IEEE 802.11 Experiments in Virginia's Shenandoah Valley

The middle of the National Radio Quiet Zone is one of the last places you'd expect to find hams engaged in active wireless experimentation.

David R. Fordham, KD9LA

As you have probably noticed, commercial IEEE 802.11 (hereinafter referred to as "802.11")[1] wireless network equipment can now be purchased from local office supply houses for less than the cost of a 1200 baud VHF terminal node controller. Some adventurous Virginia amateurs were amazed at how little money and effort it took to establish a reliable high-speed wireless data network across long distances—distances traditionally associated with voice repeaters and 1200 baud packet radio. Using off-the-shelf components, these hams easily created a network operating at Internet-compatible speeds across a span of 34 miles in a single hop. Yes, you read correctly: that's 34 miles!

Let's take a closer look at what is going on in this rural farming region—just a few hours' drive from Washington, DC.

History and Geography

Virginia's beautiful Shenandoah Valley runs diagonally along the western edge of the state. To the east are the scenic Blue Ridge Mountains, capped by the famous Skyline Drive and the Blue Ridge Parkway. To the west, the Allegheny Mountains form the border with West Virginia. Nestled down in between these two mountain ranges lies a 50 mile stretch of rolling hills, farmland and orchards that provide picture-postcard scenery.

Figure 1—All of the Amateur Radio activities described in this article occurred in the National Radio Quiet Zone. Established in 1958 to minimize possible harmful interference to the National Radio Astronomy Observatory in Green Bank, West Virginia, the Quiet Zone encompasses nearly 13,000 square miles near the state border between Virginia and West Virginia.

[1]802.11 is a trademark of The Institute of Electrical and Electronics Engineers (IEEE)

IEEE 802.11 Jargon: "Routers, Gateways and Bridges...Oh My!"

Before shopping for 802.11 gear, it helps to know something about how commercial wireless networks differ from traditional ham packet networks.

A standard AX.25 packet network starts with a user computer connected via serial port to a Terminal Node Controller (TNC), which in turn is connected to a radio transceiver. The TNC takes a data stream coming out the computer's serial port, divides it into 128-character chunks called packets, and then modulates the radio carrier in a series of individually numbered transmissions corresponding to each packet.

A typical IEEE 802.11 "access point." This routing device combines a 4-port hub, a bridge and a gateway circuit for the Internet interface. It also includes an RF transceiver—a complete radio package to the Internet. All that's needed is an Internet connection via a network cable or modem (cable or DSL) and an RF network interface card (NIC) plugged into a desk computer or laptop. Presto—wireless access!

The other end of a traditional packet circuit usually consists of one of the following three types of stations: (a) another ham with an identical setup (computer, TNC and radio), or (b) an unmanned radio connected to a TNC which receives packets and then retransmits them on the same frequency (called a digipeater), or (c) a device called a "node."

The node consists of a radio and a specially programmed TNC. The node TNC receives packets, and repeats them on the same frequency like a digipeater. But the node also handles connections and performs other administrative overhead duties, resulting in more efficient use of the radio channel than a simple digipeater.

And then there are dual-port nodes. A dual-port node receives a packet from one radio (generally on one band), and retransmits that packet on a different radio (generally on another band). This is what is known as a bridge. The dual-port TNC is taking a packet off one network (the first radio's frequency) and moving that packet to a second network (the second radio's frequency). Note that both networks are using the same AX.25 protocol.

Contrast the operation of a "bridge" with the operation of a "gateway." The term "gateway" refers to a device that takes a packet off one network, and moves it to a second network that uses an entirely different protocol and transmission medium. For example, you may be familiar with gateway stations that accept AX.25 1200-baud packets from a UHF radio and convert them to Internet (TCP/IP protocol) packets for transmission over the worldwide Internet network. The Internet uses entirely different packet structures, routing schemes, and modulation medium from AX.25. That is what makes the transfer device a "gateway" rather than a bridge.

Now let's relate this to 802.11 networks. The simplest 802.11 device uses a network card (NIC) installed in a laptop or desktop computer. This card performs the function of both a TNC and the radio. The card (along with its driver software) takes the data stream from the computer, and breaks it into packets, and the radio circuit on the card transmits the packets on the microwave band.

Two computers, each with an 802.11 NIC card, can communicate with each other exactly like a packet circuit. In the 802.11 world, this is called "ad hoc" mode.

Another, more common, 802.11 architecture uses a device called an "access point." The access point is akin to what ham packet operators call a "node." In this architecture, the access point receives the packets from one computer, and retransmits them like a digipeater, with the additional administrative overhead of connection management.

Most 802.11 access points, however, also function like a dual-port node, or "bridge." Most of them have a built-in circuit that is capable of moving wireless packets off the radio frequency and onto a wired Ethernet network. How does it do this? Look closely at the packaging of your wireless access point device. It probably has a four or five port "hub" built into it. A hub is a device used on Ethernet networks to connect wired computers together. Most 802.11 access points combine a wireless digipeating node with an Ethernet hub, and a bridge circuit between the two.

And to complete the picture, most commercial access points today also include a gateway circuit to move the 802.11 and Ethernet packets onto the Internet! These devices have a DSL or cable modem port in addition to their Ethernet ports! That is where

Johnny Appleseed began his journey across America from one of these very orchards. Abraham Lincoln's father's homestead is here, along with the birthplace of Woodrow Wilson. Jim Colter, the famous mountain man, was born here, as were three other members of Lewis and Clark's expedition. There is a large community of Old Order Mennonites, who—like the Pennsylvania Amish—still use horses and buggies. The Statler Brothers got their start here. The Shenandoah River, made famous in song, still sports a 19th century covered bridge. The river flows the full length of the valley, eventually emptying into the Potomac at historic Harper's Ferry.

Although the entire population of the valley is smaller than the number of tourists visiting Disney World on a typical day, ham radio is well represented. You would think that on-the-air activity would be limited, since most of the valley lies well within the Radio Quiet Zone surrounding the National Radio Astronomy Observatory (see Figure 1). But the observatory, located just over the mountain in Green Bank, West Virginia, has been

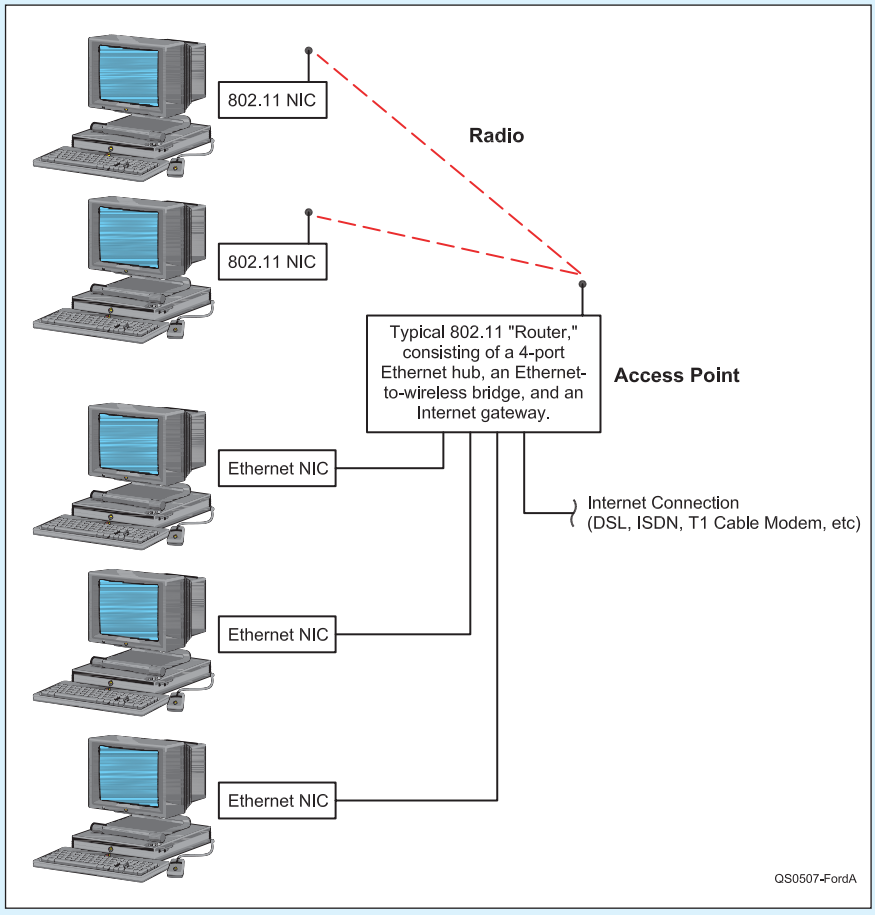

Figure A—A typical wireless network in Infrastructure mode.

the term "access point" derives. These boxes are used to provide "access points" to the Internet via radio.

So you can see these little boxes are versatile devices. They can handle packets in 802.11, Ethernet or TCP/IP formats. They listen to the various channels, and when an incoming packet is received, the packet is switched (routed) to the proper outgoing channel based on the packet's destination address. For this reason, these devices are called "routers."

A network that uses 802.11-equipped computers to communicate using an access point device is said to be operating in "infrastructure" mode. There are slight differences in how the packets are handled between "ad hoc" and "infrastructure" modes, but that is beyond the scope of this article.

The experiments conducted by N4DSL, KG4PRR and K4DJG used an Ethernet connection between the computer and access point on each end of the radio circuit. The Ethernet packets were converted to radio packets by the access point. The radio packets were then transmitted across the valley, with the access point handling the connection, acknowledgments and retries, exactly like traditional packet radio.

more than happy to work with local amateurs. Rick Fisher, KE8DH, an NRAO scientist, leads an annual tour of the facility just for hams. Wesley Sizemore (the interference coordinator) bends over backwards to allow as much ham radio activity as possible, with the absolute minimum of restrictions necessary to protect their sensitive instruments and observations. He has done an admirable job. Indeed, judging by the activity on the ham bands, you wouldn't even know the Quiet Zone was there!

There are almost 600 licensed amateurs in the valley. DX Hall of Famer and four-time DeSoto Cup winner Bob Eshleman, W4DR, got his start here. So did the venerable George Thurston, W4MLE (SK), who was longtime SEC of ARRL's Northern Florida Section. Lifetime valley resident Gerry Brunk, K4RBZ, is on the DXCC Honor Roll. Numerous national nets are called from the area, including the famous Menno-Net.

The valley is home to numerous radio clubs and repeater associations. In the northern town of Winchester, you find the Shenandoah Valley ARC. In the center is the Massanutten ARA (MARA, named after the large mountain sitting smack in the middle of the valley), and in the south is the Valley ARA (VARA). The clubs, especially MARA and VARA, cooperate closely on many projects, including a joint Field Day operation that consistently scores tops in the state and in the Top 10 in their class nationally. All of this from the "Quiet Zone"!

Experimentation = Fun

Valley hams have a rich tradition of experimentation. For such a rural and agricultural area, there is an astonishing amount of activity in new and unusual amateur technologies. For example, the e-mail address of former MARA President David Tanks, AD4TJ, is **davidmoonbounce@aol.com**. Cowles Andrus, K4EME, not only works EME, but is also one of many hams active on the orbiting satellites and meteor scatter. Vic Alger, KE4LKQ, has a satellite station that receives and filters data from EMWIN (the Emergency Managers Weather Information Network sponsored by the National Weather Service) and posts pictures from the weather satellites.

There is significant activity on PSK31, JT44 and other new digital modes. There is a repeater for fast-scan television, a monthly meeting for QRP homebrewers (surface mount, anyone?), and a lunch group that might best be called "Ham Radio Aboard Recreational Vehicles."

The W4PNT and KC4GXI repeaters are Echolink nodes, allowing users of 2 meter and 70 cm transceivers to tie directly to the Internet using Voice Over Internet Protocol. At any given time you can find more than a dozen active Automatic Position Reporting Stations, including several remote weather monitoring stations—and several hams operate APRS bicycle-mobile. The KB4OLM DX cluster is accessible from almost everywhere in the valley. Yes, all of this is right here in the "Quiet Zone"!

The experimentation is not limited to radio, either. Wind generators and solar panels provide power for packet-based weather reporting stations atop Shenandoah Mountain and the James Madison University campus. *QST* has even featured an article by valley ham Chuck Vogan, KD5KA, who crossed the country

working QRP from his motorcycle.

So it was no surprise to anyone when a local ham—who is not even old enough to buy a beer!—proposed experimenting with 802.11 gear on the 2.4 GHz microwave bands.

"To 'B' or Not to 'B'—That is the Question"

While most kids his age were learning their state capitals, Jason Armentrout, N4DSL, was building computers in his bedroom. He earned his Amateur Extra license at age 15 and was one of the first hams on PSK31. Now just barely 18 years old, he works professionally for a local cable company, designing, installing and configuring wired and wireless networks.

In early 2003, Jason became curious about the distances achievable with commercial 802.11 gear. The 802.11 equipment on the market today comes in three common flavors: 802.11a, 802.11b and 802.11g. The differences among them are transparent for most users, and lie mainly in the frequencies used and speeds of transmission. 802.11b and 802.11g advertise usable distances up to about 300 feet.

The 802.11a equipment operates on the 5 GHz band. The 802.11b and 802.11g gear is on the 2.4 GHz band, and of the 11 channels available for operation in the US, channels 1 through 6 overlap the 2.4 GHz amateur band. This makes the "b" or "g" version much more suitable for experimentation, especially for applications that might result in signals outside the allowable parameters of the gear's Part 15 certification.

The 802.11 equipment is designed to establish Ethernet-like wireless networks at high (Internet useful) speeds. Most 802.11 access point devices also are capable of patching the wireless packets into the Internet (converting 802.11 protocol packets to TCP/IP packets: a gateway). See the sidebar, "IEEE 802.11 Jargon: 'Routers, Gateways and Bridges... Oh My!'" The question Jason asked was: Can 802.11 equipment be used at distances that make it practical to establish a "valley-wide" high-speed TCP/IP-compatible network? In other words, can a "Hinterland" network really cover such a wide area using nothing more than the cheap off-the-shelf gear?

Jason conferred with amateurs at the club meetings and the daily informal lunches at the local Burger King. He learned there are two ways that hams traditionally go about experimenting with new modes. The first is to build from scratch, putting together components and basic building blocks to form a communications system. For microwaves today, this means buying magnetrons, feedhorns, transverters, waveguides, exciters and so forth. This approach might result in a good long distance RF link. But the interface between that RF link and today's TCP/IP-based packets would require significant protocol work, possibly even a custom-designed bridge or gateway. Jason was interested in building an operational network, not tinkering around with protocol layers.

The second way hams do things is to adapt surplus or commercially available gear for ham use. Being familiar with commercial network gear, Jason asked the next logical question: "Could 802.11 gear be utilized to yield useful computer communication over distances in excess of the typical radius of 250-300 feet?"

Figure 2—K4DJG mounted the dish antenna to his trailer hitch, making for a convenient "rover" microwave station at Reddish Knob at the western end of the 34 mile link.

Eureka!

On a cold Saturday morning in early 2003, Jason, Daryl Coffman, KG4PRR, and David Fordham, KD9LA, linked

Physical Description of the Link Used in the Initial Experiments

Station 1

Call sign: N4DSL
Location: Two-Mile Run Overlook on Skyline Drive
Coordinates: N 38° 17.906' latitude, W 78° 38.914' longitude
Elevation: 2770 feet ASL

Station 2

Call sign: K4DJG
Location: Reddish Knob
Coordinates: N 38° 27.946' latitude, W 79° 14.717' longitude
Elevation: 4370 feet ASL

Both stations used a laptop computer equipped with a 10BaseT Ethernet NIC adapter. The NIC was connected to a Linksys WET11 Ethernet wireless bridge, configured in ad-hoc mode.

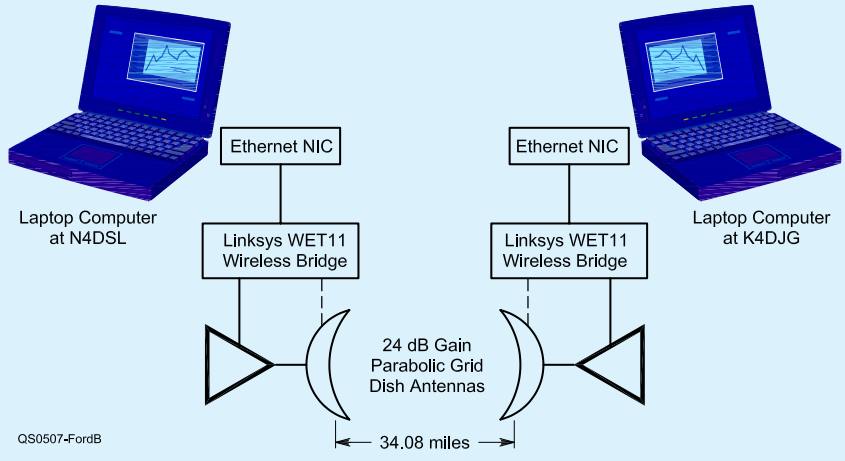

laptop computers using 802.11 equipment across a distance of about 3 miles. Using simple directional "barbecue grill" dish antennas, they drove from hilltop to hilltop, attempting to establish reliable data transfer between their laptops from high spots farther and farther apart. These first experiments were characterized by maddening, unexplainable intermittent performance. Sometimes the link worked flawlessly, and suddenly it would just disappear. Hours and hours of testing and diagnostics finally located the problem: a faulty connector on the coax feeding one of the dishes!

Then early in 2004, the Eureka Moment arrived. After obtaining a set of amplifiers, Jason and Bob Van Fossen, K4DJG, set up stations on opposite sides of the valley. Jason set up at an overlook on Skyline Drive in Shenandoah National Park. Bob drove to Reddish Knob in the George Washington National Forest. They were able to establish a solid and permanent connection using the setup shown in Figure 2 and the sidebar, "Physical Description of the Link Used in the Initial Experiments." The distance as calculated by a popular mapping program (Figure 3) was 34.08 miles!

Then for the icing on the cake: Taking the amplifiers out of the circuit, they were still able to establish a solid connection! And get this: The signal path ran directly across the city of Harrisonburg, two college campuses, a hospital, a large shopping mall, and a commercial district, all of which are known to be using active 802.11b wireless equipment!

Antennas on each end of the link were 24 dB gain parabolic grid dishes (the so-called "barbecue grill antennas"). The power with the amplifiers turned on was 800 mW (about 200 W ERP). Without the amplifiers, the power was 30 mW, or about 7.5 W ERP.

A surprising finding of the exercise was that the data transfer rate was about the same regardless of whether the amps were in or out of the circuit. With the amps on, the signal strength was "100%" as determined by the indicators on the bridge. Without the amps, the signal strength dropped to "65%." The nominal data rate remained at 1 Mbit/s regardless of the status of the amps.

A disappointing surprise was the actual data throughput rate. While the nominal data rate (the rate at which the signal modulation was pumping out the bits) was 1 Mbps, the actual data transfer rate, as measured by an FTP file transfer, was more like 150 kbit/s. This is still faster than landline ISDN service, but a far cry from the nominal. What could account for the low throughput?

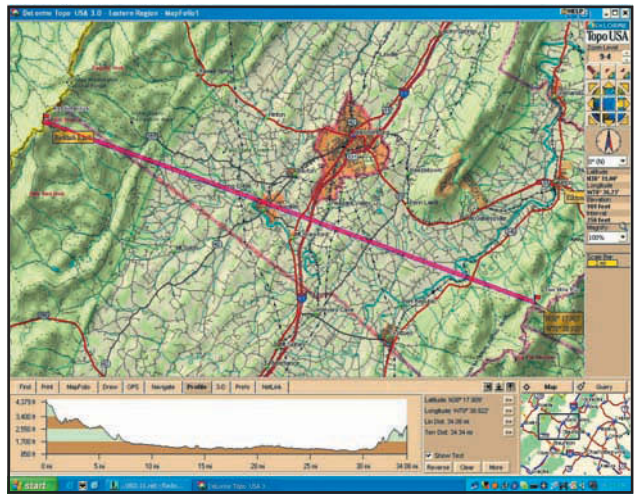

Figure 3—Initial testing showed that reliable connection paths could be achieved across the valley without power amplification, using only directional antennas on each end. The pink line shows the radio signal path, spanning over 34 miles. The western end of the path was on Reddish Knob in the Allegheny Mountains. The eastern end was at Two Mile Overlook on Skyline Drive in the Shenandoah National Park. A mapping program such as Delorme's Topo USA can be helpful in evaluating and planning possible microwave paths.

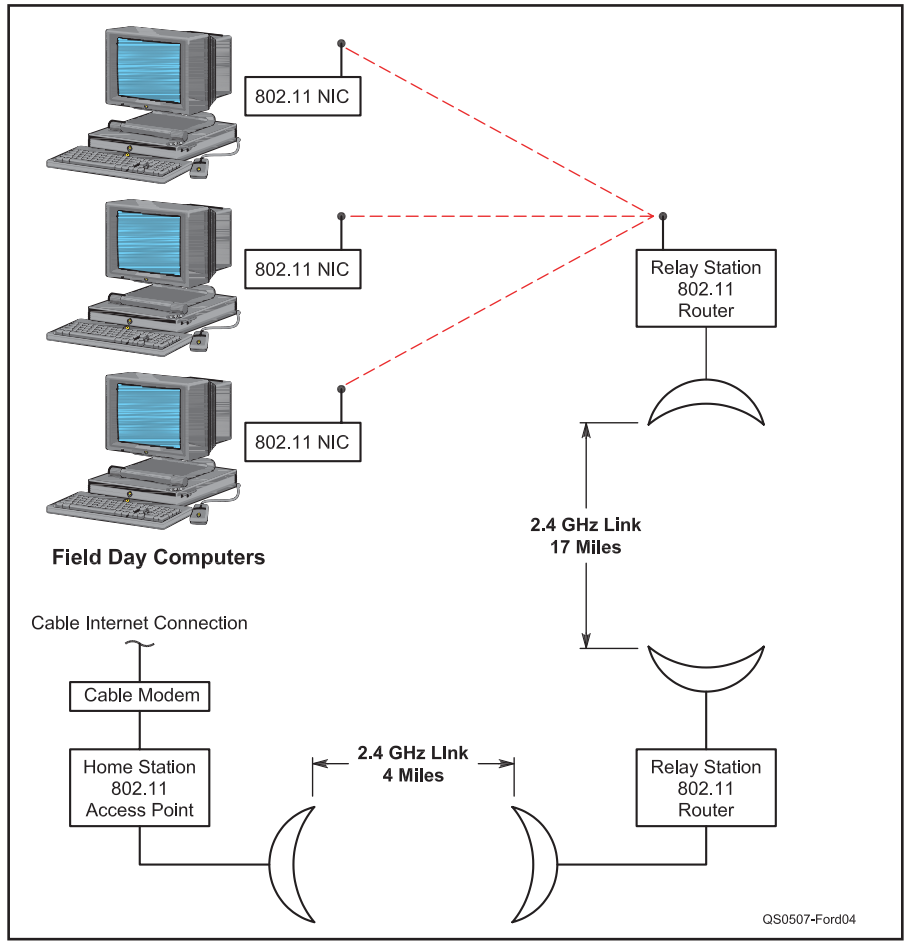

Figure 4—Diagram of the network that the MARA and VARA clubs set up for Field Day 2004.

As mentioned above, the path crossed numerous other networks. Interference from these networks (for example, carrier-sense delays) might be part of the problem. But Jason proposed another explanation.

The IEEE 802.11 standard is designed for transfer distances of 250 to 300 feet. The wait times and retry times are optimized for packet movement across such short distances. At the speed of light, the difference between 300 feet and 34 miles is significant. Jason postulated that the time interval (taken for a packet to travel the 34 mile distance, be received and decoded at the destination, and the acknowledgment sent across the 34 mile return trip) was so long, the originating station had already given up and was busy sending a retry.

Figure 5—The Field Day site in the mountain meadow featured an IEEE 802.11b wireless network, connected to the dish antenna seen atop the mast in the center of the picture.

Figure 6—The IEEE 802.11b equipment at the Field Day site was powered by solar panels located at the base of the mast. The router, designed for indoor use, had to be wrapped in plastic to protect it from the morning dew. A valley inside-joke is that "It never rains at Field Day, although sometimes we have a heavy dew."

By calculating the retry timing parameters for the 802.11 protocol, and taking into account the speed of radio signals through air, Jason determined that the maximum practical distance using the protocol was about 10-12 miles before you begin to experience major problems with the retry timing windows. What this means is, if you were using a path of only 10-12 miles, the actual data throughput rate should increase tremendously. And on a path of a mile or two, the throughput rate should approach the nominal 1 Mbit/s signal rate (assuming no interference, of course).

Even with the retry problems at 34 miles, a throughput of 150 kbit/s is significantly better than a dial-up modem or ISDN service, and over 100 times faster than 1200 bit/s packet. This means that 802.11 equipment is a suitable approach for a valley-wide Internet-compatible network.

Of course, if this equipment is being used in ways that place its operation outside the Part 15 certification parameters, care must be used to ensure compliance with Part 97 rules. Operation of the equipment described with the 800 mW amplifiers is in clear violation of the Part 15 rules, but the operation of the unamplified system clearly falls under the Part 15 and Part 2 rules as they have recently been rewritten. In fact, the described system could operate as a Part 15 system by utilizing some of the 200 mW 802.11 hardware currently available from a number of sources. The EIRP of such a system would be 50 W, which is lower than the maximum allowable power in a point-to-point link under Part 15.

If the system requires power in excess of the Part 15 limits, then Part 97 allows amateurs to operate under a less-limiting rule structure. If the equipment is to be used in a Part 97 system, then the equipment must be configured to ensure usage within the ham band (which means it must operate between channel 1 and channel 6) and provisions must be made for the 10-minute station identification. In addition, content controls must be implemented to ensure that no commercial transactions are carried by the network. But this type of experimentation is exactly what ham radio is all about.

One other consideration is avoiding interference to AO-43 satellite users. This can be accomplished by avoiding operations on channel 1.

Into the Wilderness

Point-to-point communication is good, but networks are even better. Once it was learned that long-distance connections were achievable, it was a simple matter to construct a multinode network and to tie the system into the Internet. Okay, so why would you want to construct a long-distance multinode network with Internet connectivity? Two words: Field Day!

Each June, the MARA and VARA clubs work together and build a "city" in the wilderness of the George Washington National Forest. In a picturesque mountain meadow near Flagpole Knob on the western side of the valley, nearly 100 hams and their families gather in tents and RVs, far from the interference of modern civilization. To connect such a remote location to the Internet infrastructure would provide the ultimate test for a long-distance computer network.

Figure 4 shows the basic layout of the network. At the Field Day site, Jason first established a standard 300-foot-radius 802.11g wireless network, the kind found in thousands of hotels, schools, and office buildings. All of the computers used for Field Day logging were equipped with your plain vanilla, garden-variety 802.11 net-

Figure 7—The relay station featured two dish antennas, one pointed at the mountaintop, the other at the tower located in the K4RBZ's backyard. Both dishes were connected to a single IEEE 802.11b router.

Figure 8—Daryl, KG4PRR, installing the dish antenna at the home of K4RBZ.

Figure 9—The final antenna was connected to another IEEE 802.11 access point, which was wired to K4RBZ's cable modem. N4DSL is an employee of Adelphia Cable Company, which provided the connection to the Internet infrastructure. That's K4DJG in the background.

work cards. Additionally, the mountaintop network served any personal 802.11-equipped laptops within the 300-foot omnidirectional range of the access point. Bryan Fordham, KG4JOE, even used the network with his Toshiba handheld PDA.

One of the nodes on this network was an access point, connected to the directional antenna in the middle of the Field Day site (see Figures 5 and 6). This directional antenna was aimed at a relay station 17 miles away at a hilltop repeater site down in the valley (Figure 7). Signals from the mountaintop network destined for the Internet traveled 17 miles down this path to the valley.

The relay station was an 802.11 router connected to two directional antennas. One of these was pointed at the mountaintop meadow Field Day site, and the other pointed to a dish located four miles away on a tower at the home of Gerry, K4RBZ (Figure 8).

The dish at K4RBZ was connected to a third 802.11 access point, which in turn was wired to a standard commercial cable modem on the Internet infrastructure (Figure 9). Thus, the mountaintop wilderness computers were connected via the dish antenna to the relay station, to Gerry's cable modem, using off-the-shelf components.

Talk about experimentation with practical application! It was great to sit in your hammock enjoying the cool mountain breeze of a forest clearing, notebook computer in your lap, checking your e-mail and browsing the Web while watching the white-tail deer saunter through the wildflowers! One of the most important uses of the Internet connections was checking the latest weather radar images for approaching thunderstorms. The network was also used to verify the latest elements for the orbiting space station, and the Field Day operation made contact with N1ISS!

And remember, all of this was taking place in the Radio Quiet Zone!

The Next Step?

After the initial tests were completed, the group located an 802.11 card that allows changes to the retry time parameter. With the new cards, they successfully passed the 56-mile mark with great data transfer rates. A 72-mile path is next.

These experiments demonstrate the suitability of using commercial 802.11 equipment for relatively high-speed point-to-point and network data communication across relatively long distances. But what about mobile operation? If reliable connections can be made from one side of the valley to the other, and from the peaks to the valleys, can a series of these dishes be linked together to provide a blanket of coverage sufficient to allow mobile operation?

At a distance of 12 miles, a dish with a 10° beamwidth will cover a swath of real estate more than 2 miles wide. An array of six or eight dishes located on one side of the valley might be able to offer usable links to most of the hams in Rockingham County. By using a second array for Augusta County, and linking the arrays together, a truly wide area high-speed TCP/IP compatible network should be possible. And because each of the dishes would cover a smaller area, the station-to-station interference would be reduced, somewhat akin to the cellular telephone networks.

These experiments used directional antennas on both ends of the link. What is the practical usable distance when one of the stations uses a directional dish and the other station (perhaps a mobile) is using an omnidirectional antenna? A characteristic of the 802.11 protocol is that if the signal strength decreases, reliable communications can still continue, albeit at a slower rate. This makes using traditional signal-strength calculations and tables problematic.

Reducing the distance from 34 miles to 12 miles will increase throughput. But replacement of one of the dish (gain) antennas with an omnidirectional antenna for mobile use will reduce signal strength, decreasing throughput. What will the trade-off ratio be? Will using a bigger dish (higher gain) on one end of the link compensate? What effect will the amplifiers have if put back into the circuit?

These questions will be giving Jason and the Shenandoah Valley hams an excuse to continue the fun through the coming months. In the meantime, the entire crew here in the Quiet Zone is interested in corresponding with other hams who are experimenting in this area. If you are using 802.11 gear on the ham bands in ways not covered by the Part 15 rules, please contact Jason at **n4dsl@atrs.com**.

David R. Fordham, KD9LA, has been a ham since 1974. He holds a CPA, CMA and PhD, and is currently PBGH Faculty Fellow and Professor of Information Technology at James Madison University. He is trustee of the JMU Wireless Experimenters station WN4JMU, and is past president of the Massanutten Amateur Radio Association (MARA). He edits and publishes the Monitor, *the joint monthly newsletter of MARA and the Valley ARA. You can reach him at 131 Wayside Dr, Weyers Cave, VA 24486;* **fordhadr@jmu.edu**.