

Wireless Data Network Infrastructure at Carnegie Mellon University

ALEX HILLS

Vice Provost and
Chief Information Officer
Carnegie Mellon University
Pittsburgh, PA 15213
alex.hills@andrew.cmu.edu

DAVID B. JOHNSON

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213-3891
dbj@cs.cmu.edu

In order to support mobile computing research, including the development of software which will allow seamless access to multiple wireless data networks, we are building a wireless data network infrastructure at Carnegie Mellon University. This infrastructure will allow researchers and other members of the campus community to use mobile computers to gain access to data networks while they are on-campus or while they are off-campus in the greater Pittsburgh area. The infrastructure will initially include two different types of wireless networks, a low-bandwidth wide area system and a high-bandwidth local area system, each of which will provide access to our campus computer network. Since our campus network is called "Andrew" (after Andrew Carnegie and Andrew Mellon), the new wireless infrastructure has been dubbed "Wireless Andrew."

This article describes the Wireless Andrew infrastructure we are building. We begin with an overview of the infrastructure and the characteristics of the two types of wireless networks we are using. Next, we describe in more detail each of these networks. We discuss the technology used, our implementation of it in our wireless network infrastructure, and our current deployment experience. Finally, we summarize the lessons we have learned and present conclusions.

Overview of Infrastructure

Our approach to wireless data networks is based on the notion that it is unlikely that a single wireless network able to meet all mobile computing needs will evolve. It is far more probable that many wireless networks will be available, each of which will provide service over a variety of geographical coverage areas at various speeds and at a variety of price levels. Each network will serve a niche, but none will meet all needs. Accordingly, mobile computer users will need to be able to access multiple networks in order to meet their needs.

A mobile computer, in order to take advantage of the highest speed or most capable network available, will need to be able to work with any of two, three, or even more wireless networks. Further, a mobile user would in general like to be connected to the wireless network able to meet the application's needs at

lowest usage cost. When moving beyond the range of one wireless network, a user would generally like to be automatically connected to another.

As part of the Wireless Andrew infrastructure at Carnegie Mellon, we are building a high-speed wireless data network on campus, using AT&T WaveLAN wireless local area network (wireless LAN) technology to provide high-speed (2 Mbps) service over a large portion of the campus. In addition, outside the wireless LAN service area, we are using the lower speed (19.2 kbps) Cellular Digital Packet Data (CDPD) service [5]. With both the CDPD and wireless LAN networks operational, properly equipped mobile computers will be able to use both networks. Table 1 summarizes some of the important characteristics of these two networks.

We are currently using laptop computers equipped with PC Card (PCMCIA) wireless LAN interfaces and on-board CDPD modems, and we expect to continue their use in the near term. We are developing software for use on this infrastructure, with which mobile computer users will be able to seamlessly move between the wireless LAN and CDPD networks [11]. The infrastructure is intended to be used both by researchers developing the software and by the faculty, staff, and students of the Carnegie Mellon community.

Low-Bandwidth Wide Area Component

In order to support mobile connections to our campus network from throughout the greater Pittsburgh area, we have elected to use the CDPD service [5] offered by Bell Atlantic NYNEX Mobile. Like other wide area wireless services, it operates at a low data bit rate. Since CDPD supports the Internet Protocol (IP) [16], CDPD service can be easily integrated into an existing IP network.

Details of Technology Used

Cellular Digital Packet Data (CDPD) is a relatively new service for wide-area data communication with wireless mobile hosts. The original version of the CDPD System Specification (Release 1.0) [5] was developed by IBM and released in July 1993 by a consortium of the major United States cellular telephone carriers consisting of Ameritech Mobile Communications, Bell Atlantic Mobile Systems, Contel Cellular, GTE Mobile Communications, McCaw Cellular Communications, NYNEX Mobile Communications, PacTel Cellular, and Southwestern Bell Mobile Systems. The CDPD standard is now controlled by an open trade organization known as the CDPD Forum, which released a revision of the CDPD standard (Release 1.1) in January 1995.

CDPD uses the existing infrastructure of the analog AMPS (Advanced Mobile Phone Service) cellular telephone network to transmit data, utilizing idle channels in the cellular system to provide a connectionless

Table 1 Carnegie Mellon “Wireless Andrew” Infrastructure Network Characteristics

	CDPD	WaveLAN
Base station coverage radius	1–10 miles	100–800 feet
Raw transmission data bandwidth	19.2 kbps	2 Mbps
Typical user data throughput	9–11 kbps	0.5–1.5 Mbps
Typical roundtrip latency between mobile and stationary computer	450–600 msec	5–30 msec

digital data packet service. When a channel is not being used for voice service, it can be used to send and receive packets from a CDPD-equipped mobile station. Existing cellular channels can thus be shared between voice and data service, although CDPD uses only time *between* voice calls, not silence periods *within* any voice call for sending data; voice service always has priority over CDPD service for access to a channel. A CDPD cellular carrier may also remove some channels from voice service and dedicate them exclusively to CDPD service, although this is not the normal configuration. CDPD can be used whenever one is within range of a CDPD-equipped cell site. The coverage area of each cell site is typically in a range from one to ten miles in radius, depending on the configuration installed by the cellular carrier. An array of cell sites normally covers all of a metropolitan (or larger) area.

CDPD operates at layers 1 and 2 of the OSI model (Physical and Data Link) and provides service to either of two layer 3 (Network) protocols: IP [16] or the ISO Connectionless Network Protocol (CLNP). The service operates at a raw data bit rate of 19.2 kbps, although actual throughput is well below that figure. The deterministic protocol overhead bits required for additional levels of headers (beyond IP) and for Forward Error Correction (FEC) coding and other functions reduces the available bandwidth to a theoretical limit of 13.3 kbps from the mobile user to the base station and 12.8 kbps from the base station to the mobile user [7]. This bandwidth is further reduced by contention in gaining access to the channel. In our own measurements in Pittsburgh, the usable throughput of CDPD is between 9 and 11 kbps.

The current CDPD standard includes two types of compression that may improve the usable transmission throughput. The original CDPD standard (Release 1.0) provided for use of “header compression” for packets carrying transport and network level protocols that support header compression. For example, TCP header compression [10] removes the fields from the combined TCP [17] and IP [16] headers that do not change between sequential packets on the same TCP connection; the 40 bytes of combined TCP and IP headers can typically be reduced to 3 bytes. The headers for other IP transport protocols, such as UDP [15], are not compressed, and TCP header compression is not supported by all CDPD implementations. In particular, the Bell Atlantic NYNEX Mobile system in Pittsburgh which we are using does not yet support TCP header compression. In addition to header compression, Release 1.1 of the CDPD standard added support for V.42bis data compression, as is used on recent circuit switched telephone modems. V.42bis can compress some data by as much as a factor of 3 or 4, but few CDPD systems or modems, including the Bell Atlantic NYNEX Mobile system in Pittsburgh, have yet been upgraded to support Release 1.1 of the standard.

A CDPD network serves a mobile computer (Mobile End System, or M-ES) through a radio link to a nearby cell site equipped with a Mobile Data Base Station (MDBS), as shown in Figure 1. The MDBS, in turn, communicates with a Mobile Data Intermediate System (MD-IS), which routes packets to and from the MDBSs and is typically located at the cellular carrier’s mobile telephone switching office (MTSO). Each MD-IS serves a number of MDBSs, and there may be one or more MD-ISs in a CDPD network. Also part of the CDPD network may be a number of “off the shelf” routers (Intermediate Systems, or ISs), which connect to conventional IP or CLNP networks on which are located stationary host computers or servers (Fixed End Systems, or F-ESs). A “CDPD modem” is typically used to provide the CDPD functions of an M-ES; when a customer initially establishes CDPD service with a cellular carrier, the carrier assigns an IP (or CLNP) address to the M-ES (to the CDPD modem), which is used for routing packets to the M-ES.

Costs experienced by CDPD users are based on the number of packets sent and the number of kilobytes of data sent between the M-ES and MDBS. This differs significantly from the pricing model used by cellular carriers for conventional circuit switched voice service, where prices are based on a charge per minute of connect time. There are no connect time charges on CDPD, as there is no connection. To use CDPD, the CDPD modem does not dial a telephone number and connect to a remote system. Instead, individual data packets (either IP or CLNP) are sent over the air to the MDBS using the CDPD “air interface” protocol, for forwarding through the MD-IS to other M-ESs registered with CDPD, or to any F-ESs or ISs on any wired network connected to the CDPD system.

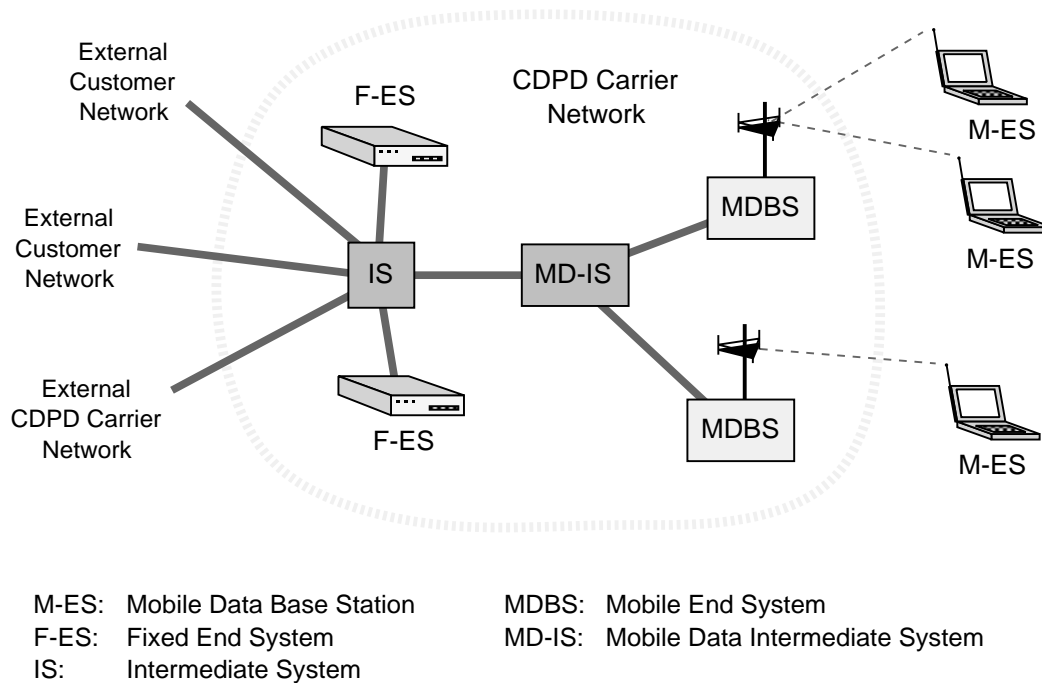


Figure 1 CDPD network architecture

In each cell, the MDBS is responsible for channel usage. The MDBS selects the channel to be used for CDPD within the cell, and “hops” from one channel to another if the channel it is using is taken by the cell for voice service. The M-ESs in the cell will follow the MDBS to the new channel, either by receiving an explicit channel hop command from the MDBS or by scanning all cellular channels to find the one being used by CDPD. When an M-ES initially begins use of CDPD (such as at boot time), it must also scan all cellular channels to find the one being used by CDPD.

Figure 2 depicts a simplified view of the processing steps performed in sending a packet over the air in a CDPD system. First, the header of the packet may be compressed, such as with TCP header compression [10], as described above. The packet is then divided into *segments*, and a segment header is added to each; the segment header includes an identifier indicating the type of header compression applied to the packet (present in all segments, but only meaningful in the first segment of a packet), and a protocol number to indicate which type of network protocol (IP or CLNP) is contained in the segment. The data of each segment is also encrypted using the RC4 encryption algorithm. Next, each segment is encapsulated with the addition of a data link header to form a *frame*; the frame header includes sequence numbers to maintain the order of frames in transit, and an identifier for multiplexing separate data link connections (such as to separate data packets from CDPD administrative packets). The resulting frames are then concatenated together, with flag bits added to delimit the individual frames, and the resulting bit sequence is divided into a series of fixed-sized *blocks*; each block also has Reed-Solomon Forward Error Correction (FEC) codes added, allowing errors in transmission to be corrected by the receiver. The resulting stream of blocks is then sent digitally over the 30-KHz AMPS channel being used by CDPD.

The *forward link channel* (from MDBS to M-ES) is a broadcast channel with no contention. The MDBS simply transmits blocks on this channel for any of the M-ESs registered in the cell. Interspersed in this data stream are “channel idle” and “decode status” flag bits, used to approximate the *carrier sense* and *collision detection* functions of CSMA/CD in IEEE 802.3 (Ethernet) networks, using a protocol called DSMA/CD

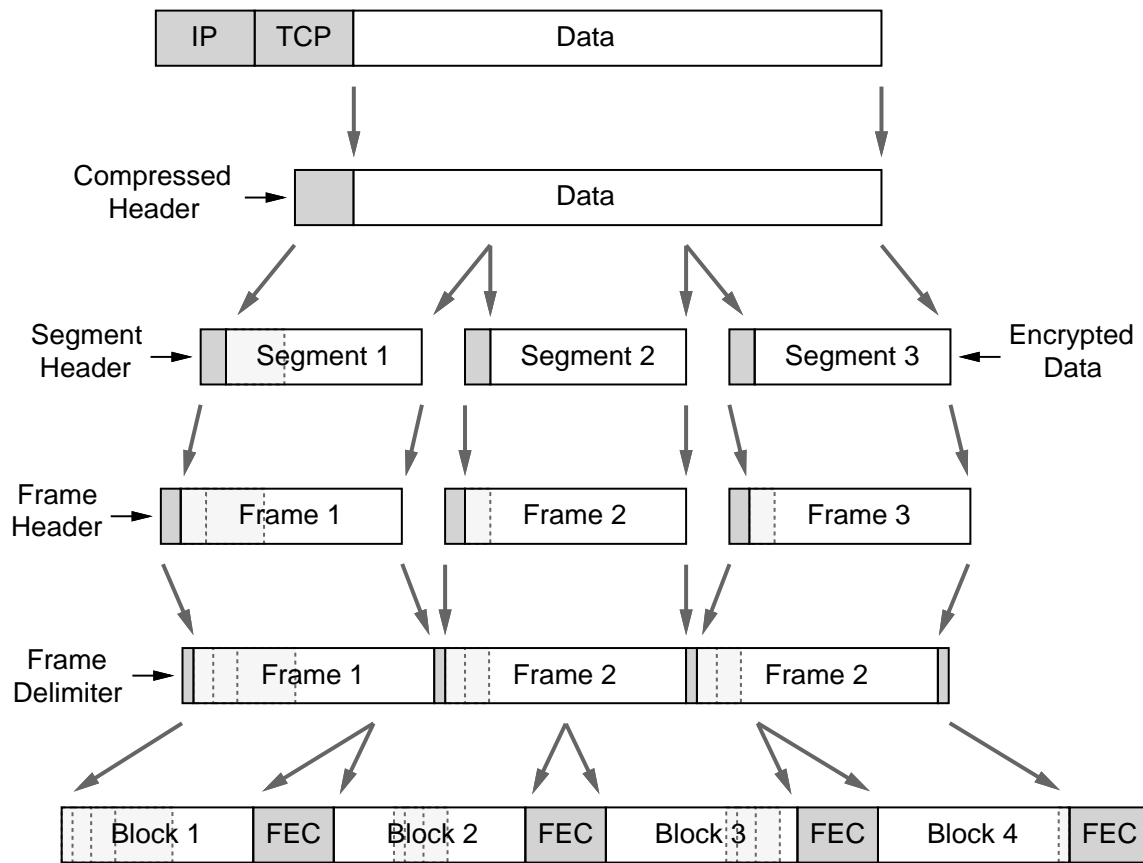


Figure 2 CDPD packet transmission processing

(Digital Sense Multiple Access with Collision Detect). An M-ES wishing to transmit a block on the *reverse link channel* (from M-ES to MDBS) listens to the forward link channel until a flag bit from the MDBS indicates that the reverse link channel is idle. It then transmits its block on the reverse link channel and checks the following decode status flag bit on the forward link channel. If the flag bit indicates that the block was not successfully received by the MDBS, a collision or other wireless transmission error has occurred, and the block is retransmitted later by the M-ES. After transmitting one block, an M-ES may continue to transmit consecutive blocks (up to a maximum number of blocks) without needing to reacquire the channel, as long as the decode status flag bit indicates that the previous block was successfully received by the MDBS.

Implementation Details

To use CDPD, the service must be ordered from a cellular carrier, and CDPD modem equipment for the M-ESs must be purchased; furthermore, to allow communication between M-ESs and F-ESs on an existing data network, a wired connection from the CDPD system to the existing network must be installed. We initially selected the Pacific Communications Sciences Inc. (PCSI) Ubiquity 1000 modem when we began using CDPD in October 1994 because it was the only CDPD modem available at that time capable of on-board installation in a portable computer. The Ubiquity 1000 is designed to be installed in the internal floppy disk drive bay, replacing the floppy drive, on IBM ThinkPad 750 series, 755C, or 755Cs laptop computers; it includes a small telescoping antenna that integrates with the ThinkPad. We are now also using

the IBM Wireless Modem for Cellular/CDPD, which is designed to be installed in a PC Card (PCMCIA) Type II slot and includes a small external radio and antenna unit.

The CDPD network can provide communication between M-ESs without use of wired data networks, since an MD-IS can route packets between the MDBSs serving two M-ESs. To provide connections to a fixed host computer or to a wired data network, however, it is necessary to connect a router associated with the CDPD network to the fixed host or to a router which is part of the wired network. Since CDPD can carry IP packets, it is natural to imagine a CDPD network connected directly to the Internet. However, this is not generally possible yet for many CDPD carriers due to the restrictions placed on the regional Bell operating companies as a result of the 1982 AT&T divestiture. This limitation is expected to change at some point in the future, and CDPD customers are currently able to connect their own wired networks into the CDPD system through their CDPD cellular carrier.

Since we wanted to provide our mobile computers with access to all of Carnegie Mellon's campus network, we needed to install a wired connection to Bell Atlantic NYNEX Mobile's CDPD network. We purchased a Cisco router specifically for this purpose and linked it to a Wellfleet router already installed at Bell Atlantic NYNEX Mobile, as illustrated in Figure 3. The connection between the routers is accomplished using a 56 kbps Frame Relay permanent virtual circuit. Although the 56 kbps service is adequate for our immediate needs, we plan to upgrade it to a T1 (1.544 Mbps) link as our CDPD usage increases.

On the mobile computers, we are currently running primarily Microsoft Windows, and are using FTP Software's PC/TCP implementation of the TCP/IP networking protocols. PC/TCP provides a number of Internet applications, including telnet, FTP, and electronic mail, and is compatible with the Windows Sockets (WinSock) API [8]. We are also running the NetBSD version of Unix on the mobile computers using CDPD in some research projects.

The Ubiquity 1000 modem provides a serial port interface to the IBM ThinkPad through the floppy disk drive connector inside the ThinkPad's case. The modem implements the Serial Line IP (SLIP) protocol [18] on this serial port for communicating IP packets to and from the ThinkPad, and we are using the built-in SLIP implementation in PC/TCP and NetBSD. The IBM Wireless Modem for Cellular/CDPD includes an NDIS [1] device driver for Microsoft Windows, allowing PC/TCP to interface to the modem on the PC Card (we are not yet using the IBM modem under NetBSD). With both modems, though, the protocol used to interface to the CDPD modem is separate from the protocol used over the CDPD network, which is defined by the CDPD standard.




Figure 3 Interconnecting CDPD and the campus network

Deployment and Experience To Date

With the establishment of the CDPD-to-campus Frame Relay link in October 1994, wireless access to our campus network became possible from anywhere in the greater Pittsburgh area. When cellular carriers put roaming arrangements in place for CDPD, the coverage area will, of course, be much larger. The new service has, so far, been used primarily for email, and it is quite adequate in that application. Subjectively, the response time and speed are close to what one would experience using a 9.6 kbps dial-up telephone connection through a modem, although the high latency (about 500 msec) is noticeable in interactive usage such as when typing over a telnet connection.

Coverage and capacity planning are serious issues for cellular carriers, but since a CDPD user is buying service from a carrier, the user need only ascertain that the service provided by the carrier is adequate. Although the Pittsburgh area is composed of hilly terrain, the base stations (cell sites) in the Bell Atlantic NYNEX Mobile network have been located appropriately for this terrain. CDPD coverage is thus comparable to what one experiences in Pittsburgh or another city when using a cellular telephone. In Pittsburgh, the CDPD coverage extends over most of the Bell Atlantic NYNEX Mobile territory around the city and surrounding areas.

When signals are weak or voice service usage is heavy, a user may experience a significant delay as the M-ES seeks to access the CDPD channel. As CDPD becomes more popular, this could become a serious problem, although we have not yet found it to be so. There are occasional delays of at least several seconds, though, when the MDDBS hops to a new channel and the M-ES must scan the channels to find the MDDBS again. Such delays can be somewhat more common during times of the day having high demand for cellular voice service (such as in the afternoon as commuters drive home from work) but we have not yet found this to be a problem.

The CDPD standard itself is still quite new and cellular carriers are still working to fully deploy it and gain experience with it. Based on our own experience to date with CDPD, we believe it will be a valuable component of our wireless data network infrastructure, allowing us to extend the coverage of our wireless infrastructure to seamlessly provide support for mobile users in the greater Pittsburgh area.

High-Bandwidth Local Area Component

To complement this low-bandwidth CDPD service and to provide an infrastructure with much higher speed capability, we are building a high-speed wireless network on the Carnegie Mellon campus as part of Wireless Andrew. We are using AT&T WaveLAN wireless LAN equipment [20] to provide this high-speed service.

Details of Technology Used

Wireless LANs were originally intended to allow LAN connections where premise wiring systems are inadequate to support conventional wired LANs. Now, because the equipment is available in PC Card (PCMCIA) form factor, wireless LANs have come to be identified with mobility. Wireless LANs can be used to provide service to mobile computers throughout a building or throughout a campus.

Many wireless LANs operate in the unlicensed ISM (Industrial, Scientific, and Medical) bands at 915 MHz (902–928 MHz), 2.4 GHz (2.4–2.4835 GHz), and 5.7 GHz (5.725–5.85 GHz). Spread spectrum techniques are used in these products, typically direct sequence spread spectrum (DSSS) at 915 MHz or frequency hopping spread spectrum (FHSS) at 2.4 GHz or 5.7 GHz. This use of spread spectrum is required for license-free operation in the ISM bands by the FCC in the United States and its equivalent in other countries. Other wireless LAN products are available which operate in the licensed 18 GHz band and at infrared frequencies [4].

Wireless LANs operate at speeds up to a few Mbps, with radio ranges typically from 50 feet to over 1000 feet, depending on the specific product and the environment in which it operates. Many of these products can be interfaced directly to IEEE 802.3 (Ethernet) or 802.5 (token ring) wired LANs. A new standard, IEEE 802.11, is being developed to allow interoperability between wireless LANs [9], but since the IEEE 802.11 standard has not yet been adopted, wireless LANs currently use proprietary protocols and generally do not interoperate with each other. The 802.11 standard, however, is near adoption, and purchasers of wireless LAN equipment should take care to understand whether there is a path by which their equipment can become IEEE 802.11-compatible.

We selected AT&T's 915 MHz WaveLAN wireless LAN equipment for use in our network. WaveLAN equipment is also available for the 2.4 GHz ISM band, although this version has not yet been released as a product in the United States. Hardware compatible with AT&T WaveLAN is also sold by Solecetek Corporation and by Digital Equipment Corporation.

WaveLAN uses direct sequence spread spectrum (in both the 915 MHz and 2.4 GHz products) to spread its transmissions over a wide bandwidth [2], allowing its transmissions to be robust against various kinds of interference and multipath effects. With direct sequence spread spectrum, a "pseudo-random" spreading code is used, consisting of a sequence of bits (called "chips"), to modulate the transmitted data signal. At the receiver, the received signal is correlated using the known spreading code to recover the data. WaveLAN operates at a raw data bit rate of 2 Mbps, which is spread using a spreading code of 11 chips per bit, for a transmission rate of 22 megachips per second. All WaveLAN units use the same spreading code, which is fixed in the WaveLAN hardware.

Transmissions use a medium access scheme called CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) [2], which is similar to the CSMA scheme used in IEEE 802.3 (Ethernet) LANs. With wireless transmissions, the Collision Detect (CD) technique used in wired LANs cannot be done effectively, since a transmitter's signal strength at its own antenna will be so much stronger than the signal from any other transmitter. Instead, CSMA/CA adds a number of features to the basic CSMA scheme to greatly reduce the number of collisions that might occur if only CSMA (without CD) were used. With CSMA, a transmitter listens for "carrier" (the presence of another transmission) before attempting to transmit its own packet; if carrier is present, the transmitter defers until the end of the transmission in progress before again attempting transmission. CSMA/CA adds a further random "back-off" delay before this new transmission attempt, to avoid the likelihood of a collision as multiple deferring transmitters all attempt transmission at the same time once the transmission in progress completes.

Like other wireless LANs, the AT&T WaveLAN product line includes both "access points," which AT&T calls WavePOINT units, and host network adapters, which AT&T calls WaveLAN units. The WaveLAN host network adapter is available as a Type II PC Card (PCMCIA) with a small external radio and antenna unit. The WavePOINT access point is designed to be mounted in a fixed position and connected to a wired LAN, allowing WaveLAN-equipped mobile computers to communicate with the wired LAN. The access point operates as a transparent MAC-level bridge to forward packets to and from the wired LAN as needed [14].

Each WavePOINT is assigned a unique 16-bit Network Identifier Designator (NWID), used to distinguish packets transmitted by one access point from those transmitted by another. The WaveLAN transmission format consists essentially of a standard Ethernet frame, with the NWID prepended to each frame. The receiver filters received packets in the WaveLAN host adapter based (in part) on the NWID of the access point it is using. This technique avoids the logical interference from packets transmitted in overlapping coverage areas of different WavePOINTS, although such transmissions can still physically interfere and do consume bandwidth in both coverage areas.

Like some other wireless LANs, the AT&T WaveLAN product allows "roaming," the ability for mobile computers to continue to receive service as they move from the coverage area of one access point to another. In order for this service to be provided, it is necessary that the wireless LAN system know in which access

point coverage area each mobile computer is located. Thus, the tables of the bridges contained in each access point must be updated as mobile computers move from one access point coverage area to another.

The WaveLAN roaming functionality is called WaveAROUND, and involves the addition of a beaconing and a sign-on protocol to the system [2]. Each WavePOINT periodically broadcasts a “beacon” packet, which mobile computers using WaveLAN monitor in order to detect when they have moved from the coverage area of one WavePOINT to another. When the signal strength received by a mobile computer from its current access point falls below a defined threshold, the mobile computer places its WaveLAN unit into NWID-promiscuous mode to allow it to receive the beacon signals also from other nearby access points. When a suitable new access point is identified, the WaveLAN unit is set to again filter received packets based on the NWID of the new access point, and the mobile computer uses the sign-on protocol to register with the new WavePOINT. The sign-on protocol also serves to update the forwarding tables in the old and new WavePOINTS and in any other learning MAC-level bridges on the wired network connecting the two WavePOINTS [14].

In wireless LANs, peer-to-peer (mobile-to-mobile) communication can be provided in one of two ways. In some wireless LANs, it is possible for a mobile to communicate directly with another mobile. In other wireless LANs, two mobiles, even though they are both within the coverage area of the same access point, can communicate only by having their transmissions relayed by the access point. The AT&T WavePOINT/WaveLAN products use the former approach.

In laying out a multiple access point wireless LAN installation, one must take care in locating the access points to assure that adequate radio coverage will be provided throughout the service area. There is an additional issue in a wireless LAN in which direct peer-to-peer communication between mobile computers is supported, without relaying by the access point. In this case, one must reduce the coverage area of each access point in a way that will insure that any two mobiles within the coverage area will be within radio range of each other. This reduction in coverage areas is in general done by reducing the distance between neighboring access points by half, and means that more access points will be needed to cover a given service area. Such reduction in coverage area is not necessary where peer-to-peer service is provided by relaying communication through the access point. In our installation, we have worked with AT&T to modify the WavePOINT access points such that peer-to-peer service can be provided through the access points, thus giving each access point the maximum possible coverage area.

Implementation Details

We are installing WavePOINT access points in twelve buildings on campus in order to provide complete WaveLAN coverage for all areas in these buildings and their adjacent outside spaces. These buildings constitute most of the campus classroom and administrative buildings. We estimate that approximately 200 WavePOINT units will be required to reach this level of campus coverage.

Figure 4 illustrates how the WavePOINT units are connected to our campus network. We are installing a new IEEE 802.3 backbone network on campus to connect the WavePOINT units in each building with the rest of the campus wired network through a separate router located in the main campus data communications facility. In each building, each WavePOINT is connected to a Synoptics IEEE 802.3 10BaseT hub located either in the building’s master wiring closet or in a remote wiring closet on a floor of the building. These hubs, in turn, are connected by multimode fiber cable to a 10BaseT hub which is central to the high-speed wireless network. This hub connects to a Cisco router. The WavePOINT units and hubs are powered by the campus 110 VAC system.

This structure allows us to operate the high-speed wireless network independently of the campus network and to disconnect the two if necessary. We consider this arrangement necessary at least until the wireless network is stable and considered a fully operational system. This structure also allows us to separate traffic on the wireless LAN backbone network from traffic on the rest of the campus network. The Cisco

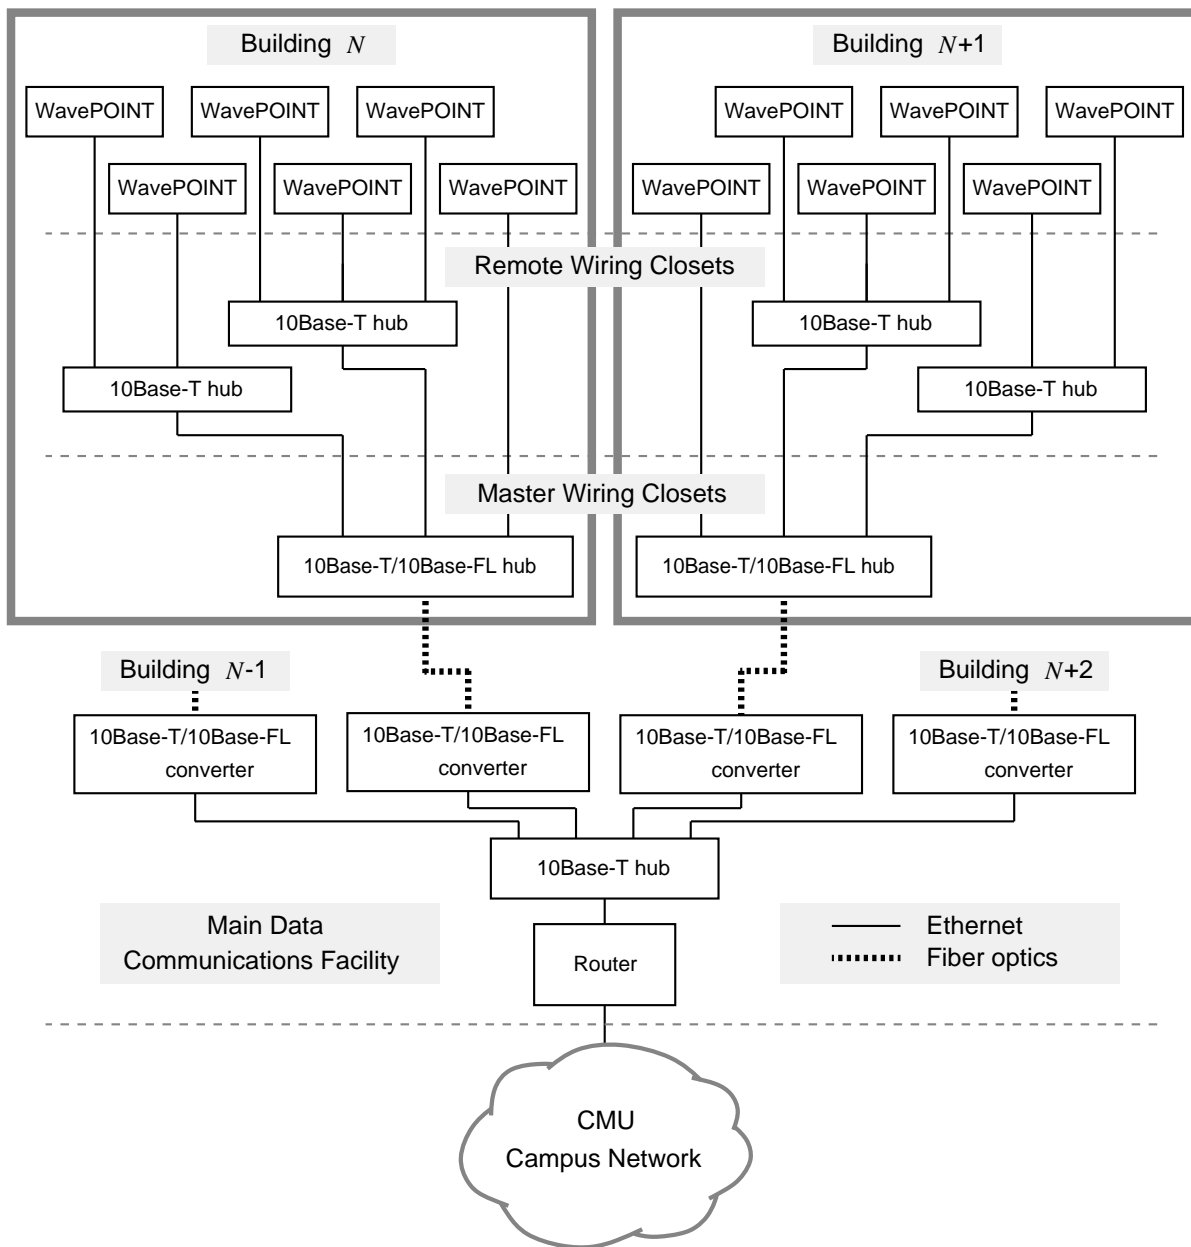


Figure 4 High-speed local area wireless infrastructure connections

router connecting the wireless backbone to the rest of campus filters packets based on destination address and only passes packets to and from the wireless backbone needed. For example, this allows us to avoid passing the packets used for the WaveAROUND sign-on protocol between a mobile computer's old and new WavePOINT access point, over the main campus network backbone.

As with CDPD, we are primarily using the FTP Software PC/TCP implementation of the TCP/IP protocols on our IBM ThinkPad laptop computers running Microsoft Windows. The WaveLAN units include an NDIS [1] device driver for interfacing to the TCP/IP protocol stack. We are also using the WaveLAN units with NetBSD Unix running on the ThinkPads and have written a Unix WaveLAN device

driver and PC Card support for NetBSD to support this. The Unix device driver supports the full WaveLAN protocol including WaveAROUND roaming.

Deployment and Experience To Date

In contrast to CDPD service, where coverage is primarily the responsibility of the cellular carrier, a user of wireless LAN equipment must take responsibility for installing access points carefully, in a way that provides complete coverage of the service area. The layout should be based on measurements, not just on “rule-of-thumb” calculations. These measurements involve extensive testing and careful consideration of radio propagation issues when the service area is large, such as when covering an entire campus as in our installation.

Because the coverage area of an access point is relatively small, terrain is not a propagation issue. Rather, the layout and construction of buildings determine the coverage area of each access point. AT&T rates the WaveLAN/WavePOINT transmission range at up to 800 feet in an open environment, but this rated range may be reduced to 100–200 feet through walls and other partitions in typical office environments. Wood, plaster, and glass introduce little barrier to the WaveLAN radio transmissions, but brick and concrete walls can be a significant barrier; the greatest barrier to radio transmissions commonly found in office environments is metal, such as desks, filing cabinets, reinforced concrete, and elevator shafts. Even a carefully considered access point design layout may have to be modified after installation is complete in order to remedy coverage gaps.

Because wireless LAN equipment is rarely used in such a large system, we are proceeding cautiously to introduce the service to the campus community to be sure that the wireless LAN service is reliable and stable before usage is attempted on a large scale. We plan to announce the service as available for research support and experimental use after we have accomplished the following tasks:

1. *Coverage Verification.* We will carry out coverage testing to be sure that targeted areas are, in fact, receiving complete coverage.
2. *Trouble Reporting and Tracking System.* We will establish a system to report and track problems on the wireless LAN system.
3. *Network Management Tools.* We will establish a rudimentary network management system for the access points, hubs, and routers associated with the wireless LAN system.
4. *Hardware and Software Tools.* We will procure and make available to the campus the hardware (e.g., host network adapters) and software they will need to support a variety of application software running on a number of specified hardware platforms.

We began installation work for the WavePOINTS and new backbone network for connecting them in May 1995, and as of November 1995, have completed installation and testing in five buildings. Our testing revealed some coverage gaps in a few buildings, which we are currently working to correct with the installation of a small number of additional WavePOINTS. The system is currently capable of supporting PC-compatible computers but does not yet support Macintosh platforms. Since Macintosh computers are so commonly used on our campus, AT&T is working with us to develop support for these as well. It is anticipated that such support will be available by summer 1996.

Related Work

A number of other universities are currently also planning or deploying wireless data network infrastructures. Purdue University is developing a campus-wide wireless network, called Crosspoint [6], based on the

Solectek AIRLAN wireless LAN product. Solectek AIRLAN uses the same hardware technology as the AT&T WaveLAN system we are using, and is built under license from AT&T. In Crosspoint, each AIRLAN unit is connected to a base station computer, and the base stations are connected using an ATM backbone. They run their own protocols between base stations to support sign-on of a mobile computer with a base station, providing roaming of mobile computers between base stations. Crosspoint uses an ATM backbone in order to support the large number of routing messages that may be generated between base stations, for example, as students move from one class to another.

At Stanford University, a new project called MosquitoNet [3] is studying operating system and application issues in mobile and wireless computing, and is building an experimental platform for their research. They are utilizing the Ricochet microcellular wireless data network service being deployed by Metricom in the San Francisco Bay Area, and planned for deployment in other major metropolitan areas in the United States. The Metricom service uses frequency hopping spread spectrum in the 915 MHz ISM band to transmit wireless data packets at a raw data rate of 100 kbps. Ricochet uses small radio units mounted on street lights, utility poles, and buildings, each with a range of about 2–5 miles. A subset of the radio units also have a wired network connection, and the Ricochet routing software forwards packets at a rate of about 2–10 hops per second between wireless radios to reach one of these wired access points. The approximate user data throughput of the system is 10–40 kbps.

The InfoPad Project at the University of California at Berkeley is building InfoNet, a wireless data network that will allow multimedia communication with small portable terminals, called “pads” [12]. The wireless link for InfoNet is currently provided by two different types of networks, each serving one direction of the communication with the pads. For communication from a base station to a pad, InfoNet uses commercial radio components manufactured by Plessy, operating at 700 kbps, and for communication from a pad to a base station, wireless LAN units manufactured by Proxim are used, operating at 244 kbps. The backbone network connecting the base stations is currently Ethernet, although an ATM backbone is planned.

Also at the University of California at Berkeley, the BARWAN Project (Bay Area Research Wireless Access Network) is developing a wide-area wireless data service featuring a number of different wireless networks each operating over different scales, including regional-area, wide-area, metropolitan-area, campus-area, in-building, and in-room. The BARWAN Project seeks to integrate these different networks together seamlessly, sharing much the same vision as our Wireless Andrew infrastructure and some of the research being conducted on it [11]. Like the MosquitoNet Project at Stanford University, BARWAN is using the Metricom Ricochet microcellular data network service. They are also integrating use of the Hughes Direct Broadcast Satellite system into the network. The Hughes “DirecPC” service provides downlink facilities only (wireless data sent to a mobile computer), but operates at a raw data rate of 12 Mbps. The DirecPC system thus must be used together with a separate (slower) network link for the uplink traffic (from the mobile computer).

Conclusion

The Wireless Andrew wireless data network infrastructure which has been described is expected to be invaluable in its ability to support research in mobile computing and wireless networking at Carnegie Mellon University. Some of the research projects that are already using this infrastructure include the following:

1. Automatic connection to the most appropriate network and handoff to another network when appropriate [11]. The appropriateness of a network depends on its coverage area, speed, error rate, and usage cost in relation to the needs of a particular application.

2. Scalable and transparent support for IP routing in an environment in which hosts are mobile. This is the “Mobile IP” problem [11]. CDPD and WaveLAN each handle host mobility internally, but IP address assignment differs between these two networks, and we are using Mobile IP to switch routing between these networks as well as between base stations on other types of networks that do not support their own protocols for internal host mobility.
3. Mobile access to information, including the caching and consistency of files in a distributed file system environment [19]. These are challenging issues with wireless links because of their unreliable characteristics relative to wired links.
4. Image and video compression techniques which can adapt to the bandwidth of the wireless link, including when switching between networks with different connection qualities [13]. For example, when switching from a 2 Mbps wireless LAN link to a 19.2 kbps CDPD link, the compression algorithms should adapt to the change in available network bandwidth.

This infrastructure will also enable faculty, staff, and students at Carnegie Mellon to access the facilities of the campus network and the Internet from around campus and around the greater Pittsburgh area. This will also serve as a source of empirical data on how mobile computers and wireless data network service are actually used by people.

Acknowledgments

We wish to thank John Leong, Richard Hovey, Robert Baron, Charles Bartel, and Pete Bronder for their contributions to this work. The work reported here is supported by Carnegie Mellon’s Information Networking Institute, the National Science Foundation, Bell Atlantic NYNEX Mobile, and Bell Communications Research.

References

- [1] 3Com Corporation and Microsoft Corporation. Microsoft/3Com LAN Manager Network Driver Interface Specification. Version 2.0.1, October 1990.
- [2] AT&T Wireless Communications and Networking Division. Data manual: WaveLAN air interface. Document Number 407-0024785, Revision 1 (draft), June 1995.
- [3] Mary G. Baker. Changing communication environments in MosquitoNet. In *Proceedings of the Workshop on Mobile Computing Systems and Applications*, pages 64–68, December 1994.
- [4] David F. Bantz and Frédéric J. Bauchot. Wireless LAN design alternatives. *IEEE Network*, 8(2):43–53, March/April 1994.
- [5] CDPD Consortium. Cellular Digital Packet Data system specification. Release 1.0, July 1993.
- [6] Douglas E. Comer, John C. Lin, and Vincent F. Russo. An architecture for a campus-scale wireless mobile internet. Technical Report CSD-TR 95-058, Department of Computer Science, Purdue University, 1995.
- [7] Antonio DeSimone and Sanjiv Nanda. Wireless data: Systems, standards, applications. Technical report, AT&T Bell Laboratories, December 1994. To appear in ACM/Baltzer *Wireless Networks*.
- [8] Martin Hall, Mark Towfig, Geoff Arnold, David Treadwell, and Harry Sanders. Windows Sockets: An open interface for network programming under Microsoft Windows. Version 1.1, January 1993.
- [9] IEEE. Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications. Draft Standard 802.11 D2.0, July 1995.
- [10] Van Jacobson. Compressing TCP/IP headers for low-speed serial links. Internet Request For Comments RFC 1144, February 1990.
- [11] David B. Johnson and David A. Maltz. Protocols for adaptive wireless and mobile networking. *IEEE Personal Communications*, 3(1), February 1996.

- [12] My T. Le, Fred Burghardt, Srinivasan Seshan, and Jan Rabaey. InfoNet: The networking infrastructure for InfoPad. In *Digest of Papers: COMPCON Spring 1995*, March 1995.
- [13] José M. F. Moura, Hirohisa Shiojiri, Jyh-Cherng Lin, and Radu S. Jasinchi. Video over wireless. *IEEE Personal Communications*, 3(1), February 1996.
- [14] Radia Perlman. *Interconnections: Bridges and Routers*. Addison-Wesley, Reading, Massachusetts, 1992.
- [15] J. B. Postel. User Datagram Protocol. Internet Request For Comments RFC 768, August 1980.
- [16] J. B. Postel, editor. Internet Protocol. Internet Request For Comments RFC 791, September 1981.
- [17] J. B. Postel, editor. Transmission Control Protocol. Internet Request For Comments RFC 793, September 1981.
- [18] J. Romkey. A nonstandard for transmission of IP datagrams over serial lines: SLIP. Internet Request For Comments RFC 1055, June 1988.
- [19] Mahadev Satyanarayanan. Mobile access to information. *IEEE Personal Communications*, 3(1), February 1996.
- [20] Bruce Tuch. Development of WaveLAN, an ISM band wireless LAN. *AT&T Technical Journal*, 72(4):27–37, July/August 1993.

Biographies

ALEX HILLS is Vice Provost for Computing Services and Chief Information Officer at Carnegie Mellon University, where he is also a Distinguished Service Professor of Engineering and Public Policy. He is responsible for the University's academic computing, telecommunications, and data network facilities, and is also involved in wireless communications research activities. Professor Hills joined Carnegie Mellon after serving as Executive Director of the University of Alaska Computer Network. He was extensively involved in developing the telecommunications networks in rural Alaska and was Deputy Commissioner of the Alaska Department of Administration, the state's chief telecommunications official. He holds a B.S. in electrical engineering from Rensselaer Polytechnic Institute, an M.S. in electrical engineering from Arizona State University, and a Ph.D. in engineering and public policy from Carnegie Mellon University. He is a Senior Member of the IEEE.

DAVID B. JOHNSON is an Assistant Professor in the School of Computer Science at Carnegie Mellon University, and also holds a courtesy appointment as an Assistant Professor in the Electrical and Computer Engineering Department. His research interests include network protocols, distributed systems, and operating systems. He has worked actively within the Mobile IP Working Group of the Internet Engineering Task Force (IETF) for the past three years and is one of the principal designers of the current IETF Mobile IP protocol. Prior to joining the faculty at Carnegie Mellon in 1992, he was on the faculty at Rice University for three years as a Research Scientist and Lecturer in the Computer Science Department. He holds a B.A. in computer science and mathematical sciences, an M.S. in computer science, and a Ph.D. in computer science, all from Rice University. He is a member of the IEEE Computer Society, IEEE Communications Society, ACM, USENIX, Sigma Xi, and the Internet Society.