

Experimenting with a Multi-Radio Mesh Networking Testbed

Joshua Robinson, Konstantina Papagiannaki, Christophe Diot, Xingang Guo, Lakshman Krishnamurthy

Abstract—Experimentation is an important next step to take in the study of multi-radio mesh networks, but it involves many obstacles preventing a stable and reproducible testbed environment. Using off-the-shelf equipment we demonstrate the impact that specific design choices for access points may have in the performance of a 2-hop network. We find that only after careful calibration with simulation and analytical results can we identify baseline node configurations with predictable behaviour. In particular, we observe that multi-radio wireless platforms face limitations due to increased interference among the several radios inside a single node. Because of this, our mesh routers cannot hold more than 2 radios while requiring a minimum antenna separation of 35db. Throughout this paper we present a methodical approach that, despite being simple, offers a way in which wireless testbeds can be calibrated before experimental deployment and evaluation.

I. INTRODUCTION

Infrastructure wireless networks based on the IEEE 802.11 protocol have become a popular choice as a network access technology. The fact that they require no wires thus allowing for increased mobility has led to their widespread acceptance in home and office environments. Within this context a wireless access point (AP) attaches a multitude of wireless devices to an infrastructure network through a single wired connection. Potential disadvantages of this configuration are i) each AP needs to feature one wired connection which constitutes the most significant part of the network cost, and ii) the range of the access network is limited by the range of the wireless medium.

These limitations gave birth to a new area of wireless communication referred to as “mesh networking”. Within this area each AP no longer needs to feature its own wired connection and may relay traffic generated by other APs. The benefits of the proposed solution are i) the network can now extend beyond the range of a single access point, and ii) the expensive wired connection attaching the wireless network to the Internet now serves more traffic than that of a single AP. [1], [2] discuss further motivation for research in this area.

The shift from “single-hop” wireless networks to “multi-hop” wireless networks leads to many possible design choices for the architecture of an AP. Each AP can simply feature one radio and forward traffic not destined to itself (as in [3]). Alternatively, it can feature several radios to form wireless point-to-point links with neighbouring nodes. When these radios are operated on non-interfering frequencies the capacity of the network increases as a function of the number of radios[1], [2].

The evaluation of the potential of mesh networks has typically been addressed through simulations. However, the research community recently started experimentally evaluating mesh routers [3], [4], [2]. Nonetheless, given the several degrees of freedom in the design of mesh routers (e.g. number of

radios, technologies of radios, etc.) comparison among these efforts is impossible. For instance, in [3] an outdoors multi-hop wireless network is described. Each node operates a single radio, and the results included refer to the performance of the network when packets are broadcast (no multi-hopping is actually tested). In [4] a testbed of 23 nodes is described. The nodes feature 2 radios, one operating at 5GHz using 802.11a and the other at 2.4GHz using 802.11g. The authors look into alternative ways to route traffic across such a network. Lastly, in [2] the authors simulate a 2-radio wireless router using two laptops connected through an Ethernet cable, each one featuring an 802.11b wireless card. They then look into the performance improvement achieved when using 2 radios and algorithms for setting up and tearing down paths through the multihop network.

The focus of these efforts was on the performance evaluation of protocols at the MAC layer and higher. This is the first work we know of to take a step back and study the issues involved in experimenting with a multi-radio wireless platform. Using an off-the-shelf platform we investigate the different choices one may face in the design of a multi-radio access point. We then *methodically quantify the impact that specific platforms choices may have on network performance*. Using simulations as a best-case performance in simple networks, we show that despite the inherent flexibility of open platforms they may suffer from limitations that arise at the RF layer (typically ignored in mesh networking research) and we proceed to quantify those limitations. We demonstrate that *simultaneous activation of multiple radios inside the same node leads to degradation in performance* due to: i) board crosstalk, ii) radiation leakage, and iii) inadequate separation between the several antennas. Through experimental work we show that *more elaborate shielding of wireless cards could potentially reduce the impact of radiation leakage*. On the other hand, the cancelation of board crosstalk caused by interference through a common bus may actually necessitate a custom made platform. Lastly, we show that *antenna separation has a dramatic impact on the performance of a multi-radio wireless node*. We quantify the performance improvement of our node to as much as 100% once antennas are separated by approximately 35db. Insufficient antenna separation essentially prevents a node from simultaneous transmission and reception, while offering no non-interfering communication channels.

Our experience suggests that a wireless testbed suffers from several limitations and reproducibility cannot be taken for granted. Therefore, *testbed results can only be trusted after careful calibration*. We perform this calibration by building up our testbed from simple configurations that can be validated using simulation and analysis. We proceed in this manner be-

cause simulation, analysis, and experimentation can only match in simple scenarios. Experimentation alone can verify reality in more complicated settings and topologies. Consequently, open platforms are likely to prove invaluable in the study of wireless networks but will need to be carefully designed for multi-radio systems¹.

The remainder of the paper is structured as follows. In Section II we present the details of our testbed, along with its monitoring capabilities and our simulation environment. In Section III and IV we establish a baseline performance for a single and 2-hop 802.11b networks respectively. In Section V we summarise and discuss our experimental findings.

II. EXPERIMENTAL SETUP

In this paper we build a wireless access point using off the shelf equipment. Our hardware comprises a Dell Precision 360 workstation with 4 (vertical) PCI slots, allowing up to four wireless cards. The operating system is Linux, kernel version 2.4.26. The wireless cards are Netgear MA311 PCI cards with the Prism2.5 chipset using version 0.1.3 of the HostAP driver. Extra equipment includes cables used for the antenna separation experiments and attenuators used for reducing the carrier sensing range of the different nodes.

All the experiments carried out for this paper were performed indoors. We have deployed 3 of our wireless APs along a corridor in our lab where the first and second nodes are 12.5 metres apart, and the second and third 20.6 metres apart. To ensure that our results are not affected by other 802.11b or bluetooth activity, we run our experiments at night and switch off all 802.11b access points that normally operate in the building where we conduct the experiments. We also use NetStumbler to further ensure no other access points will interfere. Each experiment is repeated 5 times with a duration of 10 seconds. The traffic source is always a backlogged UDP flow transmitting packets of a prespecified size between two explicitly identified nodes. All wireless cards are set to full 11Mbps transmission rates (e.g. autorate is disabled).

To get a complete picture of our access points, we collect measurements at three different levels: IP, MAC, and physical. Currently, our measurement suite includes i) tcpdump to capture all IP packets transmitted and received on each interface, ii) kismet, used to capture all 802.11b control and data frames, and iii) per-packet signal, noise, and data rate information provided by the driver (with trivial modifications).

In parallel to our experiments we also simulate the network behaviour using ns-2, version 2.27, employing the default two-ray ground propagation model. Simulation findings will act as the ideal case, since for the simple topologies tested they match analytical models. Simulations essentially provide us with the best case performance one should expect from the testbed.

III. ESTABLISHING A BASELINE - SINGLE-HOP EXPERIMENT

Before investigating the impact of different features of a wireless access point, we evaluate the performance of our

¹We are aware of two academic research groups that actually design their own multi-radio wireless platforms (Rice with the TAPS platform, and CTTC in Barcelona, Spain).

testbed under the simplest scenario, i.e. that of a single-hop wireless network. Two nodes are configured with a single wireless card operating on the same channel.

A. Single-hop (node1 -> node2: 1 channel)

In Fig. 1 we present the average throughput achieved by the flow when packet size increases, along with the results predicted by simulation. We notice that the experimental and simulated results match very closely, with a small difference most likely due to non-ideal channel conditions not modelled by simulation.




Fig. 1. Baseline Single-hop Throughput vs. packet size.

We have repeated the same experiments with the 802.11 “channel reservation” mechanism, which uses RTS/CTS (Request to Send/Clear to Send) messages as short control frames to reserve the channel and prevent packet collisions. The experimental results again match the simulation results with sufficient accuracy. Note that across all 5 experiments performed for each different packet size the minimum, average, and maximum achieved throughput are indistinguishable while instantaneous rates can vary significantly.

B. Impact of multiple wireless cards inside a node

Having established a stable and reproducible single-hop testbed we proceed to investigate the impact of multiple wireless cards on the achievable throughput of a UDP flow. We begin our experiments with a single wireless card and proceed to insert a second wireless card in the next available PCI slot. The card is powered inside the node, but does not transmit; we simply set it to monitor (passive) mode and tune it to an orthogonal channel. We then re-run our experiments and collect the new throughput values, iteratively incrementing the number of cards until all 4 PCI slots are occupied. Fig. 2 presents the relationship between the achieved throughput and the number of PCI slots occupied in the wireless node.

The results in Fig. 2 clearly show that having more than two PCI cards in the same access point introduces significant radio noise. It should be emphasised that each additional card has been placed in a passive state, and therefore any interaction between the cards is due to radiation leakage from the chipset, connectors on the card, and the antennas. This leaked radiation sums algebraically at the active card; this is why one additional card does not affect throughput. The absolute reductions in throughput are hardware dependent and should not be taken

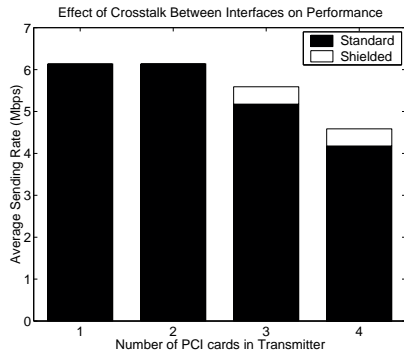


Fig. 2. Throughput reduction due to additional PCI cards (UDP packet size = 1472 bytes).

as universal. Different vendors may provide different levels of shielding or emit less radiation. The general pattern, however, should be consistent across all types of wireless cards.

We should note at this point that the impact of multiple cards has been briefly described in [4] where it was reported that no two 802.11a and 802.11g cards could coexist in the same box without significant performance impact. In addition, in [5] the authors mention that their experiments revealed that Netgear cards require a separation of 6 inches, while Cisco cards appeared to generate interference in the vertical plane, and a vertical separation of 3 inches was necessary for no interference. However, to the best of our knowledge, this is the first quantitative study on the cumulative impact of additional cards within a wireless access point.

Further, our results are representative of a best-case scenario as any transmissions from the other cards will produce even more radiation. This leaked radiation from the cards increases the amount of carrier energy, causing the transmitter to defer unnecessarily before sending a packet. Additional experiments show that the number of wireless cards on the receiving node has no significant impact on system performance.

C. Impact of shielding

One of the reasons why we observe the aforementioned results is that the transmitting card is not shielded enough against the leaking radiation from the other, otherwise dormant, cards. We thus next attempted to add shielding to the PCI cards and evaluate any performance improvement. Our shielding is made out of aluminium foil wrapped around the single transmitting card. In such a way, the card should be more robust to the surrounding radiation. Indeed, as shown in Fig. 2, shielding leads to a slight (5-10%) performance improvement. In the case with all four slots occupied (rightmost bar in Fig. 2), the achieved throughput rises from approximately 4.1 Mbps in the unshielded case to 4.6 Mbps with shielding (as shown in Fig. 2). This confirms that at least a significant part of the lost throughput is due to over-the-air leakage between cards. We expect that more thorough shielding may reduce even further the performance discrepancy, but it should not eliminate it, unless care is taken such that the board crosstalk is also minimised. In fact, there already exist integrated multi-radio solutions² that

²refer to www.engim.com

address the multi-radio functionality at the RF level using noise cancellation techniques. These boards are reported to be able to accommodate the simultaneous transmission of three on board radios tuned to orthogonal channels.

IV. ESTABLISHING A BASELINE - TWO-HOP EXPERIMENTS

In the previous section we found that our node cannot hold more than two wireless interfaces without a performance impact. We now move on to evaluating design choices when a two-hop network is taken into account.

A. Issues with RTS/CTS

The simplest scenario in a 2-hop configuration is for each node to operate one radio on a common channel. We created a 3-node chain in our lab's central corridor. The distances between nodes are not long enough to require multi-hopping, so we instead forced the intermediate node to be used through static routes. The two edge nodes are able to hear and decode all packets – a slight divergence from common experimental setup where the two edge nodes are only in carrier sensing range of each other.

Fig. 3(a) shows that our experimental results match simulated results when the RTS/CTS mechanism is not enabled. When all the nodes use RTS/CTS, our results deviate significantly from simulation. Using the statistics collected at Layer3 by *tcpdump* we focus on the first 140 ms of the experiment and present the packet transmission and reception times for each node in Fig. 3(b)³. Fig. 3(b) shows that *node0* makes 52 packets available for transmissions within the first 50 ms. Only 5 of these packets depart within the first 10 ms. Upon reception, *node1* immediately transmits them to *node2*. Then the medium remains idle for approximately 40 ms until 4 more packets are received by *node1*. Periods of inactivity appear frequently and last up to 20 ms, while the sending queue is always full. Such a behaviour would not be expected. Further observation of the *kismet* logs shows that there are no control messages lost during the idle periods.

To pinpoint the reasons that could influence (but still not justify) the observed behaviour we perform two additional tests. First, we install enough attenuation at the two edge nodes such that they cannot hear each other (the packet loss rate exceeded 95%). We repeat the same experiment and notice that the performance degradation originally witnessed has now been rectified and the simulated performance agrees with our experiments. Second, we install a second card in the middle node and tune both cards on the same channel as the rest of the nodes. We remove the attenuation at the edge and restore the state where all nodes are within range. We rerun our experiment and we observe that now RTS/CTS performs exactly as expected (Fig. 3(c)).

Upon completion of the two experiments we are still not quite clear why our testbed behaves the way reported when the middle node features one wireless card and employs the RTS/CTS mechanism. We speculate that it may be due to the

³Recall that *tcpdump* captures the packets when they become available to the network driver and therefore does not account for the time until the medium becomes available for transmission.




Fig. 3. Throughput vs. packet size for 2-hop experiments, when all nodes use the same frequency.

implementation of RTS/CTS on our wireless cards, especially given that our results are as expected across all configurations when we do not use RTS/CTS. The only actual difference between *Setup1* and *Setup2* is the fact that one card shares the functionality of two cards. At this point we should also note that we have observed other issues with RTS/CTS on our cards. For instance, only after a warmup of 2-15 minutes following a soft reset does RTS/CTS offer the expected throughput values. If the cards are not allowed enough warmup time, then single-hop throughput never exceeds 500Kbps.

B. Channel separation

The reasonable next step to take is to assign the two wireless cards in the middle node to orthogonal channels and evaluate the performance impact. Before doing so, we look into the effects of this selection. The 3 orthogonal, non-overlapping, channels in the 802.11b standard are channels 1, 6, and 11. A common assumption is that if two wireless interfaces in one device are placed on two orthogonal channels, they can operate independently and at full rate.

We reconfigure our three node chain and set the middle node to transmit to the two edge nodes using separate interfaces set on different channels. This allows us to place each link on a unique channel. More specifically, *node1a* uses channel 1 for its communication to *node0*, and *node1b* varies the communication channel to *node2* from 1 up to 11. We measure the aggregate throughput and present our results in Fig. 4(a) where the x-axis captures the separation between the two active channels⁴. The dashed line marks the experimental throughput value for our single-hop scenario, while the white and black bars show the individual performance for each flow.

Fig. 4(a) shows that at a separation of zero, both flows share the same channel and can therefore coordinate channel access allowing the channel to be fully utilised (relative to single flow throughput). A channel separation of between one and three channels causes a reduction in aggregate throughput. We expect this result since the active channels will still overlap and therefore the two flows compete for the same channel, yet are unable

⁴Note that this value corresponds to the 802.11b channels and would require multiplication by 5MHz to get the absolute frequency difference. Furthermore, because of channel overlap in 802.11, a minimum distance of 5 channels is required for orthogonality.

to coordinate intelligently. As channel separation increases to five or six channels, the aggregate throughput rises to about 10 Mbps. This is less than double the single-hop throughput which was expected given that both channels are orthogonal. Surprisingly, as the channel separation continues to increase, the aggregate throughput begins to decrease, eventually declining below the single hop baseline. At this point the two active channels are completely orthogonal, and without any radio interference aggregate throughput should simply continue to increase until a maximum of 12.8 Mbps (twice the measured single-flow throughput). Note last that both flows achieve approximately equal throughput across all experiments.

In Fig. 4(b) we present results for the case when the middle node is the sink of the 2 flows from the two edge nodes. The aggregate throughput never increases beyond a single receiver's throughput. This indicates that even though both interfaces may be on orthogonal channels, they are not operating independently. Given how closely the aggregate throughput matches the available throughput on a single channel, our results suggest that the two interfaces cannot simultaneously receive packets. The drop in aggregate throughput at a channel separation of one is due to the signal overlapping enough such that the error rate at each interface greatly increases.

We then examine the case when one interface constantly receives packets while the other constantly transmits. The aggregate performance does not significantly surpass the single-hop rate (Fig. 4(c)). In fact, we observe the same general trend of an initial dip and then levelling off as seen in Fig. 4(a) and 4(b).

Our results on the 2-hop experiments so far suggest the absence of orthogonality for the channels used by the two interfaces of *node1*. When both interfaces either transmit or receive, resources are equally shared among interfaces. But when one interface transmits and one receives, the transmitting interface is able to almost starve the receiving interface. This unfairness is due to the fact that the transmitting interface causes additional interference at the receiving interface, while the receiving interface does not affect the transmitting interface.

C. Impact of antenna separation on channel orthogonality

In the previous section we have shown that our basic configuration for the multi-radio node *node1* prevents the simultaneous

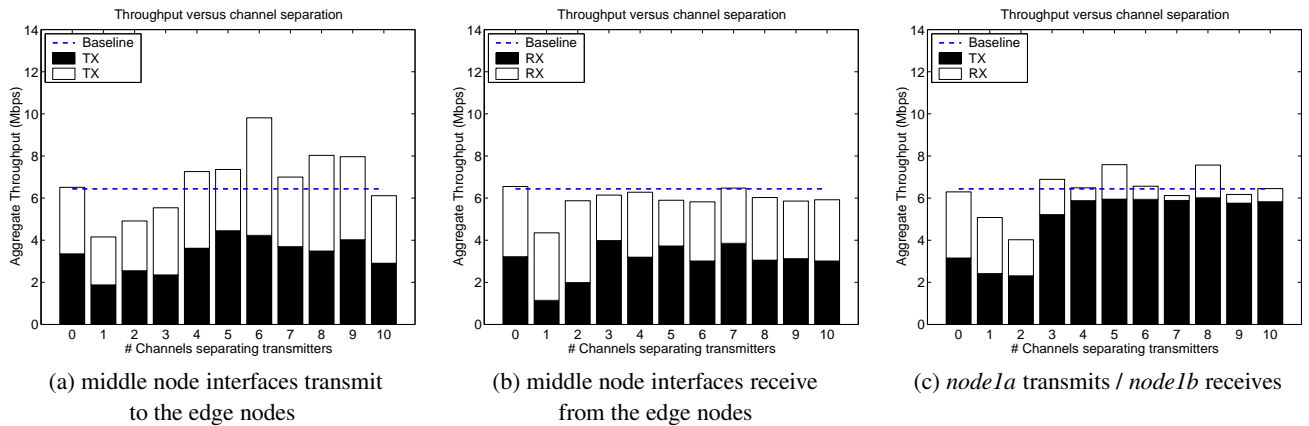


Fig. 4. Throughput vs. channel separation. Setup3: 2-Hop, Two Interfaces, 2 Channels. *Node1a* uses channel1 for its communication with *node0* and *node1b* uses channel2 to communicate with *node2*.




Fig. 5. Throughput vs. channel separation in a 2-hop wireless network using two different channels. *Node1a* uses channel1 for its communication with *node0* and *node1b* uses channel2 to communicate with *node2*.

operation of the two interfaces at full capacity. In addition, regardless of channel separation the obtained results are nowhere close to the expected performance.

As mentioned in Section III one of the reasons for performance degradation in a multi-radio system is the separation between the different antennas, which needs to be 35db according to the IEEE standard. Antennas too close to each other can interfere in ways that significantly limit their operation [6]. We thus modify our testbed to feature connectors on the middle node that allow a separation of 35db (a distance of approximately 1 metre for the kind of antennas in our testbed). Using this configuration, we rerun the experiment with both interfaces transmitting. A comparison of Fig. 4(c) with the results presented in Fig. 5(a) shows that a distance of 1 metre between antennas allows the interfaces to operate independently of each other on non-overlapping channels. Two observations can be made on Fig. 5(a). First, with four channels of separation, both interfaces almost achieve full rates. This is despite the fact there is still some small overlap between the channels. Second, the aggregate throughput begins to decrease as the channel separation reaches its maximum value. This suggests imperfect behaviour in the signal processing on the wireless cards.

The aforementioned results describe the case when we set one channel to 1 and then vary the second channel from 1 to 11. We generalise our previous results performing experiments

for all possible combinations between the frequencies used for the two channels. Our findings are shown in Fig. 5(b) and Fig. 5(c). The difference is striking. When the two antennas on the middle node are separated by 1 metre, one can identify orthogonal channels which will not interfere. On the other hand, when one uses the default configuration that necessitates the two antennas to be in close proximity, then there is essentially no interference-free channel pair and orthogonality is no longer deterministic.

D. Impact of antenna separation on throughput

Having calibrated our testbed, we can now report our findings on the 2-hop performance of our testbed. In what follows, we compare the experimental throughput achieved with the single-hop throughput. If the two hops are truly independent then each hop should be able to operate at full capacity. Fig. 6(a) and Fig. 6(b) present the relationship between throughput and packet size when the two antennas are close or far from each other respectively.

Fig. 6(a) reinforces the fact that the two interfaces cannot operate independently despite being on orthogonal channels. The transmitted energy from one interface is strong enough to distort the internal filters and amplifiers in the receiver. This results from the fact that, due to cost, almost all wireless cards initially accept the entire 2.4 GHz band, filtering the desired 22 MHz




Fig. 6. Throughput vs. packet size in a 2-hop wireless network operating using 2 orthogonal channels (1 and 11).

channel on a second pass[6]. Because the path-loss exponent is at minimum two and tending more towards four, the received signal power will be several orders magnitude less than the received signal from the other card in the node. This huge disparity in power causes the analog circuits to function incorrectly and therefore distort the received signal causing packet errors.

From Fig. 6(b) we have a clear indication of the performance improvement when the two antennas are separated by 35db. Distancing the antennas reduces inter-card interference. Nonetheless, we most likely have not completely eliminated the influence of each interface on the other, and hence the throughput falls slightly short of that predicted by simulation.

V. SUMMARY/DISCUSSION

In this work we outline some possible pitfalls of multi-radio experimentation and the impact of the mesh router architecture on the performance of a 2-hop wireless network. Using commodity hardware we quantify the impact of different design choices. We show that multi-radio systems need to address three issues: i) board crosstalk, ii) radiation leakage, and iii) antenna separation. Our experiments show that using an off-the-shelf platform:

- One cannot accommodate more than 2 cards inside the same box without more elaborate shielding. Even under such circumstances techniques are needed for the cancellation of the crosstalk. Custom made hardware may be needed for multi-radio systems with more than 2 radios.
- Even if a single node contains 2 wireless cards alone, these cards are not going to be able to receive or transmit traffic at the same time, unless their antennas are separated by

more than 35db. Only with significant antenna separation can one identify non-interfering communication channels.

- Errors in the implementation of protocols may lead to erratic behaviour, such as warmup periods for proper functionality and possibly unexplained behaviour.
- The ability to collect accurate and complete statistics at the MAC and physical layers is crucial to troubleshooting problems.

Throughout the course of this work we realised that wireless experimentation can sometimes be more misleading than careful simulation. Reproducibility is difficult to achieve and lack of persistence in identifying the reasons driving the behaviour of a testbed can lead to false findings. Only by starting from small, carefully calibrated experiments can we have confidence in our results and move toward more complicated and interesting scenarios.

Our experience demonstrates that multi-radio systems face limitations and experimentation in this domain may actually need to be a cross-disciplinary action where RF experts work together with networking researchers in order to address the several problems. Off the shelf equipment is attractive due to its flexibility but custom made hardware may actually be a necessity (essentially we are currently simulating systems that have not been built). Until then simulation can provide a measure against which experimental testbeds can be calibrated *in their simplest form* before embarking into the evaluation of sophisticated mechanisms addressing alternative MAC and routing protocols in large scale. Our work highlights such a process, the limitations of the studied system and offers potential avenues for the implementation of mesh routers using general purpose hardware.

VI. ACKNOWLEDGMENTS

We would like to thank Ratul Mahajan for providing us with the driver patch that allows reporting of signal and noise measurements on a per-packet basis, and Adrian Stephens for comments throughout the course of this work.

REFERENCES

- [1] R. Karrer, A. Sabharwal, and E. Knightly, "Enabling Large-scale Wireless Broadband: The Case for TAPs," in *HotNets*, Cambridge, MA, 2003.
- [2] S. Lee, S. Banerjee, and S. Bhattacharjee, "The Case for a Multi-hop Wireless Local Area Network," in *IEEE Infocom*, Hong-Kong, China, 2004.
- [3] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11 Mesh Network," in *ACM Sigcomm*, Portland, OR, 2004.
- [4] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *ACM Mobicom*, Philadelphia, PA, 2004.
- [5] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *Broadnets*, 2004.
- [6] T. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Prentice Hall, Dec. 2001.