

A Geography-Aware Scalable Community Wireless Network Test Bed

Bow-Nan Cheng, Shivkumar Kalyanaraman
Rensselaer Polytechnic Institute
chengb@rpi.edu, shivkuma@ecse.rpi.edu

Max Klein
Stanford University
maxklein@stanford.edu

Abstract

Wireless mesh networks have increasingly become an object of interest in recent years as a strong alternative to purely wired infrastructure networks and purely mobile wireless networks. Given the challenges that have arisen in construction, deployment, and maintenance of wireless mesh networks, we outline a broad experimental research program in the area of medium-to-large scale community wireless networks. Our research is conducted in the context of an operational community network built in our test bed laboratory with continual plans to expand to the town of Troy, NY (up to hundreds of nodes in a 1-2 mile radius around RPI campus). Leveraging Global Positioning System (GPS) receivers and Geographic Distributed Addressing (GDA), a novel and intuitive addressing assignment, geographic-based forwarding algorithms such as GPSR and TBF can be easily tested and traffic engineering theories implemented in a real-world environment. Our paper documents several design considerations and contributions in implementing community wireless networks including auto-configuration, addressing structure, and antenna characteristics among other items, in addition to describing our novel test bed lab where RF effects of distances of thousands of meters can be simulated with server, antenna, and variable attenuator clusters.¹

1. Introduction

The “last-mile” broadband infrastructure problem is perhaps the most important long-standing techno-economic challenge faced by the telecommunications industry [1]. While DSL and Cable Modem technologies are the dominant contenders, they have ultimate speed limitations which can be alleviated only through increased penetration of optical fibers[1]. Moreover, these wired alternatives tend to admit a duopoly (a cable company and a local bell company)

in terms of market structure. A variety of wireless technologies have been proposed as facilities-based competitors to wired infrastructure. These include point-to-multipoint WMAN technologies like 802.16 and millimeter wave technologies, fiberless or free-space-optical networks [2, 3], 2.5G/3G single-hop wireless access [4], 802.11 hot-spot operators and commercial multi-hop wireless techniques (eg: Nokia Rooftop [5], Mesh Networks [6]).

In addition to these commercial broadband initiatives, “grassroots” community wireless networks (CWNs) based upon multi-hop IEEE 802.11b [9, 10] technology are emerging in various parts of the world [7, 8] as an interesting paradigm to connect a community of users to the nearest broadband wire, and provide broadband wireless connectivity within the community. These ongoing grassroots efforts have contributed innovative designs (eg: pringles-can directional antenna [10]) and have had initial small-scale success stories. Current CWN efforts have been highly operations-oriented. A number of long-term challenges such as metro-scale deployment and coordination issues, address space management, auto-configuration and the seamless merging of community networks remain open.

From a research standpoint, we believe that community wireless networks (CWNs) have important characteristics different from both mobile ad-hoc wireless networks (MANETs) and traditional IP-based internetworks (eg: OSPF/IS-IS/BGP routed networks). MANETs are mobile and unstructured; while traditional IP-based networks are administration-heavy and do not leverage any geographical awareness. In contrast to MANETs, CWNs can leverage their structured, static nature, nearby infrastructure (eg: power, telephone lines, ISM wireless bands) and focus on scalability and configurability issues. Compared to traditional IP-based internetworks, CWNs must be easier-to-configure, seamless (to grow and merge organically) and can be geographically-aware.

Our project has thus far investigated the following areas:

- Antenna characterization, optimization, and usability
- Innovative hardware/software systems design and integration at the physical/link layer focusing on ease-of-

¹ This work was supported by NSF Information Technology Resesarch Program (NSF-ITR 0313095), and a grant from Intel Corporation

configuration and low-cost.

- A novel geographically based address scheme which allows for interchangeability between physical location and IP addressing and an autoconfiguration framework.
- A unique test bed laboratory utilizing variable attenuators to provide scales of anywhere from 2 meters to 2 kilometers of distance.

The our research concerns the nascent community wireless network (CWN) paradigm. We propose a unique mix of cross-cutting activities including hardware design, software design, real-world experimental test bed deployment, measurement and modeling. Success of this project could have a fundamental impact on the economics of the critical last-mile problem.

The rest of the paper is organized as follows: In Section 2 we survey related work in this area. Section 3 presents the details of our physical layer hardware, Section 4 presents the network layer while Section 5 describes our laboratory test bed in detail. Finally, Section 6 describes our future work in this direction while Section 7 presents the concluding remarks.

2. Related Work

There has been a large body of research in several aspects of ad-hoc networks, especially routing scalability [11, 12], geographic and trajectory routing [13, 14], and fundamental capacity limits [15, 16, 17]. A large fraction of these proposals have been evaluated through simulation. In contrast, there has been lesser attention devoted to large-scale experimental activities in ad-hoc networks. However, a few recent experimental efforts deserve prominent mention. The Monarch Project in CMU/Rice [18, 19] (based upon an earlier 2 Mbps version of 802.11) investigated mobile ad-hoc routing, primarily from the perspective of mobility (eg: moving cars etc). Recently MIT's Grid Project [20] has built a small-scale rooftop ad-hoc network primarily for research [21]. Their experience (which we have observed independently) suggests that *real-world deployment and performance differs substantially* from assumptions made in simple analysis and simulation of MANET protocols [21, 22].

GPS or other location information can be incorporated into routing tables to facilitate efficient inter-cluster routing (thus providing scalability with network size) or to approximate shortest path routing (e.g., DREAM [23], LANMAR [24], GLS [25]). While DREAM requires proactive flooding of location information, LAR requires reactive flooding. LANMAR employs a hierarchy to avoid global flooding, but is susceptible to nodes at the top of the hierarchy being mobile. GLS gets rid of global flooding by intro-

ducing a subset of nodes designated as location servers. L+ extends GLS to be more auto-configurable.

In contrast to these approaches, our project focuses on real-world deployment activities and testing through a street-level (as opposed to rooftop) RF-Aware network, a unique hashing algorithm that eliminates global and even location-oriented routing information flooding, and a novel laboratory test bed environment that scales from 2m to 2km through variable attenuators. We propose Geographic Distributed Addressing (GDA) which embeds location information into IP address structure providing the ability to route geographically without maintenance of a complete routing table. Our proposed hardware/software design/integration/real-world deployment activities are also a big differentiator from prior work that have largely been limited to simulation-based evaluation.

3. Physical Layer: A "Street-level" Network

One major drawback of wireless radios is the need for line-of-sight between the transmitter and receiver to effectively communicate information between two points. Obstructions such as trees, buildings, signs, among other things, significantly degrade signal strength and quality, forcing increased delays in retransmission or even a break in network connectivity. CWN test beds have traditionally tackled this problem by placing antennas on Rooftops ([7, 20]) utilizing predominantly omnidirectional wireless antennas to avoid as much obstructions as possible. While this approach has its strengths, several issues can be readily seen. First, the need for roof-top (or near roof-top) mounting complicates installation because roof-tops are often inaccessible to users and connections between antenna and router hardware require a significant amount of RF cabling resulting in heightened signal degradation. Second, in buildings where height is significantly lower than the surrounding buildings, installation of an antenna even on the top of the roof forces individuals to face recurring obstruction problems. Third, while omnidirectional antennas provide a uniform coverage radius, they are often times weaker in signal strength than comparably priced directional models resulting in the need for more antennas to complete coverage. While it is true that several wireless networks operate without line-of-sight (LOS), LOS and near-line-of-sight (NLOS) enhances capacity dramatically and we will explain in the next few subsections how we plan to enhance LOS in our network.

Our project focuses on an "under-roof/street level" approach to wireless antenna placement leveraging the directionality of streets and directional yagi antennas to mitigate the problems provided by obstructions. Mounting our yagi antennas on telephone poles, street signs, street lamps,



Figure 1. A 12dBi directional yagi antenna looking down a clear street. Directional antennas, bolstering higher gains, increases capacity even in NLOS situations.

trees, and other “above ground” objects near streets allows us to utilize existing, relatively obstruction-free, under-roof area to transmit information. Because roads are usually relatively straight, we’ve chosen to use directional yagi antennas, which boast high signal to noise ratio in the forward direction, to pierce through any minor obstructions on the street level and ensure strong connectivity between nodes.

While it is usually desirable to make sure antenna coverage blankets entire areas (which would be rather difficult with directional antennas), the chief purpose of our node and antenna placements are to provide an infrastructure overlay rather than ubiquitous network access. Each node runs a DHCP server that allows users to share network access in whatever way they deem necessary. With respect to the link layer, we have chosen to utilize 802.11’s MAC protocol in our test bed though any MAC protocol is fine as long as it is kept consistent across the board.

3.1. Antenna Data

One of the initial designs of directional yagi antennas we explored was the Pringles can antennas suggested by [10]. By utilizing off the shelf items and purchasing some relatively inexpensive parts, we were able to mass produce dozens of these directional yagi antennas for a fraction of the cost it would take to purchase equivalent directional antennas. After several trials and studies, we standardized our design to the one mentioned in [10] with the N-type connector replaced with a SubMiniature version A (SMA) connector to lessen insertion loss and the traditional aluminum tube replaced with an-oxide coated aluminum tube to lessen corrosion. Despite the modifications, our total achieved gain was approximately 6.0143dBi, which, although rather nice for such a low cost, was surprisingly below our expectations.

Taking into considering factors that come with real-world deployment such as rain and wind, it was decided that while the Pringles can antenna design was suitable for indoor usage, outdoor applications would require something a little more resilient. Hyperlinktech’s 12dBi Radome Enclosed Yagi antenna met all the requirements we were looking for, featuring a weatherproof chassis and a rear mount that could withstand winds of upwards of 150mph. The beam width for both vertical and horizontal polarization is approximately 45° which gives us a good coverage for streets that curve or are at higher elevation than others.

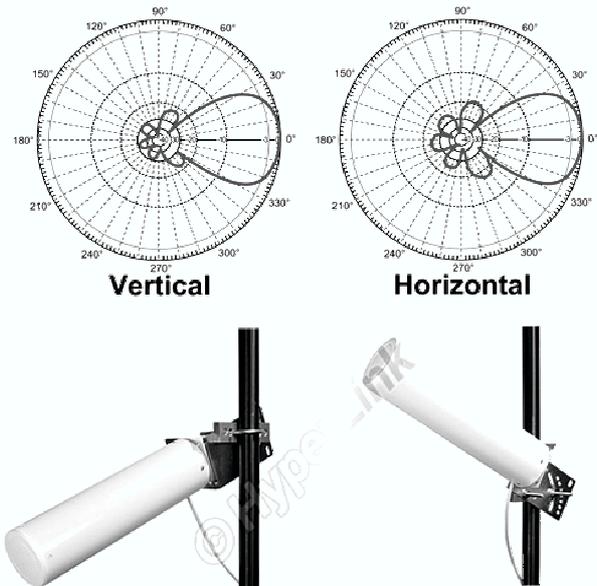


Figure 2. Hyperlinktech’s 12dBi Radome Enclosed Yagi antenna featuring a weatherproof chassis and a rear mount that can withstand winds of upwards of 150mph.

4. Network Layer: Addressing Framework and Autoconfiguration

Network administrators have traditionally solved the problem of addressing and autoconfiguration through services such as DHCP that temporarily lease IP addresses to hosts that connect to a network. But centralized approaches such as DHCP work most effectively with wired networks where packet losses are rare and range is effectively unlimited. We propose an addressing framework (explained later) that is both autonomous and decentralized. In addition, with the notable exception of the GRID project [20], there has not been significant experimen-

tal studies of geographic and location-based routing. We conjecture that there are several synergies between geographic routing and traditional topologically aware routing, and an experimental deployment in a Community Wireless Network (CWN) is an attractive way to study these synergies.

In particular, topology-aware methods can capture rich topology characteristics (eg: map/graph of a local cluster of nodes, nodes' GPS-location, link's antenna characteristics, modes (802.11a etc), power and performance statistics). These characteristics offer tremendous scope for local traffic engineering optimization within the cluster. But the topology-aware traffic engineering approach becomes complex across clusters, because of problems in aggregating and announcing such information in addition to announcing basic reachability information.

In contrast, *pure* geographically-routed algorithms have no sense of network topology, and hence encounter problems when the geographical chosen path encounters obstacles. Though innovative solutions like GPSR [13] and location proxies [26] have been proposed, these are suboptimal compared to autonomous *local* traffic engineering optimization that can be coupled with *global* geographic routing. An attractive hybrid method would propagate topological (link-state) information only within a cluster of nodes, and use geographical routing between clusters. Geographic-style routing, in general, forms a good basis for seamlessly merging CWNs that organically develop, because of the common GPS reference.

4.1. Geographically Distributed Addressing

A key step in integrating geographic and topologically aware routing is to use a *common addressing scheme*. We've developed Geographically Distributed Addressing (GDA), a common addressing scheme that implements a low-overhead hash function based on each individual node's GPS latitude/longitude information along with a default IPv4 or IPv6 prefix assignment (eg: 10.*.* IPv4 private space or a geographically-assigned IPv6 public address space [27]). Leveraging the inherent nature of latitude lines which run horizontal and change in value from north to south much like the Y-axis in a coordinate plane and longitude lines which run vertical and change in value from east to west much like the X-axis in a coordinate plane (see Figure 3), the hash function for GDA produces a geographically aware IP mapping based on the physical location of the node. Though the surface of the earth is curved, for all practical purposes, a flat plane is assumed for areas of deployment and latitude and longitude is measured in degrees. In addition, we assume that all longitude lines *west* of the prime meridian and latitude lines *south* of the equator are negative

allowing for longitude values from -165° to $+165^\circ$ and latitude values from -90° to 90° .

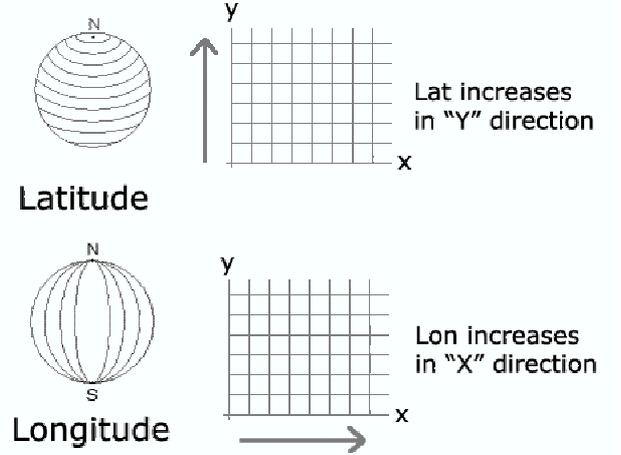


Figure 3. Longitude lines correspond nicely with the X-axis in a coordinate plane while Latitude lines correspond with the Y-axis

In GDA, a geographic coordinate plane grid is formed as a geographic "origin" (*degrees lon, degrees lat*) is selected (σ_x, σ_y) and the axes are spread in the positive X and positive Y directions. GPS degrees for each node's longitude and latitude coordinates (λ_x, λ_y) are hashed to an IP address 10.X.Y.Z using the following scheme:

$$a = (\lambda_x - \sigma_x) \times S$$

$$b = (\lambda_y - \sigma_y) \times S$$

$$X = \lfloor a/16 \rfloor \text{ (} x \text{ - coordinate)} \quad (1)$$

$$Y = \lfloor b/16 \rfloor \text{ (} y \text{ - coordinate)} \quad (2)$$

$$Z = \left((\lfloor a \% 16 \rfloor) \times 16 \right) + \left((\lfloor b \% 16 \rfloor) \right) \quad (3)$$

Reverse Hash:

$$\lambda_x = \left(\frac{(X \times 16)}{S} + \frac{(Z/16)}{S} \right) + \sigma_x \quad (4)$$

$$\lambda_y = \left(\frac{(Y \times 16)}{S} + \frac{(Z \% 16)}{S} \right) + \sigma_y \quad (5)$$

S is the scaling factor in powers of 10 which gives us flexibility in assigning addresses over large or short distances. We define precision as the inverse of the distance required to hash unique address assignments. As precision is increased, the distance required to hash unique addresses

is decreased. The expected area of coverage is $C_x \times C_y$ where $C = (256 \times 16)/S$ in degrees and C_x and C_y represent C converted from degrees into a unit of length of longitude and latitude respectively. The minimum distance required to generate unique “X” and “Y” octets is approximately $16/S$ in degrees and must be converted to a unit of length longitude and latitude-wise. The minimum distance for unique addresses as the “Z” octet is added is approximately $1/S$ degrees converted to a unit of length longitude and latitude-wise. Our intention upon deployment is to utilize $S = 10^4$ which provides a rectangular coverage of approximately $33.469km \times 45.56km$ with a minimum node separation of roughly 11m going north and south and 8m going east to west. Forwarding by comparing only the “X” and “Y” octets alone would require a separation distance of approximately 177m going north and south and 128m going east to west. Modifying the scaling factor S results in one important tradeoff: as S increases, the area of coverage decreases but the precision increases.

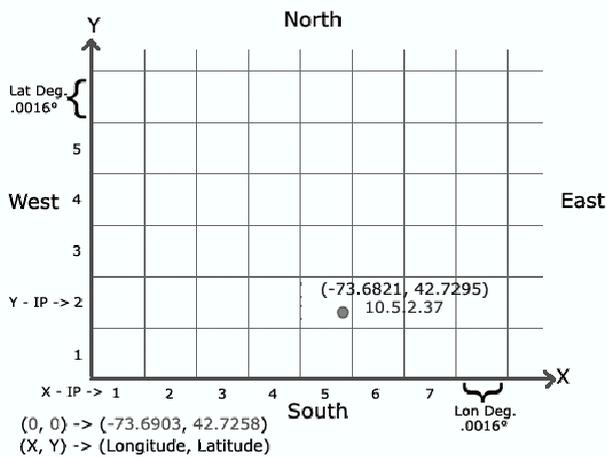


Figure 4. GPS Coordinates in degrees around the RPI campus mapped to a coordinate X-Y plane

With GDA, all nodes on the network agree to the hashing function which translates IP information to physical geographic location (given via GPS receivers). Each node simply maintains a table of its immediate neighbors and when a packet is forwarded, each router in the path examines the packet, unpacks the destination address, hashes it into a geographic location, compares the location with the nodes in its neighbor table, and forwards accordingly. The overhead associated with calculations at each node is negligible since most of the time, the hashing from IP to geographic location

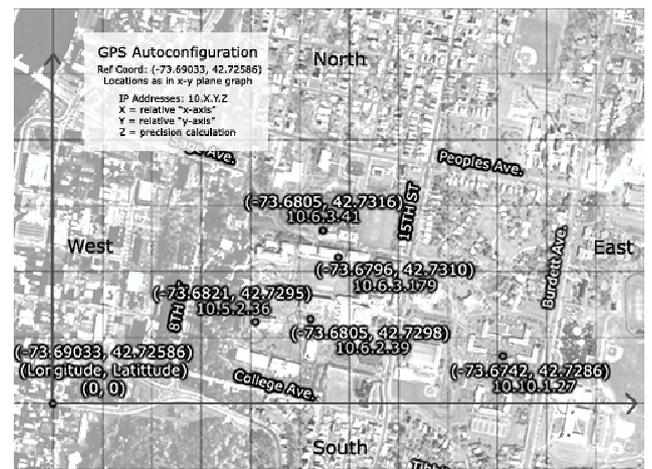


Figure 5. An Aerial view of the area around the RPI campus overlaid with GPS to IP mappings in the X-Y plane

is not required as IP addresses are geographically aware. By simply comparing the “X” and “Y” octets, a node can easily determine whether the destination is to the North, East, South, or West of itself and forward to the neighbors in the general direction accordingly. If the “X” and “Y” octets are equal for both source and destination, then the precision octet “Z” is hashed back into the GPS coordinates for closer approximation.

Our actual implementation of GDA utilizes MIT’s Click Modular Router [29] to transform simple Linux boxes into highspeed routers. Click works on configuration files that dictate what to do with incoming and outgoing packets. By simply generating a configuration file with the appropriate IP and MAC addresses and loading that configuration file into Click, one can easily and quickly setup a router to fit their needs. When the machine is powered up, a GPS daemon service is launched that communicates with the GPS receiver, polling it for the machine’s physical geographic location using NMEA 0183 standards until a fix is established. When a location fix is established, a hashing service is initiated that grabs the latitude and longitude coordinates and hashes it into a unique IP address which it uses to generate the Click Modular Router configuration file. Once the configuration file is generated with the hashed IP address and auto-sensed MAC address as well as elements for routing, the router is started and neighbor discovery algorithms available in Click are initiated.

Given such dual-sense addressing, a number of hybrid options become possible. For example, the source of a flow could encode a loose trajectory as a fixed-length *global geographic hash* value to be used for *inter-cluster routing*.

This geographic hash would complement an *intra*-cluster explicit routing strategy that would use a separate fixed-length *local topological hash* value valid only within that cluster. Note that this approach *does not* assume global visibility of topology information, *unlike* protocols like PNNI, NIMROD, Landmark [24]. Alternately, we could aggregate and distribute topological information, but *in a location-aware* manner, unlike the location-agnostic approaches in the aforementioned protocols. Hence our approach should be less complex, more scalable, auto-configurable and flexible compared to PNNI, NIMROD or Landmark routing. In summary, our abstract fixed-length path encoding strategy admits a wide range of explicit routing strategies.

5. Microcosm Testbed Lab

While one of the eventual project goals is to produce an auto-configuring “standard box” that provides end users with transparent and relatively painless network access, one of the major setbacks over the past year of work has been the physical instability of any test network we have deployed. Coupling the knowledge that people move in and out of apartments and dorms quite frequently and the fact that landlords generally have little interest in a purely developmental, as opposed to permanent network solution, we have revisited and extended a solution traditionally used in wired networks: utilizing multi-vlan switches connected to clusters of multi-homed Linux systems to simulate, in close quarters, entire network topologies. Our wireless network prototype works in much the same way, occupying about the same amount of space as the wired Linux simulator equivalent, plus normally unused ceiling space for an antenna grid.

5.1. Test Bed Theory

The wireless test bed must be capable of simulating large (10’s-100’s of meters between nodes) outdoor wireless networks within the confines of our relatively small lab (about 10 meters total). To figure out how to accomplish such a feat, it is necessary to understand RF path loss and RF budgeting. The path loss is a measure (usually in dB) of the amount of RF signal that is lost, faded, or attenuated between two points in space. For our initial system, we have assumed that all paths are line-of-sight and obstruction free (see the “Future Work” section for non-free space / non-line-of-sight propagation information). In [10] it is that free space RF propagation loss (in dB) assuming operating frequencies in the 2.45GHz range is:

$$40 + 20 \times \log_{10}(\text{distance in meters}) \quad (6)$$

Our antenna grid is approximately 2 meters per space, which would equate to $40 + 20 \times \log_{10}(2m)$ or about 46dB of signal loss between adjacent antennas. However,

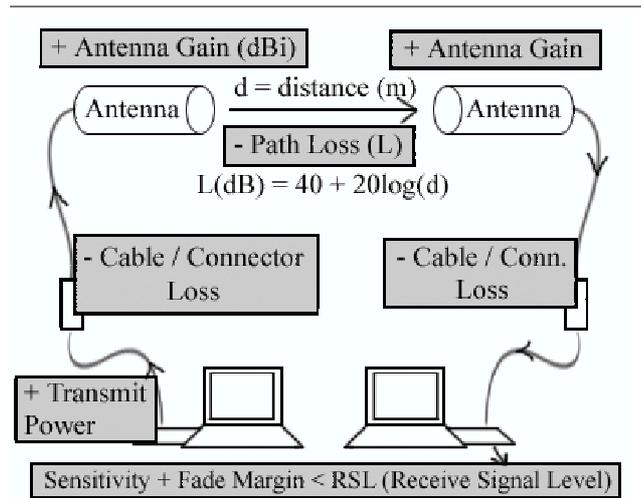


Figure 6. RF Budget Considerations

by adding a variable attenuator with a maximum attenuation of $60dB$ on just one side of the hop, the free space plus attenuator loss would be anywhere between $46dB$ and $106dB$ ($46dB$ path loss + $60dB$ attenuator loss). Again, using Equation 6, $106dB$ would equate to a 1995-meter effective node separation. By using a $30dB$ attenuator on each antenna in our grid, allowing an extra $60dB$ of path loss per hop, our test bed could scale from 2m to nearly 2km of distance between adjacent grid spacings. Increasing the range of our variable attenuators would continue to exponentially increase the maximum node separation of our test bed. Additionally, since each variable attenuator is independently controlled, it is possible to simulate a pocket of active nodes, with additional nodes separated by a much larger distance. In short, each node can be scaled by a different factor.

RF budget, however, includes not only the free space path loss, but also the antenna gains, cabling and “insertion loss” (this is the loss due to connectors between cables, cards, antennas, etc), transmitter power, and receiver sensitivity. To ensure that our test bed models the path’s impact on the system performance, we taken the following steps to mitigate these other RF factors:

- Antenna Gains - The antennas are part of the system we are testing, and therefore are not to be considered an extraneous variable. To start, we are using only Pringles can antennas, however there is no restriction on using different antennas or combining antenna types. When a mix of different antennas is used, however, it will be necessary to carefully consider how to weigh attenuation between each side of a hop if non-uniform attenuations are used. Using uniform attenuations will work fine without any extra consideration, and non-uniform attenuations are inherently more

complex to begin with, so the added complexity due to varying antennas should not be overwhelming, but must be considered during test planning.

- **Insertion and Cabling Loss** - Each antenna in our grid is wired into the interconnect panel with the exact same length and type of cable, and the same type, model, and number of connectors is used. By having these items remain constant, the associated losses can be more easily calculated across the entire system.
- **Transmitter Power and Receiver Sensitivity** - To ensure that no unintended variations in transmitter power or receiver sensitivity are introduced, we have standardized our radio platform within the test bed to the Orinoco Gold PCMCIA 802.11b cards. Although these cards cost slightly more than the silver cards (about 5–10), they support more advanced features such as variable transmitter power and better interoperability with the Lucent access points that we use. Once again, by utilizing identical hardware on each node, these factors can be kept constant across the system. Intentional variations in transmitter power, for example in future work on power-aware routing, will still be testable in the lab test bed without any modifications.

5.2. Hardware and Software Setup

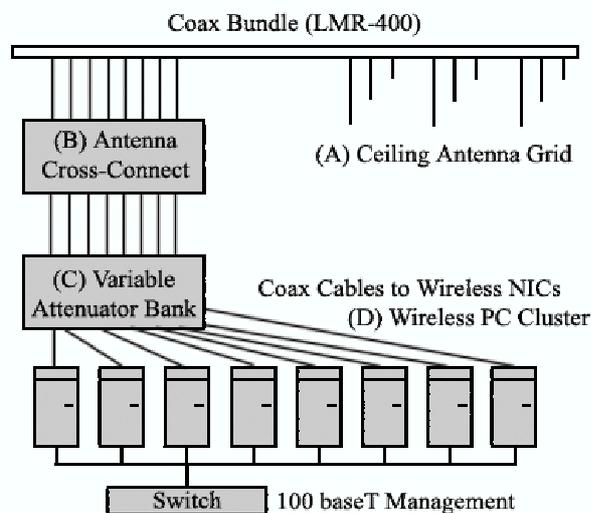


Figure 7. Test Bed intended setup

Figure 7 depicts our intended test bed setup (our current setup varies only in node count and “neatness” and ease-of-use items, such as a rack mounted cross-connect panel). The system consists of a ceiling mounted antenna grid (A) fully populated with antennas. This grid is currently formed

by attaching lab “ring stands”, or rods of aluminum, to the electrical conduit for the overhead lighting with standard chemistry lab components. The antennas are then attached to these rods with a double spring made from a coat hanger. One spring wraps around the rod, while the other, offset by 90°, goes around the antenna. This dual spring design allows the antennas to be rotated in any direction, with the ability to hold its new position without further user intervention. Each antenna is then wired (with identical wire, wire lengths, connectors, etc) to a cross-connect panel (B). For our initial system, we were able to use four identical pieces of cable (roughly LMR-400 sized, but heliax) that we had on hand from donated yagi antenna install kits. Identical N to SMA adapters were used at each antenna.

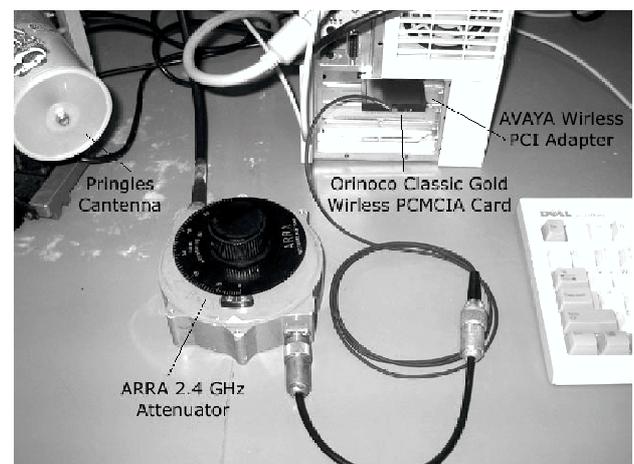


Figure 8. Test Bed hardware equipment used: a Dell system running Redhat 9 Linux; ARRA 2.4 GHz range variable attenuator; Pringles Can antenna, Orinoco Classic Gold wireless PCMCIA card

Currently, we do not have a cross-connect panel, and are just directly connecting the antenna cabling to the appropriate attenuator. Ideally, we should have a two-row N-type cross-connect panel, with the top row containing one connector per antenna, and the bottom comprised of one connector per test bed node. This panel will allow the user to select which node is connected to which antenna by simply placing a jumper between the two corresponding connectors on the cross-connect. This cross-connect can also be replaced by an electrically controlled RF switch, allowing a computer to configure the node-antenna pairings as appropriate for the experiment. This upgrade is mentioned in the “Future Work” section. The other side of this panel (or switch) is connected to each node’s wireless card (D) in the test cluster through a unique attenuator (C). We cur-

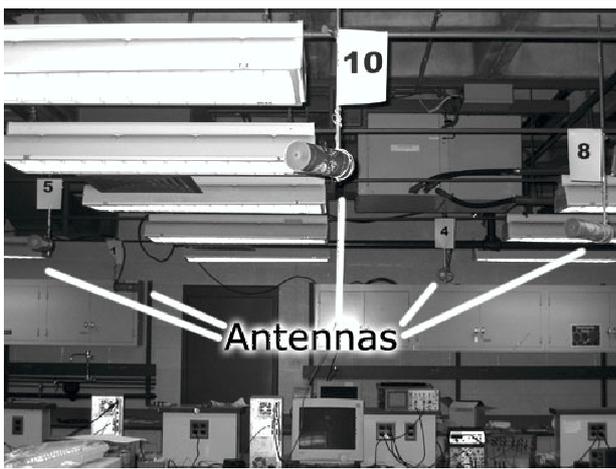


Figure 9. Test Bed Pringles antenna grid

rently have a limited supply of analog variable attenuators, and require either more analog attenuators to complete testing. These attenuators allow the lab test bed to scale the distances between each node as described in the previous subsection. The cluster of experiment nodes can be easily managed through a back-end, management only, wired Ethernet network. All of the above hardware, excluding the ceiling antenna farm, is rack mountable for easy management and containment. Additionally, all of the proposed additions under future work, excluding devices for non-line-of-sight testing, are also rack mountable.

In terms of software platform, we decided to leverage existing open source options such as Linux and GPSD, a service that communicate with GPS devices using NMEA standards, to build our router platform. Instead of manually modifying the Linux kernel code to support new routing protocols and implementations, MIT's Click Modular Router software [29] was used to turn a simple Dell machine (with Intel chips) into a flexible wireless router. Click is a software architecture for building flexible and configurable routers. A Click router is assembled from packet processing modules called elements which implement simple router functions like packet classification, queuing, scheduling, and interfacing with network devices. Writing elements that deal individually and independently with packet processing, routing, etc. allows for quick implementation of test protocols and rapid extension of the entire network system. As described in the "Auto-configuration" section of the paper, several shell and Perl scripts were also written to automatically generate hashed IPs as well as Click configuration files for usage.

Rapid deployment, upgrading, and maintenance to several nodes at once requires a vast amount of time and effort. To help mitigate some of the issues involved including software version consistency, scalability, among others, we

have been working on setting up a "diskless" Linux workstation which mounts its root filesystems via NFS [30] to a server. Though currently we are utilizing full-blown systems, NFS-root will help us transition into smaller, embedded Linux systems and help in rapid turnaround during testing and development of new protocols and setups with Click.

6. Future Work

While much of the addressing and autoconfiguration framework and preliminary software-hardware integration tasks have been completed, there are several areas of research and implementation left to be explored. The following subsections give a thumbnail sketch of work to be done.

6.1. A RF-Aware Network

Traditional wireless routing models rely on nodes exhibiting omnidirectional regions of radio coverage. While on paper this provides a nice holistic approach to "best route" analysis, real life is not always so simple. Because of the inherent nature of our test bed design which relies on directional yagi antennas, we are able to better study RF effects on routing decision in hopes of developing an "RF Aware Network". Take for example a directional antenna trying to send data to a node that is in the direction of its weakest gain. Because yagi antennas are directional, gains are significantly higher in one direction than in other directions. Sometimes, it would be better to forward packets in the direction of higher gains to ensure packet delivery (and subsequent forwarding) rather than forwarding in the direction of lower gains and wasting time with timeouts and retransmissions even though on paper, routes look shorter when forwarded in the direction of lower gains.

Antenna directionality goes hand in hand with determining a node's relationship to its neighbors. Omnidirectional antennas, while boasting a greater connectivity radius, require more nodes due to weaker signal strength while directional antennas provide stronger connectivity to fewer nodes. Current wireless routing models do not take into account the directionality of certain antenna gains and we hope to be able to study and analyze effects of antenna characteristics in routing and other applications. We intend to study several metrics to characterize RF effects on the network including antenna directionality, transmit power and receiver sensitivity, and path obstacles.

6.2. GDA Extensions

- GDA hashing modifications to include new added nodes in negative directions from the assumed origin. Currently, GDA requires a selected origin with

nodes being in the positive x and y coordinates relative to this origin. If a new node joins that happens to be in the negative x or y directions with respect to the origin, readdressing will be required. We plan to extend GDA to include all directions and perhaps clustering effects (more nodes in a certain vicinity requiring more precision) through possible uses of a smith chart-like coordinate plane.

- GDA currently hashes to IPv4 IP addresses. We plan to take the concept behind GDA and explore ways to seamlessly integrate it into IPv6 as well.

6.3. Lab Test Bed Improvements and Deployment

- Automatic Test Bed Configuration and Management - It would be desirable to have one extra computer in the test bed rack to run a GPS simulator which sends out proper NMEA encoded GPS coordinate information for each node as well as control the variable attenuators and possibly even the cross-connect. Controlling the cross-connect would require RF antenna switches (which we could possibly make on our own), but would allow for users of our lab test bed to simply enter the experimental parameters into the management machine and upload the routing algorithm to the cluster; all other setup and configuration would be done automatically. This addition would allow for experiments to be setup and run remotely by anyone, similar to existing wired simulation clusters.
- Reflectivity of the lab - Measuring the impact that the RF reflectivity of the lab walls, ceiling, hanging lights, and other artifacts in the lab has on our experimental results can be difficult. We have recently acquired several RF absorbent pads to minimize RF reflectivity and hope to devise test methods categorize effects of reflectivity on our experiments in the lab.
- Once protocols and tools have been written and hardware components standardized, it is our hope to deploy these nodes around the University. Current work is being developed in terms of locating places to setup the antennas, soliciting volunteers to house the nodes, and managing the infrastructure remotely.

6.4. Traffic Engineering and Multipath Routing Studies

- We plan on developing a broad framework and mechanisms for *connectionless* traffic engineering in Community Wireless Networks (CWNs) deployable on a wide variety of legacy as well as current protocols and platforms.

- Within the above framework, we propose to leverage the integration of GPS receivers to experimentally explore the *intersection* between the areas of multipath routing, scalable multi-cluster geographic routing and topologically-informed routing, i.e. hybrid routing schemes.

7. Conclusion

In this paper, we've explored several basic design issues associated with setting up community wireless networks including address structure, antenna characteristics, auto-configuration, and scalability. We have also outlined several key steps in establishing a full-blown community wireless network around the RPI campus through studying RF characteristics and erecting a fully scalable test bed lab to test theories before "going live". Because of the strong interest of harnessing the strengths of both topological and geographic routing, GDA was established to give IP addresses geographic significance. Simply by eye-balling an IP address, individuals can easily determine whether the destination is north, south, east, or west of itself. Similarly, geographic positions are now given network topology significance as well. GDA mitigates several issues involved in proactive and reactive protocols including need to maintain advanced routing tables and states, constant network flooding of maintenance packets, complex location services, and fear of centralized point of failure. Furthermore, by implementing code to automatically poll latitude and longitude coordinates from a GPS receiver, hash it to an IP address, and generate a Click configuration file used in the launch of the router, an auto-configurable and nearly maintenance-free node is generated. Such nodes lay the framework for rapid deployment and seamless connectivity.

As a major landmark to full deployment in and around the RPI campus, we are continuing to develop a fully scalable and testable test bed lab. Utilizing variable attenuators that go up to 60dB, we are able to simulate anywhere from 2 meters of node separation to nearly 2km of node separation simply by playing with RF budgeting principles. The Click Modular Router allows us not only to quickly transform simple Linux boxes into advanced routers, but also easily write and test various routing algorithms and network technology. Delving into technologies of NFS-root among others, we are able to quickly distribute and test several wireless frameworks as well as experiment on new protocols and algorithms. While still relatively in the early stages of development, our test bed lab represents a novel way to experiment and test various ideas and theories in various network settings and topologies in a real-life manner. It is our goal and hope that the research done here can someday benefit other grassroots CWN and industrial wireless mesh network efforts around the world.

8. Acknowledgements

- Prof. Biplab Sikdar (sikdab@rpi.edu) for his help and advice throughout the research.
- Prakash Iyer from Intel for funding and evolving the CWN to an infrastructure mesh (IMWN) framework.
- Vijay Subramanian (subrav@rpi.edu) for his work on coding GDA
- Mitesh Parikh (parikm@rpi.edu) for his work on antenna characterization
- Alex Newman (dolemite@wuilmasters.net) for his work on NFS-Root lab setup
- Roni Mazumdar (mazumd@rpi.edu) for his work on deployment logistics

References

- [1] NRC Committee on Broadband Last Mile Technology, "Broadband: Bringing Home the Bits," em National Research Council, CSTB, 2002.
- [2] H. Willebrand, B.S.Ghuman, "Free Space Optics," *Sams Pubs, 1st edition*, December 15, 2001.
- [3] Acampora, A.S. et al., "A broadband wireless access network based on mesh-connected free-space optical links," *IEEE Personal Communications*, Vol. 6, pp. 62 -65, Oct. 1999.
- [4] H. Kaaranen, S. Naghian, L. Laitinen, A. Ahtiainen, V. Niemi, "UMTS Networks: Architecture, Mobility and Services," *John Wiley & Sons, 1st edition*, August 2001.
- [5] Nokia Rooftop Broadband Solutions, <http://www.wbs.nokia.com/solution/index.html>
- [6] Mesh Networks Inc., <http://www.meshnetworks.com/>
- [7] Seattle Wireless, <http://www.seattlewireless.net/>
- [8] FreeNetworks, <http://www.freenetworks.org/>
- [9] M.S. Gast, "802.11 Wireless Networks: The Definitive Guide," *O'Reilly & Associates, 1st edition*, April 2002.
- [10] Rob Flickenger, "Building Wireless Community Networks," *O'Reilly & Associates, 1st edition*, December 15, 2001.
- [11] E.M. Royer and C.-K. Toh, "A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks," *IEEE Personal Communications*, pp 46-55, Apr. 1999.
- [12] G. Finn, "Routing and addressing problems in large metropolitan-scale internetworks," *Technical Report ISI/RR-87-180*, ISI, March 1987.
- [13] B. Karp and H.T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, August 2000.
- [14] J. Li, J. Jannotti, D.S.J. De Couto, D. Karger, and R. Morris, "A Scalable Location Service for Geographic Ad-Hoc Routing," *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, August 2000.
- [15] P. Gupta, and P. Kumar, Internets in the Sky: The Capacity of Three Dimensional Wireless Networks, *Communications in Information and Systems*, 1(1):33-49, January 2001.
- [16] J. Li, C. Blake, D. De Couto, H. Lee, and R. Morris, Capacity of Ad Hoc Wireless Networks, *ACM MobiCom'01*, July 2001.
- [17] B. Liu, Z. Liu, and D. Towsley, On the Capacity of Hybrid Wireless Networks, *IEEE INFOCOM'03*.
- [18] D. A. Maltz, J. Broch, and D.B. Johnson, "Experiences designing and building a multi-hop wireless ad hoc network testbed," *Technical Report, CMU-CS-99-116, Carnegie Mellon University, School of Computer Science*, March 1999.
- [19] D.A. Maltz, J.Broch, and D.B. Johnson, "Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed," *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.
- [20] MIT Grid Project, <http://www.pdos.lcs.mit.edu/grid/>
- [21] B.A. Chambers, "The Grid Roofnet: a Rooftop Ad Hoc Wireless Network," Master's Thesis, June 2002; <http://www.pdos.lcs.mit.edu/grid/>
- [22] D.S.J. De Couto, D. Aguayo, B.A. Chambers, and R. Morris, "Performance of Multihop Wireless Networks: Shortest Path is Not Enough," *Proceedings of the First Workshop on Hot Topics in Networking (HotNets-I)*, Princeton, New Jersey, October 2002
- [23] S. Basagni, I. Chlamtac, V.R. Syrotiuk, and B.A. Woodward, "A Distance Routing Effect Algorithm for Mobility (DREAM)," *ACM/IEEE MobiCom 98*, pp. 76-84, Oct. 1998.
- [24] P.F. Tsuchiya, "The Landmark hierarchy: A new hierarchy for routing in very large networks," *ACM SIGCOMM 88*, pp. 35-42, Aug. 1988.
- [25] J. Li, J. Jannotti, D.S.J. De Couto, D.R. Karger, and R. Morris, "A Scalable Location Service for Geographic Ad Hoc Routing," *ACM/IEEE MobiCom 2000*, pp. 120-130, Aug. 2000.
- [26] D.S.J. De Couto and R. Morris, "Location Proxies and Intermediate Node Forwarding for Practical Geographic Forwarding," *MIT Laboratory for Computer Science technical report MIT-LCS-TR-824*, June 2001.
- [27] C. Huitema, "IPv6: The New Internet Protocol," *Prentice Hall*, Second Edition, 1998.
- [28] A. Elwalid, et al, "MATE: MPLS Adaptive Traffic Engineering," *INFOCOM'01*, April, 2001
- [29] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. *ACM Transactions on Computer Systems 18(3)*, August 2000, pages 263-297.
- [30] The Linux Document Project, <http://www.tldp.org/HOWTO/NFS-Root.html>