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.

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.

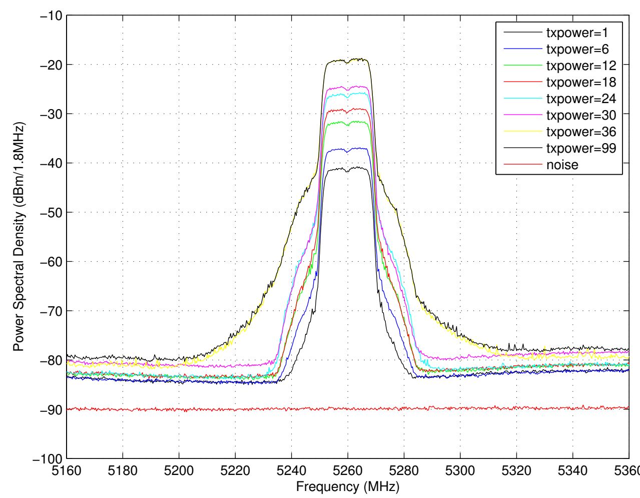

Fig. 1. **The power spectral density of an 802.11a broadcast signal on channel 52 (5260MHz).**

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.

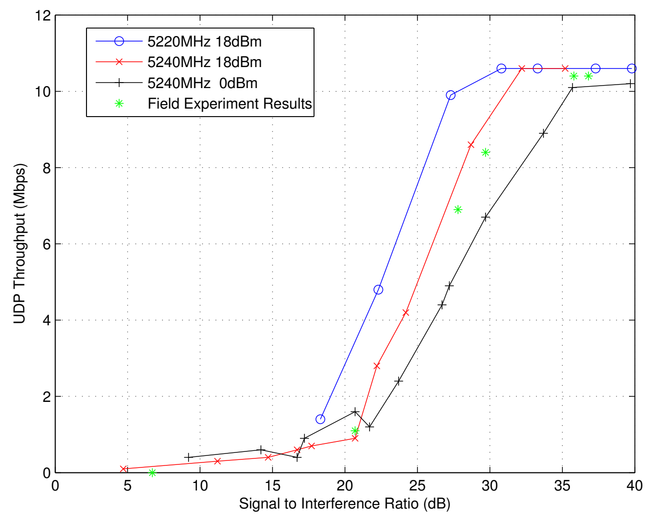

Fig. 2. **Effects of SIR on UDP throughput for a 12Mbps 802.11a link.**

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