

Comparison of MANET routing protocols using a scaled indoor wireless grid

David Johnson and Albert Lysko, Student Member, *IEEE*

Abstract—Predicting the performance of ad-hoc networking protocols has typically been performed by making use of software based simulation tools. Experimental study and validation of such predictions is vital to obtaining more realistic results, but may not be possible under the constrained environment of network simulators. This paper presents experimental comparisons of routing protocols using a 7 by 7 grid of closely spaced WiFi nodes. It firstly demonstrates the usefulness of the grid in its ability to emulate a real world multi-hop ad-hoc network. It specifically compares hop count, routing traffic overhead, throughput, delay and packet loss for three protocols which are listed by the IETF MANET working group. These are AODV, OLSR and DYMO

Index Terms—ad-hoc networks , IEEE 802.11 standard, wireless grid testbed

I. INTRODUCTION

One of the key challenges for researchers in the field of wireless networking protocol design, is the ability to carry out reliable performance measurements on their protocol. Parameters that need to be evaluated are typically, scalability, delay and throughput, network convergence in the presence of rapidly changing link quality and route optimization.

Unfortunately most of the work done so far makes use of simulations which over simplify the physical layer and even aspects of the Medium Access Control layer. There is also a lack of consistency between the results of the same protocol being run on different simulation packages [1].

Mathematical models are useful in the interpretation of the effects of various network parameters on performance. For example Gupta and Kumar [2] have created an equation which models the best- and worst-case data rate in a network with shared channel access, as the number of hops increases. However recent work done by the same authors [3] using a real testbed, employing laptops equipped with IEEE 802.11 based radios, revealed that 802.11 multi-hop throughput is still far from even the worst case theoretical data rate predictions.

After 6 hops, the mathematical model used by Gupta and Kumar shows a 30% decrease in throughput whereas the testbed they used showed a 95% decrease in throughput. This illustrates the importance of verification using physical testbeds.

A recent Network Testbeds workshop report [4] highlighted the importance of physical wireless testbed facilities for the

research community in view of the limitations of available simulation methodologies. This was the motivation for the ORBIT project [5] which describes a wireless grid similar to the one that will be discussed in this paper.

The ORBIT mesh lab has an 8x8 grid and a 20x20 grid, which makes use of 802.11 wireless equipment based on the same Atheros chipset used at Meraka. The ORBIT laboratory makes use of Additive White Gaussian Noise (AWGN) to raise the noise floor, while Meraka makes use of attenuators. It allows researchers anywhere in the world to run an experiment on the lab by making use of a scheduler. Researchers can change everything from the routing protocol to the entire operating system that will be run on the nodes.

These mini scale wireless grids can emulate real world physical networks due to the inverse square law of radio propagation, by which the electric field strength will be attenuated by 6.02 dB for each doubling of the distance in any free-space loss scenario.

Most of the indoor testbeds, such as the one used by Microsoft's Research lab [6], have been created by placing computers with wireless cards in offices and relying on the walls of the building structure to attenuate the signal sufficiently to create a multi-hop environment. Although these setups have been useful, they generate results that will be very difficult to repeat and verify due to the chaotic nature of signal propagation in an office environment.

In this paper, experimental comparisons of the performance of routing protocols, using a 7x7 grid of closely spaced WiFi nodes, are presented. The usefulness of the grid in its ability to emulate a real world multi-hop ad-hoc network is demonstrated, comparing hop count, routing traffic overhead, throughput, delay and packet loss for three protocols which are listed by the IETF MANET working group, *viz* AODV, OLSR and DYMO.

II. BACKGROUND ON AD-HOC NETWORKING PROTOCOLS USED

An Ad hoc network is the cooperative engagement of a collection of wireless nodes without the required intervention of any centralized access point or existing infrastructure. Ad hoc networks have the following key features: they are self-forming, self-healing and do not rely on the centralized services of any particular node.

The IETF Mobile Ad-hoc Networks (MANET) working group oversees the process of standardizing IP routing protocols for wireless ad hoc networks within both static and dynamic topologies. Wireless link interfaces have some unique routing interface characteristics and that node topologies within a wireless routing region may experience increased dynamics, due to motion or other environmental factors.

As a consequence three main categories of ad-hoc routing protocols have surfaced over the past decade, these are reactive routing protocols, proactive routing protocols and hybrid routing protocols. This paper only concerns itself with reactive and proactive routing.

A. Pro-active routing protocols

Pro-active or table-driven routing protocols maintain fresh lists of destinations and their routes by distributing routing tables in the network periodically. The advantage of these protocols is that a route is immediately available when data needs to be sent to a particular destination. The disadvantage of this method is that unnecessary routing traffic is generated for routes that may never be used. The Pro-active routing protocol that this paper will investigate on the testbed is called Optimized Link State Routing (OLSR) [7]

OLSR reduces the overhead of flooding link state information by requiring fewer nodes to forward the information. A broadcast from node X is only forwarded by its multi point relays. Multi point relays of node X are its neighbors such that each two-hop neighbor of X is a one-hop neighbor of at least one multi point relay of X. Each node transmits its neighbor list in periodic beacons, so that all nodes can know their 2-hop neighbors, in order to choose the multi point relays (MPR).

Fig. 1 illustrates how the OLSR routing protocol will disseminate routing messages from node 3 through the network via selected MPRs.

The RFC for OLSR makes use of hysteresis to calculate the link quality between nodes in order to stabilize the network in the presence of many alternative routes. A new improved routing metric, the Expected Transmission Count (ETX) [8], has also been incorporated into the source code for OLSR but it is not formally part of the RFC. All the MANET RFC's prefer to use hop count as a routing metric for the sake of simplicity.

Link hysteresis is calculated using an iterative process. If q_n is the link quality after n packets and h is the hysteresis scaling constant between 0 and 1 then the received the link quality is defined as:




Fig. 1. OLSR routing protocol showing selection of MPRs

$$q_n = (1 - h)q_{n-1} + h \quad (1)$$

For each consecutive unsuccessful packet the link quality is defined as:

$$q_n = (1 - h)q_{n-1} \quad (2)$$

When the link quality exceeds a certain high hysteresis threshold, $q_{HYST_THRESH_HIGH}$, the link is considered as established and when the link quality falls below a certain low hysteresis threshold, $q_{HYST_THRESH_LOW}$, the link is dropped.

Fig. 2 shows a graph for 7 consecutive successful packets followed by 7 unsuccessful packets with $h = 0.5$, $q_{HYST_THRESH_HIGH} = 0.8$ and $q_{HYST_THRESH_LOW} = 0.3$.

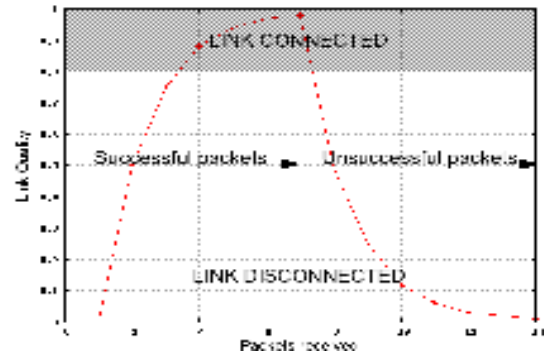


Fig. 2. Link Hysteresis in the OLSR routing protocol

Hysteresis produces an exponentially smoothed moving average of the transmission success rate and the condition for considering a link established is stricter than the condition for dropping a link.

The alternative metric, Expected Transmission Count (ETX) [8], calculates the expected number of retransmission that are required for a packet to travel to and from a destination. The link quality, LQ , is the fraction of successful packets that were received by us from a neighbor within a window period. The neighbor link quality, NLQ , is the fraction of successful packets that were received by a neighbor node from the sender within a window period. Based on this the ETX is calculated as follows:

$$ETX = \frac{1}{(LQ \times NLQ)} \quad (3)$$

In a multi-hop link the ETX values of each hop are added together to calculate the ETX for the complete link including all the hops.

Fig. 3 shows the ETX values for 7 consecutive successful packets followed by 7 consecutive unsuccessful packets assuming a perfectly symmetrical link and a link quality window size of 7.

A perfect link is achieved when ETX is equal to 1. ETX has




Fig. 3: ETX Path metric values for successive successful and unsuccessful packets

the added advantage of being able to account for asymmetry in a link as it calculates the quality of the link in both directions. Unlike Hysteresis ETX improves and degrades at the same rate when successful and unsuccessful packets are received respectively. OLSR with ETX will always choose a route with the lowest ETX value.

B. Reactive or on-demand routing protocols

Reactive or on-demand protocols find routes on demand by flooding the network with Route Request packets. This allows only the routes that the network needs to be entered into a routing table. The disadvantage of this method is that there will be a startup delay when data needs to be sent to a destination to allow the protocol to discover a route. The two reactive protocols will be investigated in this paper are Ad hoc On-demand Distance Vector (AODV) [9] routing and its recent successor called Dynamic Manet On-demand Routing (DYMO) [10].

AODV employs destination sequence numbers to identify recent and up to date paths. Source node and intermediate nodes only store the next-hop information corresponding to each flow for a data packet transmission. A node will update its path information only if the destination sequence number of the current packet received is greater than the last destination

sequence number stored at the node.

If an intermediate node already has a valid route to a destination it will send a gratuitous route reply otherwise it forwards the route request. Route errors are determined using periodic beacons to detect link failures. Link failures cause a route error message to be sent to the source and destination nodes.

Fig. 4 shows AODV discovering a route from node 1 to node 10 using Route Requests (RREQ) and Route Replies (RREP).




Fig. 4. AODV routing protocol showing the route discovery process

DYMO is the most recent ad hoc networking protocol proposed by the MANET working group. It seeks to combine advantages of reactive protocols, AODV and DSR together with some link state features of OLSR. It makes use of the path accumulation feature of Dynamic Source Routing (DSR) by adding the accumulated route, back to the source, to the Route Request packet. It retains the destination sequence number feature of AODV but HELLO packets are an optional feature and are normally left out by default. It also does away with the gratuitous RREP feature of AODV. Routing information is kept up to date by expiring unused routes after a specific time interval. DYMO is also able to make use of periodic beacons to monitor link status and send route errors when failures occur.

Fig. 5 shows how DYMO creates the full path back to node 1 in the routing packet as the RREQ is forwarded towards the destination node 10. The RREP is sent back along this accumulated path.

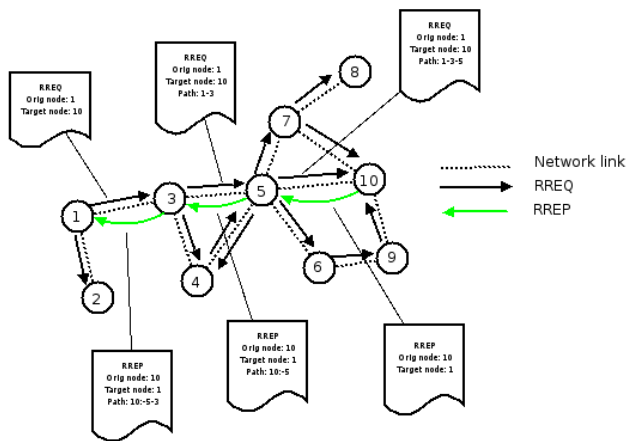


Fig. 5. DYMO routing protocol showing how path accumulation is used during route discovery

III. LINUX IMPLEMENTATIONS OF AD-HOC NETWORKING PROTOCOLS

A crucial part of comparing a set of ad-hoc networking protocols on a physical testbed is finding implementations of the protocol that are well written and are as close as possible to the original published RFC.

Currently there are approximately 110 known ad hoc routing protocols that are widely known and of these only approximately 14 have an implementation which can run on a physical network. There are however many more which have implementations, which can run in a simulation environment such as NS2. All the MANET protocols have been implemented on a UNIX platform. AODV has 10 implementations, OLSR has 7, DSR has 4, DYMO has 2 and TBRPF has 1.

The choice between a multitude of implementations of the same protocol was based on whether the particular implementation claimed to be RFC compliant, and if there was a strong developer community supporting the code base. Preference was also given to cases where the same code base was used for simulations and running the code on a physical network as this would make future comparisons of simulations and physical network results very simple.

The following implementations were chosen for the protocols used on the testbed.

1. For OLSR, the implementation developed by Tønnesen [11] was used. This implementation is commonly called `olsr.org` and is now part of the largest open source ad hoc networking development. Version 0.4.10 is used in the massive mesh version. This implementation is RFC3626 compliant and is capable of using the standard RFC link hysteresis

metric or the new ETX metric for calculating optimal routes. All parameters mentioned in the RFC are implemented and can be modified through a configuration file.

2. For AODV, the implementation by Nordström was used [12]. This implementation is designated AODV-UU and the current version used in the test bed is 0.9.3. The code is compliant with the AODV RFC3561 standard. This code base also supports the use of the same C code to run NS2 simulations. All parameters mentioned in the RFC are implemented and can be modified by changing the source code.
3. For DYMO, the implementation by Ros [13] was used. This implementation is named DYMOUM and was developed out of the AODV-UU code base. The current version being used in the test bed is 0.3 and the code claims full compliance with the Internet draft version 5 of DYMO. All parameters mentioned in the Internet draft are implemented and can be modified by changing the source code.

All the implemented routing protocols were used with their default RFC suggested configuration parameters.

IV. BUILDING THE MESH TESTBED

A. Physical construction of the 7x7 grid

The mesh testbed consists of a wireless 7x7 grid of 49 nodes, which was built in a 6x12 m room as shown in Fig. 6. A grid was chosen as the logical topology of the wireless testbed due to its ability to create a fully connected dense mesh network and the possibility of creating a large variety of other topologies by selectively switching on particular nodes as shown in Fig. 7.

Each node in the mesh consists of a VIA 800 C3 800MHz motherboard with 128MB of RAM and a Wistron CM9 mini PCI Atheros 5213 based WiFi card with 802.11a/b/g capability. For future mobility measurements, a Lego Mindstorms robot with a battery powered Soekris motherboard containing an 802.11a (5.8GHz) card and an 802.11b/g (2.4GHz) can be used.




Fig. 6. The architecture of the mesh lab. Ethernet is used as a back channel to connect all the nodes to a central server through a switch. Each node is also equipped with an 802.11 network interface card.




Fig. 7. Layout of the 7 by 7 grid of WiFi enabled computers, the line following robot is an option which will be explored in the future to test mobility in a mesh network.

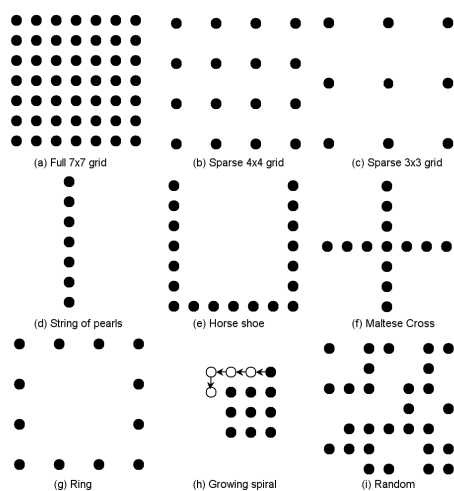


Fig. 8. Various topologies that can be tested on the 7x7 grid, diagrams (a) to (c) demonstrate various levels of density in a grid, diagram (e) is used to create a long chain to force routing protocols to use the longest multi-hop route, and diagram (g) is used to test route optimization

Every node was connected to a 100Mbit back haul Ethernet network through a switch to a central server as shown in Fig. 8. This allows nodes to use a combination of a Preboot Execution Environment (PXE), built into most BIOS firmware, to boot the kernel and a Network File System (NFS) to load the file system.

The physical constraints of the room, with the shortest length being 7m, meant that the grid spacing needed to be about 800 mm to comfortably fit all the PC's within the room dimensions.

At each node, an antenna with 5dBi gain is connected to the wireless network adapter via a 30 dB attenuator. This introduces a path loss of 60dB between the sending node and the receiving node.

Reducing the radio signal to force a multi-hop environment is core to the success of this wireless grid. The wireless NIC's

that are being used in this grid have a wide range of options that can be configured.

The output power level can be set from 0dBm up to 19dBm. 802.11g and 802.11b modes are available in the 2.4GHz range. 802.11b allows the sending rate to be set between 1Mbps and 11Mbps and 802.11g allows between 6Mbps and 54Mbps. The receive sensitivity of the radio, which is the level above which it is able to successfully decode a transmission, depends on the mode and rate being set. The faster the rate, the lower the receive sensitivity threshold.

Fig. 9 shows free space loss curves over the distance of the grid to illustrate what the received signal will be at any particular node. This Fig. also shows the receive sensitivity of the radio at various modes and data rates. In theory, where the curve line rises above the horizontal lines, there will be connectivity but as will be seen later, there are other factors other than free space loss which affect the signal propagation.

This network was operated at 2.4GHz due to the availability of antennas and attenuators at that frequency, but in future the laboratory will be migrated to the 5GHz range which has many more available channels with a far lower probability of being affected by interference.

B. Electromagnetic modeling

In order to understand the stochastic behavior of the wireless nodes in the grid, the underlying electromagnetic properties need to be understood..

The test-bed was modeled using numerical electromagnetic (EM) modeling, based on the method of moments [14]. This modeling was used to obtain the values of the coupling coefficients (scattering matrix elements) between nodes S_{ij} , where $i \neq j$, $i, j = 1, N$, and N , is the total number of nodes in the test-bed.

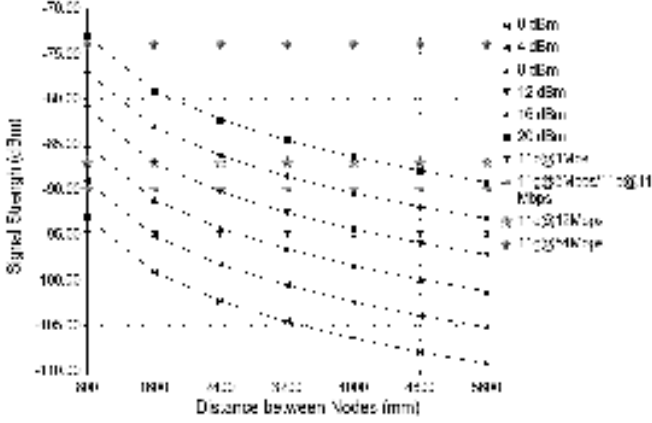


Fig. 9. Received signal strength vs. distance between nodes in the grid spaced 800mm apart. The horizontal lines show the receive sensitivity of the Atheros 5213 wireless network card, if the received signal strength curve is above this line, there will be connectivity between the nodes.

The elements of the scattering matrix are defined as follows [14]:

$$S_{ij} = \left. \frac{V_i^-}{V_j^+} \right|_{V_k^+ = 0 \text{ for } k \neq j} \quad (4)$$

where V_k^+ is the amplitude of the voltage wave incident on port k (port of the antenna at k^{th} node) and V_n^- is the amplitude of the voltage wave reflected from port n (port of the antenna at n^{th} node), while the incident waves on all ports except the j^{th} port are set to zero. Effectively, S_{ij} is the transmission coefficient from port j to port i when all other ports are terminated in matched loads.

The single node model consists of a rectangular metallic PC case and antenna. The antenna is a typical 5 dB gain dipole antenna supplied with many wireless cards. Both measurements and numerical EM modeling confirmed that this antenna has a reasonably flat gain and is well matched in the operating frequency range of 2.4-2.5 GHz.

The EM modeling showed that, for a single node, the presence of the PC case changes the effective horizontal plane radiation pattern from omni-directional to a more complex pattern. The maximum variation from the omni-directional gain pattern was found to be 1.5 dB. This effect is due to close proximity of the PC working as an offset reflector.

Once the nodes are assembled into an array, the effective radiation patterns of individual nodes become even more distorted, with dependence on the position in the array; it also manifests itself in deviation from the line-of-sight free-space propagation loss.

In the case of a linear 1 x 7 array with 0.8 m inter-node spacing, dependence on position was found to be negligible

(within 0.3 dB). However, for a rectangular, 7x7 array, the effect of arraying became much stronger, with variations in signal strength of up to 3 dB.

It was also found that the attenuation of the signal propagating from one node to another was dependent on the direction of propagation. This anisotropy is due to the antennas installed closely to the PC cases, and can be explained in the Fresnel zone terms.

The boundaries for the Fresnel interference zones can be calculated for any two nodes in the closely spaced rectangular grid. It is clear that as the nodes are chosen further apart, the number of PC cases that can possibly lie within the first Fresnel zone, increases, with concomitant increase in interference.

It was also found that the propagation is affected by the specific position of the PC cases associated with the nodes in the test bed. In one direction the wide sides of the cases are presented, while in an orthogonal direction, the narrow sides, with the antennas partially obscured, are presented. This can affect the signal strength by as much as 1.5 dB.

Experimental tests were run on the test bed by measuring the Received Signal Strength Indicator (RSSI) value between all possible pairs of nodes, while keeping all other nodes in the network switched off. RSSI is a measure of the magnitude (not necessarily the quality) of the received signal strength in a wireless environment, in arbitrary units. It is used internally in a wireless networking card to determine when the signal is below a certain threshold at which point the network card is clear to send (CTS).

Measured values of RSSI versus distance for two models of transmitter and computer (one with cases and one without) are shown in Fig. 10.

The 7x7 test bed was simulated numerically for the case where only the antennas (no PC cases) were included. It was found that a strong correlation existed with the case where the PC cases were included in the simulation. However, the simulation including PC cases shows better agreement with the experimental data for long distances. At the shortest distance between the neighboring nodes, when there is no obstruction between the nodes, the results of two simulations match.

The boundaries of the mean values of RSSI, shown in Fig. 10, show variation in the coupling for nodes with the same separation. In practice, the signal strength between two pairs of nodes, both being separated at the same distance, may vary by as much as 10dB.




Fig. 10. Received signal strength indicator (rssi) value versus distance between nodes - measured and simulated results for a rectangular 7×7 testbed. Crosses define the standard deviation-based range of rssi with respect to mean values shown with circles, diamonds and dots.

These variations must be taken into consideration in later experiments with ad hoc routing protocols where routing paths will vary between short and long hops due to these signal strength fluctuations.

C. Challenges

The following defines specific challenges that were encountered while trying to obtain meaningful results from the wireless grid.

1) Complexity and density of grid

The mesh grid forms a highly connected dense graph which creates a difficult optimization problem for a routing algorithm. In a full 7×7 grid routing algorithms will be presented with thousands of equivalent hop length routes, OLSR using ETX will constantly be receiving new routes with changing ETX values.

2) Communication Grey zones

Communication gray zones [15] occur because a node can hear broadcast packets, as these are sent at very low data rates, but no data communication can occur back to the source node, as this occurs at a higher data rate. Fig. 11 shows how a RREQ can be broadcast to the edge of the communication gray zone but the RREP cannot get back to the source node. This problem was solved by locking the broadcast or multi cast rate to the data rate.

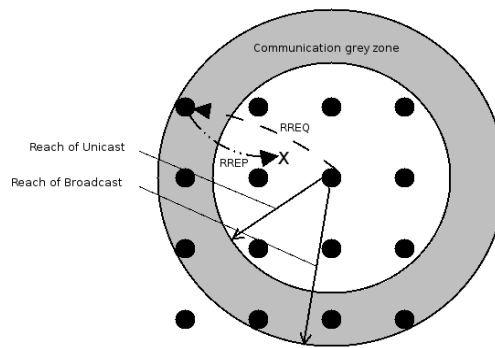


Fig. 11. Communication gray zones

3) Hardware issues

There are many physical hardware problems that one has to deal with such as faulty wireless NICs and non-uniformity of the receive sensitivity of the cards. These have been characterized in the section on electromagnetic modeling

4) Routing protocol bugs

Both AODV and DYMO gave kernel errors when the network size was greater than approximately 20 nodes. This caused the routing algorithms to freeze and not allow any packets to enter or exit the wireless interface. The particular error complained that the maximum list length had been reached. This constant was increased in the source code and subsequently the bug disappeared, but it did confirm the fact that these protocols had not been tested on networks as large as this testbed.

5) Antenna dual diversity

It was found that when dual diversity was switched on, the nodes became very unpredictable.

6) Wide choice of wireless card parameters

Finding the best combination of communication mode, data rate and transmit power was not a simple process. Using Fig. 9 gave some direction, but only trial and error eventually helped converge the settings to using 802.11b mode, an 11Mbps data rate and using a power level ranging from 0 to 8dBm.

7) Time consuming experiments

Experiments were very time consuming. Testing the throughput and delay for all permutation pairs of 49 nodes in the grid for 4 routing protocols using a 20 second test time takes approximately 52 hours.

8) Interference

Finding a channel in 2.4 GHz which is relatively free from interference is not easy. The building where the experiments were conducted has an extensive in-building wireless network operating on 2.4GHz. Even relatively weak signals close to -90 dBm are a problem when you are using a highly attenuated lab. In the future the lab will be migrated to 5.8GHz which has far more available channels.

V. MEASUREMENT PROCESS

All measurements other than throughput tests were carried out using standard Unix tools available to users as part of the operating system. The measurement values were sent back to the server via the nodes Ethernet ports of the nodes and therefore had no influence on the experiments that were being run on the wireless interface.

It was found that the lab provides the best multi-hop characteristics trade off with the best delay and throughput when the radios are configured with the following settings.

Channel = 6
 Mode = 802.11b
 Data rate = 11Mbps
 TX power < 8dBm

The following processes were used for each of the metrics being measured:

1) Delay

Standard 84 byte ping backs were sent for a period of 10 seconds. The ping reports the round trip time as well as the standard deviation.

2) Packet loss

The ping tool also reports the amount of packet loss that occurred over the duration of the ping test

3) Static Number of hops for a route to a destination

The routing table reports the number of hops as a routing metric.

4) Round trip route taken by a specific packet

The ping tool has an option to record the round trip route taken by an ICMP packet but unfortunately the IP header is only large enough for nine hops. This sufficed for most of the tests that were done but occasionally there were some routes which exceeded 9 round trip hops and no knowledge of the full routing path could be extracted in these instances.

5) Throughput

The tool *iperf* [16] was used for throughput measurements. It uses a client server model to determine the maximum bandwidth available in a link using a TCP throughput test but can also support UDP tests with packet loss and jitter. For these experiments an 8K read write buffer size was used and throughput tests were performed using TCP for 10 seconds. UDP could be considered a better choice as it measures the raw throughput of the link without the extra complexity of contention windows in TCP. This does make the measurement more complex, however; as no prior knowledge exists for the link and the decision on the test transmission speed is done by trial and error.

6) Routing traffic overhead

In order to observe routing traffic overhead the standard Unix packet sniffing tool *tcpdump* was used. A filter was used on the specific port that was being used by the routing protocol. The tool made it possible to see the number of routing packets leaving and entering the nodes as well as the size of these routing packets.

To force dynamic routing protocols such as AODV and DYMO to generate traffic to establish a route a ping was always carried out between the farthest two points in the network.

7) Growing network size

When tests are done which compare a specific feature to the growing number of nodes in the network, a growing spiral topology, shown in Fig. 12, starting from the center of the grid is used. This helps to create a balanced growth pattern in terms of distances to the edge walls and grid edges, which may have an electromagnetic effect on the nodes.

8) Testing all node pairs in the network

When throughput and delay tests were carried out on a fixed size topology, all possible combinations of nodes were tested. If the full 7x7 grid was used, this equates to 2352 (49x48) combinations.

VI. RESULTS

A. Testing a long chain with fixed routing tables

In order to establish the baseline for performance of the wireless nodes in the grid, it is useful to remove any effects of routing and establish the best possible multi-hop throughput and delay between the nodes.

Fig. 13 shows a string of pearls 49 nodes long built by creating a zigzag topology in the grid using manually configured static routes.




Fig. 13. Growing spiral topology used for test which compares a metric against a growing network size. The arrows are used to show the sequence in which nodes are turned on. Solid circles represent nodes which are turned on and dotted-circles represent nodes which will be turned on in the future.




Fig. 12. Creation of a string of pearls topology 49 nodes long using the 7x7 grid

All the radios were set to maximum power (20dBm), using 802.11b mode with a data rate of 11Mbps to avoid any packet loss.

Theoretical work done by Gupta and Kumar calculated the best- and worst-case throughput for a node n hops away where all radios share a single channel and are all within transmission range of each other. From [2]:

$$\lambda_{WORST}(n) = \left(\frac{W}{\sqrt{n \log(n)}} \right) \quad (5)$$

$$\lambda_{WORST}(n) = \left(\frac{W}{\sqrt{n}} \right) \quad (6)$$

where W = Bandwidth of first hop

However a recent study [3] by the Gupta and Kumar using laptops equipped with 802.11 based radios placed in offices revealed, using a least-squares fit, that the actual data rate versus the number of hops is

$$\lambda(n) = \left(\frac{W}{n^{1.68}} \right) \quad (7)$$

This represents a dramatic difference in throughput after a multiple number of hops for 802.11 compared to the theoretical predictions. After 10 hops the measured results were as much as 10% of the theoretical worst-case prediction.

Fig. 14 compares 7x7 grid multi-hop throughput to theoretical and measured results. The measurements revealed a less pessimistic result but one which was still less than the worst-case theoretical results.

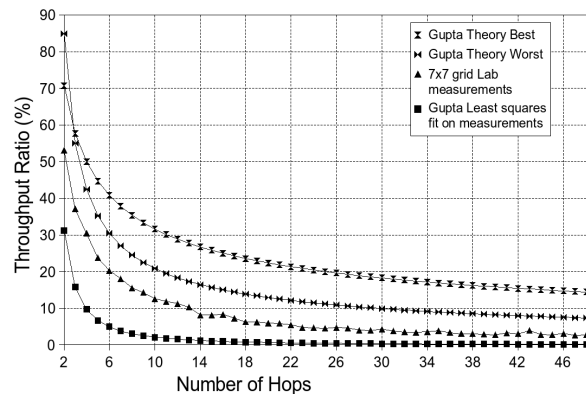


Fig. 14. Comparison 7x7 grid multi-hop throughput to theoretical other measured results

Carrying out a least squares fit on the results obtained with the testbed and using a plot of the log of both the x and y-axis reveals the following function for TCP throughput under ideal conditions for the grid.

$$\lambda(n) = \left(\frac{W}{n^{0.98}} \right) \quad (8)$$

Fig. 15 shows how the delay increases as the number of hops increases. It follows a basic linear progression with increasing standard deviation.




Fig. 15. Increasing delay with standard deviation for a 49-node string of pearls n the 7x7 wireless grid

B. Hop count distribution

The ability to create a multi-hop network in the mesh testbed is a key measure of the ability of the lab to emulate a real world wireless mesh network.

A basic test using the OLSR routing protocol with ETX as a routing was configured using a growing spiral topology as described in section V. A topology depicted in Fig. 16 at the end of the paper was created. The values in the graph are the ETX values for a node pair.

Fig. 16 shows the total number of routes in specific hop categories versus a growing number of nodes in the grid.

Up to 5 hop links were achieved with 2 hop links forming the dominant category after 16 nodes. This demonstrates that there is enough attenuation in the wireless grid to form a range of multi-hop links that one would find in a real-world network.




Fig. 16. Total number of links with a specific total number of hops increasing as the number of nodes increases




Fig. 17. Average number of hops versus distance for full 7x7 grid between all 2352 possible pairs

The tendency of a routing protocol to choose a longer or shorter path depends on the strategy of the routing algorithm.

For example AODV tries to minimize hop count whereas OLSR-ETX tries to minimize packet loss. Fig. 17 shows a comparison of AODV, DYMO, OLSR-RFC and OLSR-ETX in terms of average hop count versus distance.

This experiment was carried out using a full 7x7 grid with a test carried out between every possible pair of nodes in the grid. The set of all possible pairs is equal to $49 \times 48 = 2352$.

Average hop count values based on data shown in Fig. 17 are as follows: AODV: 1.98 hops, DYMO 3.04 hops, OLSR-ETX: 2.21 hops and OLSR-RFC: 2.95 hops. From the graph in Fig. 17 and these averages one can see that AODV is trying to minimize hop count. OLSR-RFC tends to use more hops as links with long distances between them tend to be penalized by its steep downward hysteresis curve when packets are dropped (see Section II). DYMO picks the first possible route it can obtain and doesn't try to continuously optimize for shorter hop links. OLSR-ETX has decided that shorter hops are better in the grid in terms of minimizing packet loss.

C. Routing traffic overhead

The ability of a routing protocol to scale to large networks is highly dependent on its ability to control routing traffic overhead. The following graphs show the results of measuring routing traffic as the network size grows in a growing spiral as described in section V.

Fig. 18 shows that outbound OLSR traffic rapidly increases but then begins to level off after about 25 nodes due to the multi point relays limiting router traffic forwarding.

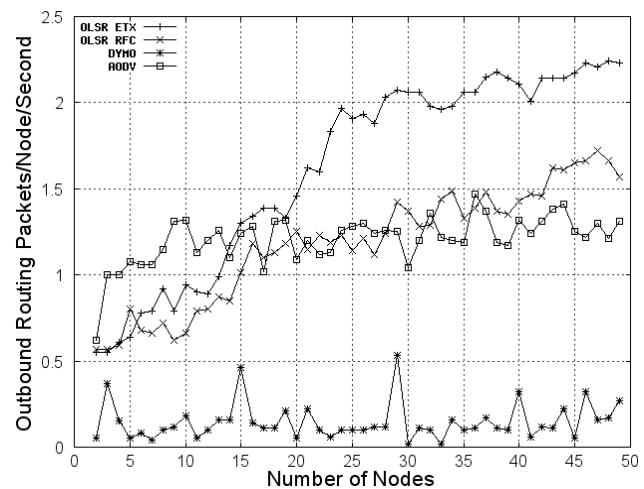


Fig. 18. Outbound routing packets per node per second versus increasing number of nodes using a growing spiral




Fig. 19. Inbound routing packets per node per second versus increasing number of nodes using a growing spiral

Outbound routing traffic should always be less than the inbound traffic as the routing algorithm makes a decision to rebroadcast the packet or not and Fig. 19 confirms this.

Outbound routing traffic consists of data packets leaving a node to communicate some routing information to other nodes. Inbound routing traffic consists of data packets entering a node in order for it to learn routing information from other nodes.

DYMO shows the least amount of routing traffic due to its lack of HELLO packets. This is also due to no further routing packets being transmitted once it has found a route to a destination. The occasional spikes in the routing traffic are for cases where it took longer than normal to establish a route.

Fig. 20 shows how routing packet lengths grow as the number of nodes increase. This is another important characteristic to analyze if a routing protocol is to scale to large networks.

As the network grows, OLSR needs to send the entire route topology in Topology Control (TC) update messages, which helps explain this steady linear increase with the number of nodes. OLSR with the ETX extension uses a longer packet length due to the extra overhead of carrying link quality metrics.




Fig. 20. Average Routing Packet length growth versus increasing number of nodes

AODV does not carry any route topology information in its packets, which explains its extremely small packet length which stays constant. DYMO makes use of path accumulation which explains its steady increase in size relative to the number of hops between the two furthest points on the network.

D. Throughput, packet loss and delay measurements

The ability of a routing algorithm to find an optimal route in the grid will be exposed by its throughput and delay measurements.

In order to evaluate their performance a series of tests were done with increasing complexity. The simplest starting case is to test routing performance for a simple string of 7 pearls, followed by three adjacent 7-node columns and finally the full 7x7 grid

Results for a string of pearls 7 nodes long.

Table 1 summarizes the results for all 42 possible pairs

	Forward HOP count	Route changes	Packet loss	Delay	Delay(stddev)	TP	No link
AODV	1.33	0.43	11.19	37.24	116.64	2723.36	1
DYMO	1.52	0	9.52	3.65	2.37	2907.67	0
OLSR_ETX	1.43	0.1	8.57	27.56	101.91	2730.69	0
OLSR_RFC	1.67	0.76	2.14	5.35	5.35	2923.64	0

Table 1. Comparison of throughput, delay and packet loss for a 7 node string of pearls topology

OLSR_RFC had the highest number of route changes and forward hops over the 10-second measurement period but had the best average throughput. The route changes, therefore, must have converged the link towards a more optimal route. DYMO achieved the best performance in terms of delay. Only AODV had 1 case where the routing algorithm could not establish a link.

Fig. 21 shows the cumulative distribution function for all possible 42 pairs.

The graphs are very similar except for the fact that AODV and OLSR-ETX have approximately 20% of their links unable to achieve any throughput in the 10 seconds that they were tested. There are also clearly noticeable discrete clusters of throughput categories around approximately 2000 KB/s and 4200 KB/s, this is due to discrete collections of single or multi-hop routes.




Fig. 21. Throughput CDF for 7 node strong of pearls

Results for 3 adjacent columns of 7 nodes (21 nodes)

The complexity is now increased somewhat, where the routing algorithms can begin to choose between many alternative routes.

Table 2 summarizes the results for all 420 possible pairs.

	Forward HOP count	Route changes	Packet loss	Delay	Delay(stddev)	TP	No link
AODV	0.97	0.46	43.62	148.17	648.86	1245.23	126
DYMO	1.54	0.09	26.88	58.41	126.02	1701.69	56
OLSR_ETX	1.28	0.08	24.05	38.92	120.47	2899.34	68
OLSR_RFC	1.9	1.66	8.76	34.57	94.13	2113.12	20

Table 2. Comparison of throughput, delay and packet loss for 3 adjacent 7 node columns

AODV was clearly the worst performer in terms of number of failed links, average throughput and average delay. OLSR with ETX achieved the best average throughput with a very low number of route changes, whereas OLSR RFC achieved the best delay with a relatively high number of route changes, an average of 1.66 changes in the 10-second test period.




Fig. 22. Throughput CDF for 3 adjacent 7 node columns

Fig. 22 shows how AODV had a 50% route failure rate when carrying out throughput tests. OLSR-RFC had the lowest route failure but most of the throughput was clustered in the lower

range, probably due to non-optimal high hop count. OLSR-ETX had a strong clustering in the upper (>4000KB/s) region.

Results for full 7x7 grid (49 nodes)

The entire grid is now used to understand how the routing protocols perform with the maximum complexity available.

Table 3 summarizes the results for all 2352 possible pairs. AODV was clearly the weakest protocol in this scenario, with more than half the links achieving no route at all. All the other protocols performance metrics were very close. On the whole OLSR-RFC was marginally better than the rest, achieving the top average throughput rate of 1330 KB/s.

	Forward HOP count	Route changes	Packet loss	Delay	Delay(stddev)	TP	No link
AODV	1.36	0.53	71.22	117.87	317.35	773.33	1425
DYMO	2.2	0.11	32.81	64.72	150.2	1165.66	413
OLSR_ETX	1.84	0.25	24.05	68.84	247.78	1187.57	453
OLSR_RFC	2.28	2.34	22.22	67.44	132.49	1330.05	381

Table 3. Comparison of throughput, delay and packet loss for 7x7 grid

Fig. 23 shows a far flatter throughput performance compared to the previous network experiments. AODV had close to 80% of its links unable to achieve any throughput whereas the rest were all around 40%.

These results demonstrate how network performance quickly degrades for all routing protocols as the network size and complexity increases.




Fig. 23. Throughput CDF for 7x7 grid

Comparison of throughput results against baseline

Finally Fig. 24 shows how the routing protocols performance compares to the ideal multi-hop network that was set up in section VI-A. AODV could not be compared due to 80% of the links failing to achieve any throughput.

This graph demonstrates how routing overhead, route flapping and non-optimal routes all contribute towards decreasing the throughput of all three routing protocols. The baseline presents the best possible throughput the routing protocols could

achieve which will be asymptotically more difficult to reach, the closer you get. OLSR-RFC performed the best and came within an average of 76% of the baseline.




Fig. 24. Throughput performance of routing protocols against theoretical and baseline measurements

VII. CONCLUSION

The results from experiments done in the wireless grid lab have shown that it is possible to build a scaled wireless grid which yields good multi-hop characteristics. Currently hop counts up to 5 are achievable with routing protocols in the full 7x7 grid when the power is set to 0dBm with 30 dB attenuators.

A grid structure does yield a worst-case complexity problem for routing protocols in terms of the number of alternative routes available between distant points in the grid. This has a severe impact on route flapping if some kind of damping is not employed.

The AODV protocol showed the weakest performance in the grid with close to 60% of possible link pairs achieving no route for the full 7x7 grid. DYMO showed good results for its low routing overhead with the least amount of delay for the full 7x7 grid and the 2nd best throughput performance in a simple string of pearls topology.

The RFC version of OLSR had the best overall performance in a gull 7x7 grid in terms of throughput achieved and successful routes but OLSR with the ETX extension performed better in medium size networks of about 21 nodes.

All these performance tests were carried out using suggested configuration parameters that are published in MANET RFC's and internet drafts; in the future it will be interesting to see how performance can be tweaked for specific topologies by changing parameters such as HELLO intervals. Some degree of node mobility and network load will also be the domain for future measurements in the wireless grid.

ACKNOWLEDGMENT

The authors would like to thank Professor J.A.G. Malherbe at University of Pretoria for his valuable comments on this paper and advice on research writing style.

REFERENCES

- [1] T. R. Andel and A. Yasinac, "On the Credibility of Manet Simulations", IEEE Computer Society, July 2006.
- [2] P. Gupta and P. R. Kumar, "The capacity of wireless networks," IEEE Trans. Inf. Theory, vol. 46, no. 2, pp. 388–404, Mar. 2000.
- [3] P. Gupta and R. Drag, "An Experimental Scaling law for Ad Hoc Networks", Bell Laboratories technical report, 2006.
- [4] NSF Workshop on Network Research Testbeds, Chicago, IL, Oct 2002. http://www.net.cs.umass.edu/testbed_workshop/
- [5] S. Ganu, H. Kremono, R. Howard and I. Seskar, "Addressing Repeatability in Wireless Experiments using ORBIT Testbed", Proceedings of IEEE Tridentcom 2005, Trento, Italy, Feb 2005.
- [6] R. Draves, J. Padhye and B. Zill, "Comparison of Routing Metrics for Static Multi-Hop Wireless Networks", in proc. SIGCOMM, August 2004
- [7] P. Jacquet, P. Muhlethaler, T. Clausen, A. Laouiti, A. Qayyum and L. Viennot, "Optimized link state routing protocol for ad hoc networks," in proc. IEEE INMIC - Technology for the 21st Century, pp. 62–68, Dec 2001
- [8] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A High-Throughput Path Metric for Multi-Hop Wireless Routing", in proc. 9th ACM International Conference on Mobile Computing and Networking (MobiCom '03), September 2003
- [9] C. E. Perkins and E. M. Royer, "Ad-hoc on demand distance vector routing," in IEEE Workshop on Mobile Computing Systems and Applications (WMCSA), February 1999.
- [10] I. Chakeres, E. Belding-Royer and C. Perkins, "Dynamic source routing in ad hoc wireless networks," IETF MANET Working Group Internet-Draft, 2005, Work in Progress.
- [11] A. Tønnesen, "Implementing and extending the Optimized Link State Routing protocol", Master's thesis. University of Oslo, Norway, 2004
- [12] AODV-UU webpage, "<http://core.it.uu.se/core/index.php/AODV-UU>"
- [13] DYMOUM webpage, "<http://masimum.dif.um.es/?Software:DYMOUM>"
- [14] B. Kolundzija and A. Djordjevic, "Electromagnetic Modeling of Composite Metallic and Dielectric Structures", Artech House, 2002
- [15] H. Lundgren, E. Nordstrom, and C. Tschudin, "Coping with communication gray zones in IEEE 802.11b based ad hoc networks", WOWMOM '02: Proceedings of the 5th ACM international workshop on Wireless mobile multimedia, pages 49–55, New York, NY, USA, 2002. ACM Press.
- [16] A. Tirumala and J. Ferguson, "Iperf 1.2 - The TCP/UDP Bandwidth Measurement Tool", <http://dast.nlanr.net/Projects/Iperf>. 2001.

AUTHORS



David Johnson was born in South Africa in 1972. He received his B.Sc. (Electronic Engineering) at Cape Town University in South Africa. He is currently completing his M.Eng in Computer Engineering at the University of Pretoria in South Africa on scaled wireless grids for benchmarking ad hoc routing protocols. David Johnson has been working in the field of wireless mesh networks for the past 4 years and in telecommunications engineering for the past 6 years. He currently leads a research team which is addressing the fundamental research questions around wireless mesh networks for rural connectivity at the Meraka Institute, CSIR, South Africa. His current personal interests are computational intelligence, benchmarking ad hoc networking protocols using outdoor and indoor scaled testbeds, graph theory, and gateway location mechanisms in mesh networks. He is an active

participant in fora such as the “Wireless World Research Forum” and the “Association for Progressive Communications” which looks at mechanisms to build sustainable telecommunications infrastructure in Africa.



Albert Lysko (M’01) was born in Leningrad, USSR in 1974. He received his B.Sc. in Technical Physics and M.Sc. in Radiophysics in 1996 and 1998, respectively, both from Saint-Petersburg State Technical University in Russia. During his studies for B.Sc. and M.Sc. degrees, he was with the Ground Penetrating Radar Laboratory, Saint-Petersburg, Russia, dealing with non-destructive testing methods for characterization of materials at microwave frequencies. Presently, Albert Lysko is completing his dr.ing. (Ph.D). at the Norwegian University of Science and Technology (NTNU), Norway, where he also co-realized a system for 3D antenna and radar measurements. Currently, he is working at the Meraka Institute, CSIR, South Africa. His research interests are in smart antennas for wireless networking, as well as in numerical methods for electromagnetic design and analysis, in ultra-wide bandwidth antennas, and antenna and radar measurements. A. Lysko is a member of the Institute of Electrical and Electronics Engineers (IEEE), and also a former member of the Antenna Measurement Techniques Association (AMTA).




Fig. 25. OLSR topology using the ETX metric showing good multi-hop characteristics. The wireless NICS were configured to 802.11b mode, 11 Mbps data rate and a 0dBm power level. The values in the graph show the ETX values for each node pair in the network