \|h_1\|_2^2 \le \cdots \le \|h_m\|_2^2$. Let $\gamma^{\text{SDT/MRC}}$ and $\gamma^{\text{MRT/MRC}}$ be the SNR for SDT/MRC and MRT/MRC, respectively. Then, as $\text{HubGap}_1(H)$ increases, $\gamma^{\text{MRT/MRC}}/\gamma^{\text{SDT/MRC}}$ approaches one at a rate of at least $\text{HubGap}_1(H)/(\text{HubGap}_1(H)+1)$.*

3.5. GSS-MRT/MRC and performance comparison with MRT/MRC

Using an analysis similar to the one above, we can derive performance bounds for a recently discovered communication system that incorporates antenna selection with MRT on the transmission side while applying MRC on the receiver side [11, 12]. This approach will be called GSS-MRT/MRC here. Given a GSS scheme that powers those k transmitters which yield the top k highest SNR values, a GSS-MRT/MRC system is defined to be an MRT/MRC system applied to these k transmitters. Let $f(1), f(2),\ldots,f(k)$ be the indices of these k transmitters, and $\tilde{H}$ the matrix formed by columns $h_{f(i)}$ of H for $i = 1,\ldots,k$. It is easy to see that the SNR for GSS-MRT/MRC is

$$\gamma^{\text{GSS-MRT/MRC}} = \tilde{\lambda}_m, \tag{30}$$

where $\tilde{\lambda}_m$ is the largest eigenvalue of $\tilde{H}^H\tilde{H}$.

Theorem 6. *Let $H \in \mathbf{C}^{n\times m}$ be a hub matrix with its columns denoted by $h_1,\ldots,h_m$ and $0 < \|h_1\|_2^2 \le \cdots \le \|h_{m-1}\|_2^2 \le \|h_m\|_2^2$. Let $\gamma^{\text{GSS-MRT/MRC}}$ and $\gamma^{\text{MRT/MRC}}$ be the SNR values for GSS-MRT/MRC and MRT/MRC, respectively. Then*

$$\frac{\text{HubGap}_1(H)}{\text{HubGap}_1(H)+1} \le \frac{\gamma^{\text{GSS-MRT/MRC}}}{\gamma^{\text{MRT/MRC}}} \le \frac{\text{HubGap}_1(H)+1}{\text{HubGap}_1(H)}. \tag{31}$$

Proof. Since $0 < \|h_1\|_2^2 \le \cdots \le \|h_{m-1}\|_2^2 \le \|h_m\|_2^2$, $\tilde{H}$ consists of the last k columns of H. Moreover, since H is a hub matrix, so is $\tilde{H}$. From the proof of Theorem 4, we note both λ_m and $\tilde{\lambda}_m$ are bounded above by $\|h_{m-1}\|_2^2 + \|h_m\|_2^2$ and below by $\|h_m\|_2^2$. It follows that

$$\frac{\text{HubGap}_1(H)}{\text{HubGap}_1(H)+1} = \frac{||h_m||_2^2}{||h_{m-1}||_2^2 + ||h_m||_2^2} \le \frac{\gamma^{\text{GSS-MRT/MRC}}}{\gamma^{\text{MRT/MRC}}} = \frac{\tilde{\lambda}_m}{\lambda_m} \le \frac{||h_{m-1}||_2^2 + ||h_m||_2^2}{||h_m||_2^2} = \frac{\text{HubGap}_1(H)+1}{\text{HubGap}_1(H)}. \tag{32}$$

□

3.6. Diversity selection with partitions, DSP-MRT/MRC, and performance bounds

Suppose that transmitters are partitioned into multiple transmission partitions. We define the diversity selection with partitions (DSP) to be the transmission scheme where in each transmission partition only the transmitter with the largest SNR will be powered. Note that SDT discussed above is a special case of DSP when there is only one partition consisting of all transmitters.

Let k be the number of partitions, and $f(1), f(2), \ldots, f(k)$ the indices of the powered transmitters. A DSP-MRT/MRC system is defined to be an MRT/MRC system applied to these k transmitters. Define $\hat{H}$ to be the matrix formed by columns $h_{f(i)}$ of H for $i = 1,\ldots,k$. Then the SNR for DSP-MRT/MRC is

$$\gamma^{\text{DSPS-MRT/MRC}} = \hat{\lambda}_m, \tag{33}$$

where $\hat{\lambda}_m$ is the largest eigenvalue of $\hat{H}^H\hat{H}$.

Note that in general the powered transmitters for DSP are not the same as those for GSS. This is because a transmitter that yields the highest SNR among transmitters in one of the k partitions may not be among the transmitters that yield the top k highest SNR values among all transmitters. Nevertheless, when H is a hub matrix with

$0 < \|h_1\|_2^2 \le \cdots \le \|h_{m-1}\|_2^2 \le \|h_m\|_2^2$, we can bound $\widehat{\lambda}_m$ for DSP-MRT/MRC in a manner similar to how we bound $\widetilde{\lambda}_m$ for GSS-MRT/MRC. That is, for DSP-MRT/MRC, $\widehat{\lambda}_m$ is bounded above by $\|h_k\|_2^2 + \|h_m\|_2^2$ and below by $\|h_m\|_2^2$, where h_k is the second largest column of $\widehat{H}$ in magnitude. Note that $\|h_k\|_2^2 \le \|h_{m-1}\|_2^2$, since the second largest column of $\widehat{H}$ in magnitude cannot be larger that than of H. We have the following result similar to that of Theorem 6.

Theorem 7. *Let $H \in \mathbf{C}^{n\times m}$ be a hub matrix with its columns denoted by $h_1, \ldots, h_m$ and $0 < \|h_1\|_2^2 \le \cdots \le \|h_{m-1}\|_2^2 \le \|h_m\|_2^2$. Let $\gamma^{\mathrm{DSP-MRT/MRC}}$ and $\gamma^{\mathrm{MRT/MRC}}$ be the SNR for DSP-MRT/MRC and MRT/MRC, respectively. Then*

$$\frac{\mathrm{HubGap}_1(H)}{\mathrm{HubGap}_1(H)+1} \le \frac{\gamma^{DSP-\mathrm{MRT/MRC}}}{\gamma^{\mathrm{MRT/MRC}}} \le \frac{\mathrm{HubGap}_1(H)+1}{\mathrm{HubGap}_1(H)}. \tag{34}$$

Theorems 6 and 7 imply that when $\mathrm{HubGap}_1(H)$ becomes large, the SNR values of both GSS-MRT/MRC and DSP-MRT/MRC approach that of MRT/MRC.

4. HUB DOMINANT MATRIX THEORY

We generalize the hub matrix theory presented above to situations when matrix A (or H) exhibits a "near" hub property. In order to relax the definition of orthogonality of a set of vectors, we use the notion of frame.

Definition 6 (frame). A set of distinct vectors $\{f_1, \ldots, f_n\}$ is said to be a frame if there exist positive constants ξ and ϑ called frame bounds such that

$$\xi \|f_j\|^2 \le \sum_{i=1}^{n} |f_i^H f_j| \le \vartheta \|f_j\|^2 \quad \text{for } j = 1, \ldots, n. \tag{35}$$

Note that if $\xi = \vartheta = 1$, then the set of vectors $\{f_1, \ldots, f_n\}$ is orthogonal. Here we use frames to bound the nonorthogonality of a collection of vectors, while the usual use for frames is to quantify the redundancy in a representation (see, e.g., [13]).

Definition 7 (hub dominant matrix). A matrix $A \in \mathbf{C}^{n\times m}$ is called a *candidate-hub-dominant matrix* if $m-1$ of its columns form a frame with frame bounds $\xi = 1$ and $\vartheta = 2$, that is, $\|a_j\|^2 \le \sum_{i=1}^{m-1} |a_i^H a_j| \le 2\|a_j\|^2$ for $j = 1, \ldots, m-1$. If in addition the remaining column has its Euclidean norm greater than or equal to that of any other column, then the matrix A is called a *hub-dominant matrix* and the remaining column is called the *hub column*.

We next generalize the definition of arrowhead matrix to arrowhead dominant matrix, where the matrix D in Definition 1 goes from being a diagonal matrix to a diagonally dominant matrix.

Definition 8 (diagonally dominant matrix). Let $E \in \mathbf{C}^{m\times m}$ be a given Hermitian matrix. E is said to be diagonally dominant if for each row the magnitude of the diagonal entry is greater than or equal to the row sum of magnitudes of all off-diagonal entries, that is,

$$|e^{(i,i)}| \ge \sum_{\substack{j=1 \\ j\ne i}}^{m-1} |e^{(i,j)}| \quad \text{for } i = 1, \ldots, m. \tag{36}$$

For more information on diagonally dominant matrices, see for example [5, page 349].

Definition 9 (arrowhead dominant matrix). Let $S \in \mathbf{C}^{m\times m}$ be a given Hermitian matrix. S is called an *arrowhead dominant matrix* if

$$S = \begin{pmatrix} D & c \\ c^H & b \end{pmatrix}, \tag{37}$$

where $D \in \mathbf{C}^{(m-1)\times(m-1)}$ is a diagonally dominant matrix, $c = (c^{(1)}, \ldots, c^{(m-1)}) \in \mathbf{C}^{m-1}$ is a complex vector, and $b \in \mathbf{R}$ is a real number.

Similar to Theorem 1, we have the following theorem.

Theorem 8. *Let $A \in \mathbf{C}^{n\times m}$ and let $S \in \mathbf{C}^{m\times m}$ be the Gram matrix of A, that is, $S = A^H A$. S is an arrowhead dominant matrix if and only if A is a candidate-hub-dominant matrix.*

Proof. Suppose A is a candidate-hub-dominant matrix. Since $S = A^H A$, the entries of S can be expressed as $s^{(i,j)} = a_i^H a_j$ for $i, j = 1, \ldots, m$. By Definition 7 of a hub-dominant matrix, the nonhub columns of A form a frame with frame bounds $\xi = 1$ and $\vartheta = 2$, that is $\|a_j\|^2 \le \sum_{i=1}^{m-1} |a_i^H a_j| \le 2\|a_j\|^2$ for $j = 1, \ldots, m-1$. Since $\|a_j\|^2 = |a_j^H a_j|$, it follows that $|a_i^H a_i| \ge \sum_{j=1, j\ne i}^{m-1} |a_i^H a_j|$, $i = 1, \ldots, m-1$, which is the diagonal dominance condition on the sub-matrix D of S. Since S is Hermitian, the transpose of the last column is the complex conjugate of the last row and the diagonal elements of S are real numbers. Therefore, $S = A^H A$ is an arrowhead dominant matrix in accordance with Definition 9.

Suppose $S = A^H A$ is an arrowhead dominant matrix. Note that the components of the S matrix of Definition 9 can be represented in terms of the columns of A. Thus $b = a_m^H a_m$ and $c^{(i)} = a_i^H a_m$ for $i = 1, \ldots, m-1$. Since $|a_j^H a_j| = \|a_j\|^2$, the diagonal dominance condition, $|a_i^H a_i| \ge \sum_{j=1, j\ne i}^{m-1} |a_i^H a_j|$, $i = 1, \ldots, m-1$, implies that $\|a_j\|^2 \le \sum_{i=1}^{m-1} |a_i^H a_j| \le 2\|a_j\|^2$ for $j = 1, \ldots, m-1$. So, A is a candidate-hub-dominant matrix by Definition 7. □

Before proceeding to our results in Theorem 10, we will first restate a well-known result which will be needed for the proof.

Theorem 9 (monotonicity theorem). *Let $G, H \in \mathbf{C}^{m\times m}$ be Hermitian. Assume H is positive semidefinite and that the eigenvalues of G and $G+H$ are arranged in increasing order, that is, $\lambda_1(G) \le \cdots \le \lambda_m(G)$ and $\lambda_1(G+H) \le \cdots \le \lambda_m(G+H)$. Then $\lambda_\kappa(G) \le \lambda_k(G+H)$ for $k = 1, \ldots, m$.*

Proof. See [5, page 182]. □

Theorem 10. *Let $A \in \mathbf{C}^{n\times m}$ be a hub-dominant matrix with its columns denoted by $a_1,\ldots,a_m$ with $0 < \|a_1\|_2 \le \cdots \le \|a_{m-1}\|_2 \le \|a_m\|_2$. Let $S = A^H A \in \mathbf{C}^{m\times m}$ be the corresponding arrowhead dominant matrix with its eigenvalues denoted by $\lambda_1,\ldots,\lambda_m$ with $\lambda_1 \le \cdots \le \lambda_m$. Let $d^{(i)}$ and $\sigma^{(i)}$ denote the diagonal entry and the sum of magnitudes of off-diagonal entries, respectively, in row i of S for $i = 1,\ldots,m$. Then*

(a) $\text{HubGap}_1(A)/2 \le \text{EigenGap}_1(S)$, *and*
(b) $\text{EigenGap}_1(S) = \lambda_m/\lambda_{m-1} \le (d^{(m-1)} + b + \sum_{i=1}^{m-2}\sigma^{(i)})/(d^{(m-2)} - \sigma^{(m-2)})$.

Proof. Let T be the matrix formed from S by deleting its last row and column. This means that T is a diagonally dominant matrix. Let the eigenvalues of T be $\{\mu_i\}$ with $\mu_1 \le \cdots \le \mu_{m-1}$. Then by Theorem 9, we have $\lambda_1 \le \mu_1 \le \lambda_2 \le \cdots \le \lambda_{m-1} \le \mu_{m-1} \le \lambda_m$. Applying Gershgorin's theorem to T and noting that T is a diagonally dominant with $d^{(m-1)}$ being its largest diagonal entry, we have $\mu_{m-1} \le 2d^{(m-1)}$. Thus $\lambda_{m-1} \le 2d^{(m-1)} = 2\|a_{m-1}\|_2^2$. As observed in the proof of Theorem 4, $\lambda_m \ge b = \|a_m\|_2^2$. Therefore, $\|a_m\|_2^2/(2\|a_{m-1}\|_2^2) \le \lambda_m/\lambda_{m-1}$ or $\text{HubGap}_1(A)/2 \le \text{EigenGap}_1(S)$.

Let E be the matrix formed from T with its diagonal entries replaced by the corresponding off-diagonal row sums, and let $\overline{T} = T - E$. Since T is a diagonally dominant matrix, $\overline{T}$ is a diagonal matrix with nonnegative diagonal entries. Let the diagonal entries of $\overline{T}$ be $\{\overline{d}^{(i)}\}$. Then $\overline{d}^{(i)} = d^{(i)} - \sigma^{(i)}$. Assume that $\overline{d}^{(1)} \le \cdots \le \overline{d}^{(m-1)}$. Since E is a symmetric diagonally dominant matrix with positive diagonal entries, it is a positive semidefinite matrix. Since $T = \overline{T} + E$, by Theorem 9 we have $\mu_i \ge \overline{d}^{(i)}$ for $i = 1,\ldots,m-1$. Let

$$S = \begin{pmatrix} D & c \\ c^H & b \end{pmatrix} \tag{38}$$

in accordance with Definition 9. By Theorem 3, we have $\sum_{i=1}^{m-1} d^{(i)} + b = \sum_{i=1}^{m}\lambda_m$. Thus, by noting $\lambda_1 \le \mu_1 \le \lambda_2 \le \cdots \le \lambda_{m-1} \le \mu_{m-1} \le \lambda_m$, we have

$$\begin{aligned} d^{(1)} + d^{(2)} + \cdots + d^{(m-1)} + b &= \lambda_1 + \lambda_2 + \cdots + \lambda_m \ge \lambda_1 + \mu_1 + \cdots + \mu_{m-2} + \lambda_m \\ &\ge \lambda_1 + \overline{d}^{(1)} + \cdots + \overline{d}^{(m-2)} + \lambda_m. \end{aligned} \tag{39}$$

This implies that $d^{(m-1)} + b + \sum_{i=1}^{m-2}\sigma^{(i)} \ge \lambda_1 + \lambda_m \ge \lambda_m$. Since $d^{(m-2)} - \sigma^{(m-2)} = \overline{d}^{(m-2)} \le \mu_{m-2} \le \lambda_{m-1}$, we have

$$\text{EigenGap}_1(S) = \frac{\lambda_m}{\lambda_{m-1}} \le \frac{d^{(m-1)} + b + \sum_{i=1}^{m-2}\sigma^{(i)}}{d^{(m-2)} - \sigma^{(m-2)}}. \tag{40}$$

□

Note that if there exist positive numbers p and q, with $q < 1$, such that $(1-q)d^{(m-2)} \ge \sigma^{(m-2)}$ and

$$p\left(d^{(m-1)} + b\right) \ge \sum_{i=1}^{m-2}\sigma^{(i)}, \tag{41}$$

then the inequality (b) in Theorem 10 implies

$$\frac{\lambda_m}{\lambda_{m-1}} \le r \cdot \frac{d^{(m-1)} + b}{d^{(m-2)}}, \tag{42}$$

where $r = (1+p)/q$. As in the end of the proof of Theorem 4, it follows that

$$\text{EigenGap}_1(S) \le r \cdot \left(\text{HubGap}_1(A) + 1\right) \cdot \text{HubGap}_2(A). \tag{43}$$

This together with (a) in Theorem 10 gives the following result.

Corollary 3. *Let $A \in \mathbf{C}^{n\times m}$ be a matrix with its columns $a_1,\ldots,a_m$ satisfying $0 < \|a_1\|_2^2 \le \cdots \le \|a_{m-1}\|_2^2 \le \|a_m\|_2^2$. Let $S = A^H A \in \mathbf{C}^{m\times m}$ be a Hermitian matrix with its eigenvalues $\lambda_1,\ldots,\lambda_m$ satisfying $0 \le \lambda_1 \le \cdots \le \lambda_m$. The following holds*

(1) if A is a hub-dominant matrix with $\|a_m\|_2 \gg \|a_{m-1}\|_2$, then S is an arrowhead dominant matrix with $\lambda_m \gg \lambda_{m-1}$; and
(2) if S is an arrowhead dominant matrix with $\lambda_m \gg \lambda_{m-1}$, and if $p(d^{(m-1)} + b) \ge \sum_{i=1}^{m-2}\sigma^{(i)}$ and $(1-q)d^{(m-2)} \ge \sigma^{(m-2)}$ for some positive numbers p and q with $q < 1$, then A is a hub-dominant matrix with $\|a_m\|_2 \gg \|a_{m-1}\|_2$ or $\|a_{m-1}\|_2 \gg \|a_{m-2}\|_2$ or both.

Sometimes, especially for large-dimensional matrices, it is desirable to relax the notion of diagonal dominance. This can be done using arguments analogous to those given above (see, e.g., [14]), and extensions represent an open research problem for the future.

5. CONCLUDING REMARKS

This paper has presented a hub matrix theory and applied it to beamforming MIMO communications systems. The fact that the performance of the MIMO beamforming scheme is critically related to the gap between the two largest eigenvalues of the channel propagation matrix is well known, but this paper reported for the first time how to obtain this insight directly from the structure of the matrix, that is, its hub properties. We believe that numerous communications systems might be well described within the formalism of hub matrices. As an example, one can consider the problem of noncooperative beamforming in a wireless sensor network, where several source (transmitting) nodes communicate with a destination node, but only one source node is located in the vicinity of the destination node and presents a direct line-of-sight to the destination node. Extending the hub matrix formalism to other types of matrices (e.g., matrices with a cluster of dominant columns) represents an interesting open research problem. The contributions reported in this paper can be extended further to treat the more general class of block arrowhead and hub dominant matrices that enable the analysis and design of algorithms and protocols in areas such as distributed beamforming and power control in wireless ad-hoc networks. By relaxing the diagonal-matrix condition, in

the definition of an arrowhead matrix, with a block diagonal condition, and enabling groups of columns to be correlated or uncorrelated (orthogonal/nonorthogonal) in the definition of block dominant hub matrices, a much larger spectrum of applications could be treated within the proposed framework.

ACKNOWLEDGMENTS

The authors wish to acknowledge discussions that occurred between the authors and Dr. Michael Gans. These discussions significantly improved the quality of the paper. In addition, the authors wish to thank the reviewers for their thoughtful comments and insightful observations. This research was supported in part by the Air Force Office of Scientific Research under Contract FA8750-05-1-0035 and by the Information Directorate of the Air Force Research Laboratory and in part by NSF Grant no.ACI-0330244.

REFERENCES

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

[2] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment," in *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 668–677, San Francisco, Calif, USA, January 1998.

[3] H. T. Kung and C.-H. Wu, "Differentiated admission for peer-to-peer systems: incentivizing peers to contribute their resources," in *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, Calif, USA, June 2003.

[4] D. P. O'Leary and G. W. Stewart, "Computing the eigenvalues and eigenvectors of symmetric arrowhead matrices," *Journal of Computational Physics*, vol. 90, no. 2, pp. 497–505, 1990.

[5] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1985.

[6] D. J. Love and R. W. Heath Jr., "Equal gain transmission in multiple-input multiple-output wireless systems," *IEEE Transactions on Communications*, vol. 51, no. 7, pp. 1102–1110, 2003.

[7] D. J. Love and R. W. Heath Jr., "Corrections to "Equal gain transmission in multiple-input multiple-output wireless systems"," *IEEE Transactions on Communications*, vol. 51, no. 9, p. 1613, 2003.

[8] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*, SIAM, Philadelphia, Pa, USA, 2000.

[9] C.-H. Tse, K.-W. Yip, and T.-S. Ng, "Performance tradeoffs between maximum ratio transmission and switched-transmit diversity," in *Proceedings of the 11th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC '00)*, vol. 2, pp. 1485–1489, London, UK, September 2000.

[10] D. J. Love, R. W. Heath Jr., and T. Strohmer, "Grassmannian beamforming for multiple-input multiple-output wireless systems," *IEEE Transactions on Information Theory*, vol. 49, no. 10, pp. 2735–2747, 2003.

[11] C. Murthy and B. D. Rao, "On antenna selection with maximum ratio transmission," in *Conference Record of the 37th Asilomar Conference on Signals, Systems and Computers*, vol. 1, pp. 228–232, Pacific Grove, Calif, USA, November 2003.

[12] A. F. Molisch, M. Z. Win, and J. H. Winter, "Reduced-complexity transmit/receive-diversity systems," *IEEE Transactions on Signal Processing*, vol. 51, no. 11, pp. 2729–2738, 2003.

[13] R. M. Young, *An Introduction to Nonharmonic Fourier Series*, Academic Press, New York, NY, USA, 1980.

[14] N. Sebe, "Diagonal dominance and integrity," in *Proceedings of the 35th IEEE Conference on Decision and Control*, vol. 2, pp. 1904–1909, Kobe, Japan, December 1996.

H. T. Kung received his B.S. degree from National Tsing Hua University (Taiwan), and Ph.D. degree from Carnegie Mellon University. He is currently William H. Gates Professor of computer science and electrical engineering at Harvard University. In 1999 he started a joint Ph.D. program with colleagues at the Harvard Business School on information, technology, and management, and cochaired this Harvard program from 1999 to 2006. Prior to joining Harvard in 1992, Dr. Kung taught at Carnegie Mellon, pioneered the concept of systolic array processing, and led large research teams on the design and development of novel computers and networks. Dr. Kung has pursued a variety of research interests over his career, including complexity theory, database systems, VLSI design, parallel computing, computer networks, network security, wireless communications, and networking of unmanned aerial systems. He maintains a strong linkage with industry and has served as a Consultant and Board Member to numerous companies. Dr. Kung's professional honors include Member of the National Academy of Engineering in USA and Member of the Academia Sinica in Taiwan.

B. W. Suter received the B.S. and M.S. degrees in electrical engineering in 1972 and the Ph.D. degree in computer science in 1988, all from the University of South Florida, Tampa, FLa. Since 1998, he has been with the Information Directorate of the Air Force Research Laboratory, Rome, NY, where he is the Founding Director of the Center for Integrated Transmission and Exploitation. Dr. Suter has authored over a hundred publications and the author of the book *Multirate and Wavelet Signal Processing* (Academic Press, 1998). His research interests include multiscale signal and image processing, cross layer optimization, networking of unmanned aerial systems, and wireless communications. His professional background includes industrial experience with Honeywell Inc., St. Petersburg, FLa, and with Litton Industries, Woodland Hills, Calif, and academic experience at the University of Alabama, Birmingham, Aa and at the Air Force Institute of Technology, Wright-Patterson AFB, Oio. Dr. Suter's honors include Air Force Research Laboratory Fellow, the Arthur S. Flemming Award: Science Category, and the General Ronald W. Yates Award for Excellence in Technology Transfer. He served as an Associate Editor of the IEEE Transactions on Signal Processing. Dr. Suter is a Member of Tau Beta Pi and Eta Kappa Nu.