EWANT: The Emulated Wireless Ad Hoc Network Testbed

Sagar Sanghani, Timothy X Brown, Shweta Bhandare, Sheetalkumar Doshi
University of Colorado at Boulder, 80309-0530
{sanghani, timxb, bhandare, doshi}@colorado.edu

Abstract— **In this paper we demonstrate a wireless 802.11 testbed. This testbed allows for direct comparison between mobile wireless routing protocols without affecting their routing or MAC layer protocols or inter-layer interaction. We have built a low cost environment to facilitate such wireless network research. The core idea is "compressing" the network and emulating mobility without actually moving the nodes. We successfully emulate the RF effects along with mobility. The viability of this test bed is checked with the help of a Click implementation of the Dynamic Source Routing protocol. Our experiences show that such a testbed provides valuable feedback not available through simulation.**

I. INTRODUCTION

A wireless ad hoc network consists of autonomous mobile nodes that spontaneously spawn a communication network between each other. The nodes are usually PDA's or laptops equipped with wireless ethernet cards. Because these nodes are constantly moving, the link states are never stable. Wired routing protocols do not have to account for rapidly changing topology or link quality and thus cannot be used efficiently over such a network. Hence ad hoc wireless specific protocols such as the Ad hoc On demand Distance Vector (AODV) [1] and the Dynamic Source Routing (DSR) have been developed [2]. While a large amount of work has gone into developing these protocols, there has never been a standard way to test their implementations. This creates the need for a standard, reliable, low cost test bed that allows for direct comparison between mobile wireless ad hoc routing protocol implementations.

Most wireless protocol testing can be done with software simulations. Software simulations are an excellent choice for initial design and an estimation of results. But, they do not realistically duplicate the physical layer. Also a software simulation cannot catch subtle bugs seen in interactions between the operating system, the system hardware, and the design environment. A software simulator also ignores inter-layer communication, which is integral to the effectiveness of these protocols. Specifically in our work, we have observed a number of significant discrepancies between simulated and hardware results. This makes hardware testing critical.

Prior work involving testbeds has been attempted mostly to test some specific routing protocol developed. Maltz et al. [3] tested their routing protocol using a drive-test tool. Mobility was generated by nodes placed in student driven vehicles. Tests carried out in such a manner are very expensive. They do not lend themselves to rapid experimentation and testing cycles. Also repeated trials cannot be carried out with the exact same conditions. Lundgren et al. [4] developed choreography scripts which directed the movement of students around a building. Although tests here were carried out on a large scale, each test run required a large number of people. Theoretically this testbed allows repeatable experiments but at high cost. Zhang et al. [5] introduced another approach of connecting the nodes using wired ethernet to a hub or a router. The routing tables on the nodes were altered to simulate mobility. This approach is both repeatable and affordable, however it imposes an artificial substitute for an RF based MAC protocol that does not reflect RF interference and capture effects. As will be demonstrated later, a number of subtle design bugs cannot be detected in a tool that ignores RF effects. Kaba et al. [6] introduced an approach similar to ours. However these nodes were hardwired to each other and ignore the effects of mobility. Thus, none of the testbeds discussed above simultaneously provide node **mobility**, an accurate reflection of **RF effects**, hardware and software **independence**, and **repeatable controlled experiments**. Our aim is to provide all the features discussed above

In Section II we describe an implementation of the testbed, consisting of Dell CPi laptops with the Cisco Aironet [7] 350 series wireless ethernet cards. The testbed has been designed to fit in a single room using shielding and attenuation. This allows emulation of RF propagation effects of a huge network spread over hundreds of meters in a very small area. The hardware testbed achieves mobility via an antenna switching mechanism.

In Section III we present an experiment demonstrating the features of the testbed. We also show the importance of "live" testing over simulation. Tests are carried out with a Click [8] implementation of the DSR protocol [9], [2].

II. HARDWARE IMPLEMENTATION

This section lists the requirements of a hardware testbed and then details our approach in satisfying the same. We also describe the manner in which mobility is achieved.

A. Testbed requirements and our approach

Based on prior work and our initial testing experiences we derived the following requirements for our testbed.

- Tests must be controlled and repeatable.
- Shared RF media effects such as interference and channel capture must be present.

0-7803-7701-X/03/$17.00 (C) 2003 IEEE

- Node mobility, it's effects on topology, and the RF media must be present.
- The protocols tested, operating system, and hardware used should be independent of the testbed.
- The testbed should not impose any restrictions on the operating system or underlying hardware.
- A variety of spatial layouts must be possible. Changing these layouts should not be time consuming.
- The testbed must be affordable, manageable, and allow rapid prototype testing.

Radio signals broadcast in all directions. Noise, interference, and clutter within radio range affect the observed network connectivity and throughput. If the nodes are static then fading is purely a static function of position. However due to the motion of the nodes, dynamic effects such as fading are seen. For example a node moving behind a building will have a sudden change in signal strength which disrupts routing. Even more interesting effects such as sporadic connectivity at the edge of communication range challenge route maintenance mechanisms. These effects affect design parameters of a routing protocol significantly. The difficulty of incorporating all these effects in simulation suggests using a wireless testbed.

In ad-hoc mode 802.11b cards can communicate over a distance of a hundred meters. Maltz et al. [3] tested a network at this scale. However as the nodes are very far away, it is difficult to manage the testbed. Also experiments cannot be recreated. For example one of the scenarios we show later is about changing a single variable in the code and repeatedly running the test. This would be impossible with the above mentioned drive test tool. Also the cost of such an approach is very high. Not only the initially investment but also the cost of running each trial. Thus the short range approach was chosen over the long range one. We choose a middle ground between simulation and a full scale network: *emulation*. We aim to "compress" a wireless network so that it fits on a single table.

EWANT uses standard 802.11b wireless LAN (WLAN) cards with external antennas. If used directly in a single room, all the cards will be able to talk to each other, but this would not be a very interesting topology. Inserting attenuators between the card and the antennas, unobtrusively shrinks the range so that an entire multi-hop network can be setup on a table top.

The attenuators are two-way, so the received as well as transmitted powers are attenuated. To emulate mobility a switch is used. The attenuated output from the wireless card is fed to the input of a 1:4 RF multiplexer. At the outputs, four antennas are connected. This setup is repeated for each node. The antennas can be placed in suitable positions to form dynamic topologies. The rest of this section describes the details of WLAN hardware, table top test range, antennas and mobility.

B. WLAN Hardware

For the choice of WLAN cards we decided to go with either Cisco or Orinoco, simply because of their ease of availability, excellent support for Linux and the ability to connect an external antenna. The external antenna jack is provided to increase the range of the card with higher gain antennas. But in our case we use it for the exact opposite reason.

The Orinoco Silver or Gold card provides an option to connect an external antenna. These cards also have a built in internal antenna. When an external antenna is connected to these cards, the signal to the internal one is cut-off. However, shielding is still required to block leaking signals [6].

The Cisco Aironet LMC 352 provides an option to connect an external antenna, but has no internal antenna. When inserted into the PC slot the card is completely covered with no part exposed. The PC slot of the laptop now acts as a shield for this card. Due to this the cards do not communicate with each other even over a distance of a few centimeters. We chose the Cisco over Orinoco since they provide transmit power control necessary for minimum energy protocols [10].

The Cisco Aironet provides two outputs with MMCX connectors for external antennas. Two antennas can be used to provide spatial diversity. In order to reduce the proliferation of untested—and thus unapproved—combinations of transmitter, amplifier and antenna in the unlicensed and shared 2.4GHz spectrum, the FCC requires vendors to use non-standard external antenna connectors. Cisco provides an RP-TNC to MMCX adaptor. A converter from RP-TNC to the standard BNC connector was then found. However this setup was found to be prone to leakage radiation and also was mechanically weak. Finally an MMCX to BNC converter was found from a third party vendor. This allowed for a direct connection to the antennas through a BNC cable. This setup is mechanically stable and also provides significant reduction of leakage radiation.

C. Shielding

It was seen that when the WLAN cards were inserted into two laptops adjacent to each other, they would not talk even with the cards facing each other over a distance of a few centimeters. However, if the MMCX to BNC cable was inserted into the card and the open end was connected to a terminator then the nodes would communicate even over a distance of a half meter. This suggested that the MMCX connector itself radiates. For our experiment this leakage radiation must be shielded properly or else all the nodes will be able to communicate with each other through uncontrolled leakage radiation rather than the intended antennas. Even though the percentage of these spurious signals was significantly lower it still can lead to erroneous measurements.

Spurious leakage signals are hard to detect and thus hard to prevent. Some solutions for this problem were covering the connector with copper tape or aluminum foil. This setup was found to be unreliable and prone to human error, simply because the foil kept coming off.

Finally, to get reliable readings we put each laptop inside a cardboard box covered with aluminum foil. The face of this box was kept open for ease of access and to prevent over heating the laptops. For even more reliability we plan to use custom made metallic boxes for each laptop.

D. Simulation of Mobility

It is difficult to simulate a continuous change in position without physically moving the antennas. We decided to ignore the smooth node movement and assumed that the worst case scenario is a sudden change in the signal strength. This can be seen in practice too, for example, when a node moves behind a building or a similar obstruction.

A computer controlled variable attenuator can be used to gradually vary the attenuation. For example, a stepper motor based attenuator can be controlled, remotely, via PC and then the attenuation can be varied in small steps repeatedly. This approach thus emulates a limited sort of continuous change in position. However we decided not to include this in our setup, mostly because the cost of such attenuators is substantially high. Also as mentioned above we were mainly interested in observing effects sudden changes in the RF environment.

We use the Agilent 34970A Data Acquisition / Switch Unit. The 34970A is a multi-purpose device that can be used for varied applications. These applications are available through the use of detachable modules. For our purpose we use the 34905 RF MUX module. This module consists of two, independent, RF, 1-to-4 multiplexers. Thus in our case the laptops were connected to the input port of the 34905 module. Now 4 antennas can be connected to the output. Thus by simply switching which antenna is connected to the input, the nodes “move”.

The Agilent 34970A is shipped with the HP-IB and RS-232 interfaces for remote control. A Win32 library file is also provided. This allows a user to control the device from a Windows based computer. We used Visual C++ to write an antenna switching script interpreter. With this interpreter we can write arbitrary scripts to precisely and repeatably switch between antennas.

Some care must be used with this switch in our environment at 2.4GHz. We measured the crosstalk between signals inside the switch. We found that adjacent channels had only 20dB of attenuation, while non-adjacent channels had 40dB of attenuation. Further, signals on different cards had 50dB of attenuation. This leads to a “near-far” problem. For instance a near antenna may have 40dB attenuation due to propagation and a far antenna may have 70dB. If these are adjacent channels, then the near antenna will never have more than 60dB of attenuation (40dB + 20dB). This means the near antenna will always have a stronger signal even when the signal is switched away to the other channel. Care was used to avoid this near-far problem in our experiments.

E. Test-Bed Layout

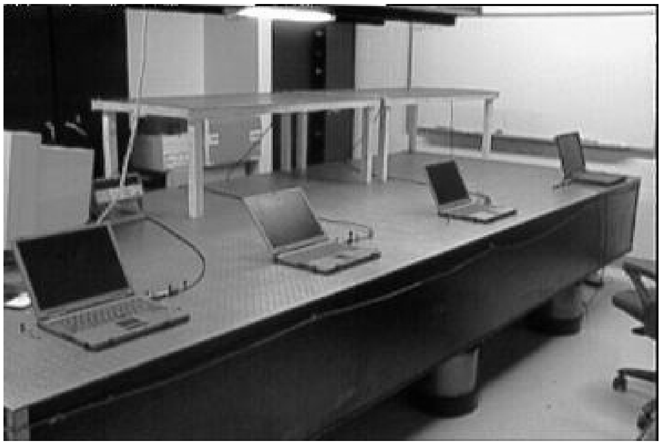

Fig. 1. Emulated Wireless Ad hoc Network Testbed (EWANT)

The actual test-bed is made of two 60cm-by-120cm tables with metallic planar surfaces. Attenuators are used to shrink the network so that multihop behavior is possible in this limited area. These tables have holes on their surface as shown in Fig. 1. Cables are run from the output of the switch to the table and are then connected to the antenna. The metallic table surface isolates the testbed from interacting with the cable runs and provides a uniform propagation environment. The main advantage of this setup is that the tables can be placed in a single isolated room. The high attenuation on each node combined with the room walls effectively isolates this environment from outside signals.

F. Antenna design

The BNC cable is brought from under the table to connect to the antenna. The BNC is converted to female SMA using an adaptor. The SMA connector fits perfectly into the holes provided on the metallic sheet. This connector is screwed to the table. A thin wire of 3cm is then inserted into this SMA connector. Thus we form a quarter wave dipole antenna as shown in Fig. 2. These holes on the table do not cause any periodic reflections or scattering as they are significantly smaller than the wavelength of the signal (12.5cm). The goal is to shrink the range by a factor of 50 to 100. These 3cm antennas placed directly above the table provide a roughly equivalent height of 1 to 2m above the ground in a full scale system. With this setup, the antennas can precisely and repeatably be arranged spatially to form different topologies

G. Attenuation

As mentioned earlier 802.11b nodes can talk over a distance of a hundred meters. Thus if used directly, in a room, all the nodes will be able to see each other. Obviously if all the nodes talk to each other, no interesting topologies can be formed. Attenuators are used to shrink the network so it fits into a single room. Thus we are simulating large scale effects on a small scale for easy access.

The choice of attenuation is driven by two considerations. The first is the range scale to work on. Too high attenuation will mean that the antennas must be placed in the near-field (within about one wavelength, 12.5cm at 2.4GHz) of each other to communicate. Near-field behavior is not typical of full-scale behavior and should be avoided. Too little

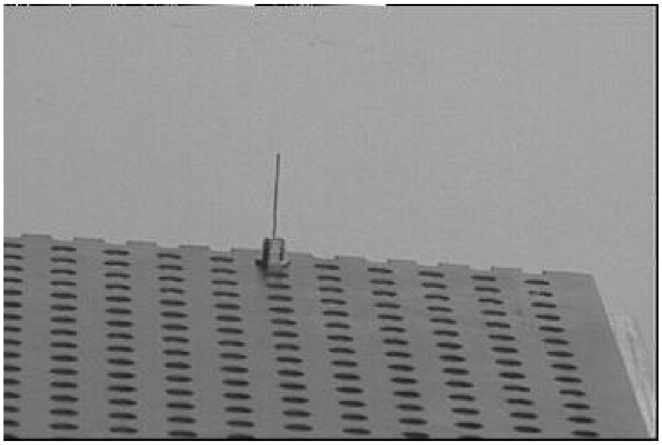

Fig. 2. Antenna used for each node of the hardware test-bed

attenuation will mean that the nodes in a room will be able to communicate without needing to go multihop. Also, low attenuation results in strong signals that can still be received via crosstalk in the switch.

The second consideration is interference. If the attenuation is too low, then signals from outside of the room may be received and interfere with the experiments. If the attenuation is too high then interference can again become a problem as follows. If we use XdB of attenuation on each node, then the attenuation is $2X$dB on a complete transmit and receive path. But, outside interfering signals only see XdB of this attenuation. The outside signals see excess pathloss due to their being further away and having to penetrate the walls of the room. Call this excess attenuation YdB. Thus in order to isolate the testbed from outside signals our requirement is $X < Y$.

We measured the penetration loss of the labs concrete block walls by placing a transmitter 4m outside of the room and measuring the received signal strength with a spectrum analyzer just outside and just inside of the room. The resulting loss is 13dB. This matches similar measurements reported in [11]. Ultimately, we wanted the testbed ranges to be 1m or less. A user outside the room would be at least 3m from our testbed. Using free-space pathloss, outside users would have an additional 10dB signal loss due to the increased distances. Together these suggest that $Y = 23$dB for an interferer directly outside the lab. In fact, the nearest known interferers are located at least 6m away and must penetrate two walls suggesting $Y > 40dB$.

In addition to explicit attenuation, the switch and the BNC-SMA cables added further attenuation. Fig. 3 shows the received signal strength plotted across all the eleven 802.11b channels. The upper plot is obtained when the signal is directly connected to a spectrum analyzer. The lower plot shows when the signal is connected to the spectrum analyzer through the switch. The error bars were obtained with 10 repeated trials. As can be clearly seen the switch does not have any significant frequency dependence, but it does attenuate the signal by approximately 4dB for each channel.

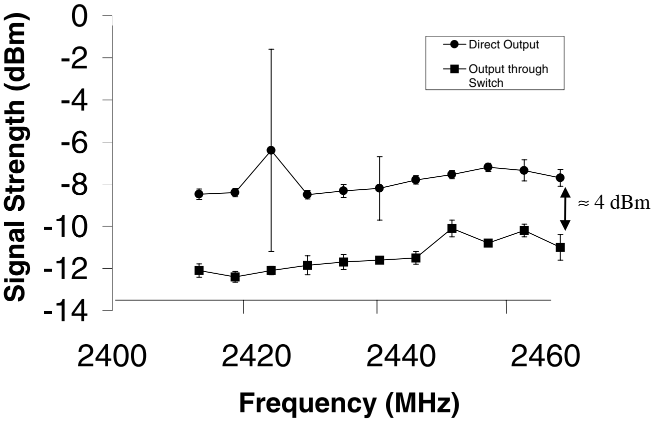

Fig. 3. Plot of Signal Strength Vs Frequency, with and without the Agilent switch

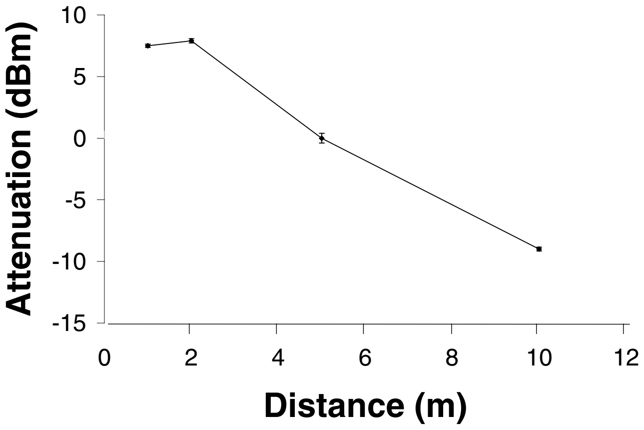

Fig. 4. Plot of the signal strength vs length of BNC cable used

The type of BNC cables used were ordinary, RG-58. These have significant losses above frequencies of 500MHz. This attenuation is due to the "skin effect". This effect describes a condition where as frequency is increased, signal is carried through a conductor further and further away from the center. Fig. 4 depicts the plot of signal strength vs. cable length as measured in the lab. From this plot attenuation of the BNC cables was found to be 2.3dB/meter. This calculated loss is higher than the rated one due to the cable and the connectors used. Therefore, for constant attenuation on all nodes we used cables of the same length. Although cables with higher operating frequencies are available we decided not to use them. This is because these cables are significantly more expensive and we do not need to minimize cable losses. The cable and switch attenuation are part of the overall excess attenuation.

Putting all these factors together we found a good set up is 3m BNC cables with an added attenuator value around 20dB.

III. DEMONSTRATION OF FEATURES

As a demonstration of the features of EWANT we present an experiment using the DSR routing protocol. DSR is an on

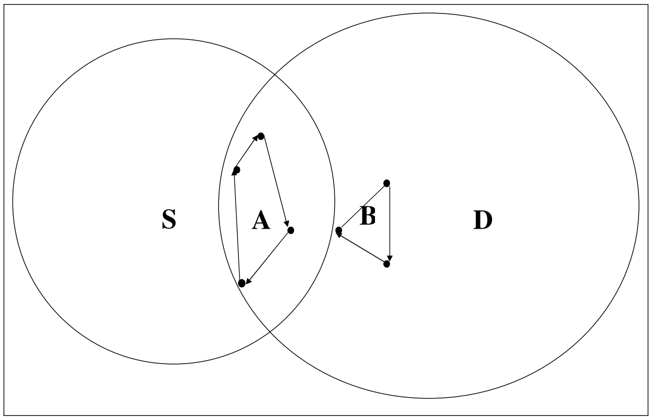

Fig. 5. Diagram showing the range and movement of nodes

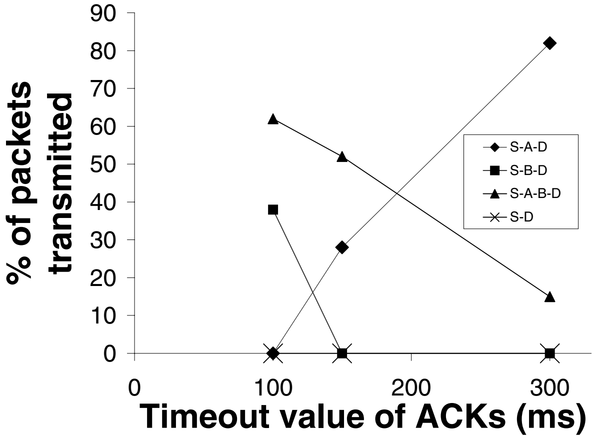

Fig. 6. Plot of percent of packets received vs. ACK timeouts for the different routes

demand routing protocol that is based on the concept of source routing. The protocol maintains a route cache in each node, which is updated as new routes are learnt. This is looked up every time a packet has to be sent, to determine if a route to the destination is present. If a route is not present it initiates a route discovery. Route maintenance is carried out in case the links in the route being used break due to channel fluctuations and node mobility causing the received power to fall below the receivers sensitivity threshold. Route error packets are generated at a node in case the link layer reports a broken link during data transmission. For a complete description of the DSR protocol please refer to [2]. For further test results or information about our energy aware DSR protocol please refer to [9], [10].

As a simple test of the testbed we measure the effect of route-level ACK timeouts on routing choice. Route level ACK timeouts generate route errors and DSR initiates a new route discovery process. We expect that channel variability will cause packet losses, MAC level retransmits, and variability in round trip times. The setup is shown in Fig. 5. Node *S* is sending to Node *D* and both are stationary.

Nodes *A* and *B* are moving in localized cycles in between *S* and *D*. The 4 nodes are placed so that the only possible routes are *S-A-D*, *S-B-D* and *S-A-B-D*. However node *B* is placed just outside the range of node *S*. Thus we expect that the route *S-B-D* will never be used. During the experiment nodes *A* and *B* are **moving** clockwise. These nodes stay at each point for 10 seconds. This movement is carried out repeatedly for 2 minutes. Because the nodes are moving none of the links in the network are stable. Such variability is seen due to **RF effects** such as multipath and fading.

We **repeat** the test with three different values of ACK timeout. Because all the nodes are easily **accessible** changing timeout values on them is quite trivial. The number of times each route is chosen was found out for different ACK timeouts. The results are shown in Fig. 6. This shows the percentage of times each route is chosen for each timeout value. This figure clearly shows that the number of routes chosen for lower timeout values is much more than the routes for longer timeout values.

The same topology was duplicated on the *ns* network simulator. In *ns* the route chosen was always *S-A-D*. *B* is never used during the entire test run. Thus in spite of the advantages, software simulators fail when dealing with the **physical layer** and an important qualitative effect, route oscillation with short timeouts is missed. These results clearly indicate the need for a wireless testbed.

IV. CONCLUSION AND FUTURE WORK

We have thus successfully demonstrated EWANT as an accessible wireless testbed. This testbed has realistic features and allows rapid development. Mobility, which was our main goal was successfully demonstrated. Although the testbed does not allow a gradual movement of nodes, this effect could be simulated by slowly “hopping” the node over short distances or by using variable attenuators.

Our aim here was not to test some specific code, but provide a level playing field for all wireless routing protocols. To achieve this goal we have designed a testbed which is independent of the underlying operating system or hardware used. Although we have implemented this testbed for 802.11b, it could easily be modified to suit similar wireless protocols.

Inspite of these advantages, there is still room for improvement. We feel a more elaborate central controller is needed. This must allow remote access to all nodes. This controller must also provide features for automating measurements, collecting data and synchronizing the start of experiments.

REFERENCES

[1] C. E. Perkins and E. M. Royer, “Adhoc on-demand distance vector routing,” in *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, Feb. 1999, pp. 90–100.
[2] D. Johnson, D. Maltz, Y.-C. Hu, and J. G. Jetcheva, “Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)ting,” IETF MANET group,” INTERNET DRAFT, November 2001.
[3] D. Maltz, J. Broch, and D. Johnson, “Experiences designing and building a multi-hop wireless ad hoc network testbed,” 1999. [Online]. Available: citeseer.nj.nec.com/maltz99experience.html

0-7803-7701-X/03/$17.00 (C) 2003 IEEE

[4] H. Lundgren, D. Nielsen, J. Nordstrom, and E. Tschudin, “A large-scale testbed for reproducible ad hoc protocol evaluations,” *IEEE WCNC*, 2002.

[5] Y. Zhang and W. Li, “An integrated environment for testing mobile ad-hoc networks,” in *Proceedings of the third ACM international symposium on Mobile ad hoc networking and computing*. ACM Press, 2002, pp. 104–111.

[6] J. Kaba and D. Raichle, “Testbed on a desktop: Strategies and techniques to support multi-hop manet routing protocol development,” 2001.

[7] “Cisco aironet 350 series client adapter data sheets,” available from `www.cisco.com`.

[8] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The click modular router,” *ACM Transactions on Computer Systems*, vol. 18, no. 3, pp. 263–297, August 2000.

[9] S. Doshi, S. Bhandare, and T. X. Brown, “An on-demand minimum energy routing protocol for a wireless ad hoc network,” *Mobile Computing and Communications Review*, vol. 6, pp. 50–66, 2002.

[10] S. Bhandare, S. Doshi, T. X. Brown, and S. Sanghani, “Comparison of two ad hoc wireless routing protocols on a hardware test-bed,” *IEEE WCNC*, March 2003.

[11] T. Rappaport, *Wireless Communications - Principles & Practice*, 2nd ed. Prentice-Hall, 2002.