Wireless Andrew: Experience Building a High Speed, Campus-Wide Wireless Data Network

Bernard J. Bennington
Information Networking Institute

Charles R. Bartel
Computer Services

Carnegie Mellon University, Pittsburgh PA 15213-3890
(benb, crb@andrew.cmu.edu)

Abstract

Carnegie Mellon University (CMU) has created a high speed wireless network, known as Wireless Andrew. This network uses wireless local area network technology, utilizing spread spectrum in the ISM band. The purpose is not only to support wireless research but also to create a campus-wide mobile computing laboratory.

The paper shows the challenge of designing and managing large scale wireless networks. The examples show the differences experienced in wireless vs. wireline networks due to the nature of mobility vs. static components, the nature of RF propagation vs. wireline connection, and the difference in sophistication of tools because of the different place in the product development cycle. The paper describes the lessons learned at completion of the first three year phase including: installation design issues, issues of RF interference and data throughput, the unique problems of wireless network management, and the release of the system.

1. Introduction

In 1994 the Information Networking Institute (INI) of Carnegie Mellon University (CMU) successfully submitted a proposal to the National Science Foundation to create a "High Speed Wireless Infrastructure." The purpose was to support wireless research [1,2,3,4] but also to attract additional research and create, in essence, a campus-wide user community as a mobile computing laboratory. The CMU campus is ideal for such research, as,

- it is concentrated,
- even non technical departments are heavy users of networked computing, and
- most staff, students, and faculty live within a one-mile radius of campus, making the surrounding area ideal for prototyping extensions to the campus systems.

The campus wired network, known as the Andrew system [5,6] consists of three major components: a high speed campus-wide network, network clients, and network services. The key network services are the Andrew File System, electronic mail, bulletin boards, distributed printing, access to library information, access to the supercomputing center, and the Internet.

The Wireless Andrew concept was to build a high speed wireless access system using wireless local area network technology, utilizing spread spectrum in the ISM (industrial, scientific, and medical) band. This access system is integrated with wireline Andrew. The whole is to be an operational utility for the use of anyone on campus with an appropriate computer and wireless device. This implies every conceivable usage from researchers with wearable computers, to mobile robots, to faculty with laptops in meetings, to students in classrooms. All should be able to draw upon their normal Andrew network and Internet services, albeit at a partly reduced data rate from that of wired access.

Not only was Wireless Andrew seen as a tool to support research but the very creation of Wireless Andrew in itself was seen as a research project. The development creates an installation larger by some orders of magnitude than networks normally built with wireless LAN technology. Indeed, to our knowledge, the Wireless Andrew network is the largest single WaveLAN system installed anywhere.

In this paper we present specific implementation experience gained over three years and show how the experience was used to improve tools and methods in designing and implementing large systems. We begin by looking at the process for selecting a product, how the original design was created, and the make up of our early adopter community. The paper goes on to examine issues like interference and throughput, changes in the design method, and network management. Finally we look at the release of the network and the measurement systems we have put in place.

2. The Project

2.1 Selection of Partner

From the beginning, we understood we were dealing with a technology and concept early in its growth cycle. We expected that specific wireless technology products would become obsolete quickly, and the architecture conceived today could be primitive within a few years. Consequently, it was our strategy to partner with a technology vendor based not just on current

MOBICOM 97 Budapest Hungary

product but on their view of the future. We were looking for a partner who saw CMU's campus representing the challenge of future markets and hence a testbed for ideas and solutions. We anticipated there would be issues of scalability of products and systems. These were not just the obvious technological ones, such as throughput and ability to roam, but included two key parameters needed of an operational network:

- Ability to design an effective network: a network that would cover such varied characteristics as lecture rooms, labs and offices; historic and new buildings; multi-story buildings; a network surrounded by a busy urban area; and serving every conceivable usage, as stated, from mobile robots to classrooms;
- Ability to manage the wireless network and provide operational service levels acceptable to wireline Andrew users.

In September 1994 a team composed of the INI (as the research organization) and the university's Computing Services group (as the implementers and future operators of the system) began by looking at published comparisons of wireless LAN product performance [7], then held informal discussions with all vendors who claimed to have products. We solicited products for testing and received them from Lucent Technologies (then AT&T), Proxim, and Xircom (now Netwave Technologies Inc.) The products tested are summarized in Table 1.

	Xircom	Proxim	Lucent
Frequency	2.4 GHz	2.4 GHz	900 MHz
Transmission	FH	FH	DS
Claimed data rate	1 Mb/s	1.6 Mb/s	2Mb/s

Table 1: Tested Products

To gain familiarity with the technology we undertook "ping" tests and FTP of large files. For example, with a single laptop, then with two and three laptops simultaneously, we executed tests such as 1000 continuous pings or FTPed a 1.5 Mbyte file. Typical results are shown in Table 2.

	Xircom	Proxim	Lucent
Latency/Failed pings.			
1 Laptop	38 ms/2	32 ms/63	27 ms/0
2 Laptops	41 ms/1	35 ms/377	28 ms/0
3 Laptops	-	43 ms/645	28 ms/1
Data rate bits/sec.			
1 Laptop	233,453	311,292	360,951
2 Laptops	135,591	202,476	347,853
3 Laptops	-	141,268	329,804

Table 2: Typical Performance Test[1]

Coverage testing focused mainly on the School of Computer Science building (Wean Hall, an eight story major campus building) but also included the INI in Hamburg Hall and the Computing Services building, Cyert Hall. As we gained familiarity with the technology, we made a preliminary technical assessment based on factors such as form factor, installation characteristics, management tools, latency, throughput, and coverage, and found Lucent and Xircom to be essentially equal and, though we rated Proxim lower, all three vendors were considered qualified and so were included in the continuing discussions.

As a result of the testing and discussions, we began to quantify what we felt were the key technical criteria: 1. coverage, 2. throughput, 3. form factor, 4. ease of use, and, 5. Apple Macintosh support as well as PC support. The last item was important because of the considerable number of Apple laptops existing on campus in addition to PC laptops.

Central to these technical criteria were two issues: Direct Sequence Spread Spectrum vs. Frequency Hopping Spread Spectrum; and 900 MHz vs. 2.4 GHz, the two available ISM bands in the U.S. Regarding the former issue, there were identifiable trade-offs between cost per unit throughput and resistance to different kinds of interference. In the second, there were issues of coverage per unit cost vs. specific known interference, future products, and industry direction. As part of the coverage assessment, for example, we compared the products using ping tests in Cyert Hall, which has sparse internal structure and few obvious obstacles such as elevator shafts or bulk filing cabinets, to determine typical linear coverage. We then took the proposed cost of access points and, assuming circular coverage with the linear distance as radius, calculated the cost per unit coverage and throughput. As Lucent had indicated that their future direction would be away from the 900 MHz product (a band useable only in the U.S.) to a global version of a 2.4 GHz product, we estimated the coverage/cost factors for the as-yet-unreleased product. The comparative ratios were a coverage cost per unit of 1.0 for 900 MHz, 1.67 for 2.4 GHz, and 3.76 for Xircom's product, with the discrepancy ratio increasing to almost 1:6 when throughput was also considered.

As our technology ideas began to focus, we began to explore with each of the vendors what we thought important from a partnership perspective: 1. resources, expertise, and commitment; 2. a strategy for ensuring the continued currency of the network; and, 3. willingness to enter into a special relationship to jointly pursue research and development. At this point we asked each of the three tested vendors to submit a formal proposal addressing technical, cost and partnership issues. Based on these proposals we excluded Proxim. As we considered the whole package of existing products and proposed improvements of the two other companies, there was not a clear choice. For example, Xircom proposed to allow external antennas, which would improve coverage but degrade form factor. Lucent's proposal included the planned 2.4 GHz product, which had the potential to improve aggregate throughput but would reduce coverage. In addition, Xircom was also taking steps to improve its product's throughput. If all of these improvements would come to pass, the two products would be judged more nearly equal. The selection began to pivot on the proposed partnerships. Here, there was also the

[1] The missing results were due to inability to get one platform working with the Xircom product. The low throughput compared to advertised was due to the FTP application and operating system overheads.

need to balance what could actually be delivered vs. what might or might not materialize. This was particularly true of a very large corporation such as Lucent whose span encompassed almost everything in the wireless industry. Deliberate care was taken to ensure real commitment that could actually be delivered by the proposing team. However, we did not foresee in any way the break up of AT&T into three separate companies and the impact it would have upon us, as the AT&T proposal included what were to become two separate corporations, Lucent and NCR.

In the final analysis, using coverage, throughput, form factor, ease of use and Mac support for the final technical scores, and implementation support, continued currency, and commitment for future research collaboration for the final in partnership scores, the results were very close: Lucent 900 MHz = 81%; Lucent 2.4 GHz = 75%; and Xircom = 75%. There was not a clear leading competitor. Both Lucent and Xircom provided a believable team and commitment. However, in considering the potential for future research collaboration, the arrangement and commitment proposed by Lucent clearly could bring more added value to CMU. In addition, the balance was swung when the School of Computer Science Coda Project [2] indicated they had decided to go ahead with the WaveLAN product. There was one final consideration of perhaps having two test beds: one Lucent, one Xircom, but this was rejected as impractical for the general user.

As a consequence, a final decision was made in March 1995 that, though each vendor had plusses, for our purposes the 900 MHz WaveLAN product from Lucent Technologies of Utrecht, Holland, [8] was considered the best overall solution for CMU.

The selected technology, WaveLAN, is composed of two main elements: a network interface card (WaveLAN card) using direct sequence spread spectrum physical layer and CSMA/CA medium access control; and a WavePOINT access point acting as an Ethernet bridge and handling roaming from cell to cell. The raw data throughput in a single cell is two megabits per second.

2.2 Installing the Buildings

As part of the final decision, a pilot installation was built covering the sixth, seventh and eight floors of Wean Hall (the Computer Science building). Wean Hall is a reinforced concrete structure and we had had concerns about penetration and coverage. The pilot was expected to test as many of the assumptions concerning coverage and operation as possible. In general, the pilot indicated that penetration between floors was higher than we had assumed and coverage was significantly better than our worst-case scenario used in initially estimating access point placement.

Following the success of the pilot, coverage testing was accomplished by AT&T Global Information Systems (now NCR), as the U.S. presence of the WaveLAN organization. This was completed in April 1995 for the first six buildings. A design was created, and installation completed by September for six buildings representing about 50 percent of the campus teaching, research and office space:

- Cyert Hall: home of Computing Services ;
- Hamburg Hall: home of the INI and the Engineering Design Research Center;
- Wean Hall: home of the School of Computer Science;
- Porter Hall: home of researchers from Electrical and Computer Engineering and Social and Decision Sciences, who had wireless work in progress;
- GSIA/Posner Hall: the twin-building home of the graduate school of business, the largest base of student laptops.

The total access points (or cells) used in this initial installation numbered 73. At that time we were considering six more buildings: Baker Hall, Doherty Hall, Hammerschlag Hall, Scaife Hall, the Hunt Library, and the Fine Arts building -- in essence almost all of the campus teaching, laboratory and office space. This was estimated at another 100 access points. This size can be contrasted with a typical wireless LAN application in the business world of two or three access points in a retail store.

The access points are interconnected via a dedicated network using SNMP (Simple Network Management Protocol [9]) managed 10BaseT hubs. These hubs are connected with fiber to a wireless backbone network located in our central network facility. A router then connects the wireless backbone network and the campus backbone. This method provides a dedicated wireless subnet routed into the Andrew system (see Figure 1.)

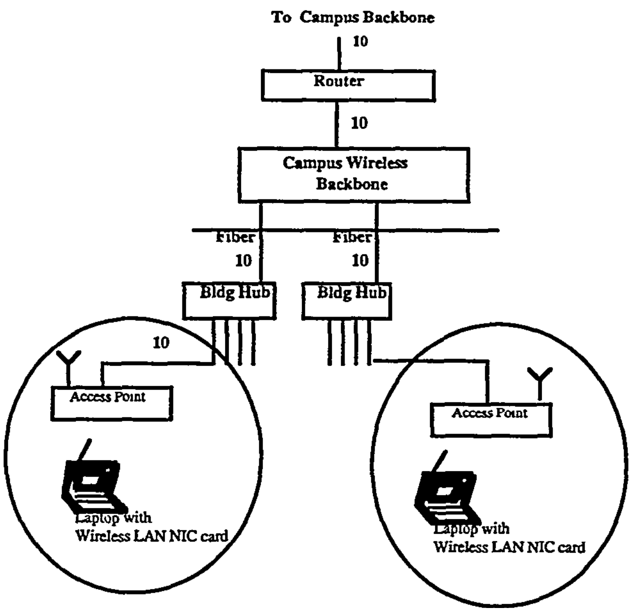

Figure 1: The Wireless Andrew Subnet

Installation of every wireless access point involves design and measurement of the RF propagation pattern to properly locate the individual access point, then provision of Ethernet and power wiring to the selected location. This physical infrastructure is itself a large investment, constituting about one quarter of the overall costs. The wish to preserve this investment in physical wiring during future upgrades of the network is central to one of the arguments detailed below in Section 3.3.

By the end of 1995, the network of the first six buildings was in place [10]. The plan was to exercise this network with a selected, controlled population of users to identify two things:

- issues that had to be resolved before turning the network over to public operation, and
- issues concerning further expansion to the whole campus.

The former concerned our ability to meet expectations of operational reliability established by the wired network. The latter considered inevitable changes in product direction and strategy over time in such a fast-changing technology.

2.3 Early Adopters

An initial "early adopter" user group of 80 users was identified, evenly distributed across the six buildings. It was important to us that each of these had:

- purchased their laptops already for other purposes and so had no unrealistic expectations of early return-on-investment in Wireless Andrew, and
- were sophisticated users who would work with us to identify and help solve problems.

Our intention was to loan each of them the necessary WaveLAN card to reinforce the first point above, i.e. no personal money invested in the early system. This was planned to include seeding 40 WaveLAN cards and to purchase an equal number of Apple adapters which had been developed by Digital Ocean. In fact, though our target initial community was 80, a number of research groups went ahead and purchased WaveLAN cards themselves to support their research efforts. This created a total pool of nearly 100 early adopters.

During this first phase, we deliberately tried to keep the existence of the network quiet, though some publicity about the project being underway was inevitable. However, the policy was to not raise expectations until we were ready to release the first buildings as an operational network. Then and only then would the campus retail store stock WaveLAN and Digital Ocean devices, and hence by implication begin to attract users who would make a significant investment in attaching to Wireless Andrew.

3. Findings

By September 1996, the anniversary of the first installation of the initial six buildings, we had solved several issues, some anticipated and some not. Four of these we considered significant enough, that having solved them jointly with our partner, we made the decision to completely redesign and reinstall the system, moving almost all access points. These major issues included: 1. interference, 2. throughput, 3. design method, and 4. network management. This section discusses these issues, and in Section 4 the resulting redesign and reinstallation is discussed.

3.1 Interference

The ISM band is a shared band and hence there is the potential for different users to interfere with each other. This comes from two general sources:

- Interference by a strong foreign signal. We had expected this source;
- Contention for usage of the same spectrum by overlapping cells in the same subnet. We had not expected this to the extent that we experienced it.

First, regarding foreign signal interference: The 902 to 928 MHz ISM band accommodates many applications including wireless stereo speakers, industrial heaters, welding equipment, food preparation equipment, medical instruments such as magnetic resonance imagers and diathermy machines (perhaps significant due to CMU's contiguous location to the massive hospital complexes of the University of Pittsburgh), as well as military radar, law enforcement video surveillance cameras, and commercial location and monitoring services [11]. A general scan was made of the 902-928 band using a spectrum analyzer and nothing significant was noted at that time. One potential growing source of foreign signal interference was the use of 900 MHz cordless phones. To test the possible impact of these, we tested throughput of a small number of laptops in the region of a single access point for signs of excessive retransmission of packets while a 900 MHz cordless phone was in use. We looked for consequent degradation of throughput but none was observed. On the contrary, the effects of interaction only seemed to impact the signal quality of the telephone. This observation has recently been confirmed for narrowband 900 MHz cordless phones by more extensive tests, although some interference was observed from the newer direct sequence spread spectrum 900 MHz phones [12].

Another example of interference from a foreign signal was encountered from a nearby commercial paging transmitter sited near the University. There are two sets of signals both associated with paging. One is occasionally seen at 928 MHz and one is at 931 MHz. Although technically outside the 902-928 MHz ISM band, the power of the 928 MHz paging antenna is sufficient at times to swamp the top end of the ISM band down as far as 926.25 MHz, with signal strength of the same order as that experienced by some WaveLAN users. This causes interference with access points at the far western walls of buildings. The most obvious approach to resolving this would be to add additional access points in the affected areas to overcome the interfering signal. This however contributes to the second facet of interference, throughput, which is addressed in Section 3.2.

There was one concern at the lower end of the ISM band. During implementation of the wireless network, Bell Atlantic NYNEX Mobile (BANMS) approached CMU about locating a new cell site on Wean Hall. There were concerns that the large signal strength could spill into the lower end of the band and impact Wean Hall users. We decided to undertake formal rigorous tests prior to allowing the cell site. Lucent indicated that if possible the tests should reflect a -40 dBm signal strength from the BANMS transmitter measured at the

WaveLAN access point and mobile unit. (The BANMS designed value was -50 dBm.) The chosen frequencies representing the upper cellular band were 892.99 and 893.97 MHz [13]. Testing was done at both peak (4:30 to 6:30 p.m.) and off peak calling (5:30 a.m. to 6:30 a.m.) times. The tests consisted of using the WaveLAN point-to-point tool (ptpdiag) to determine signal level, packet loss, signal quality, and signal to noise ratio, and the WaveMonitor testing tool (wmonitor) to log signal strength and local and remote noise levels. In addition, ping tests were used to examine actual throughput. The tests were accomplished by locating hot spots in Wean Hall using a spectrum analyzer to determine areas where there were BANMS signals of -40dBm coinciding with WaveLAN signals of approximately -80 dBm to represent an assumed worst case. No significant difference in WaveLAN performance was noted so the cell site was allowed.

During the cell site tests, however, it was noted that other unknown sources of interference were in fact stronger than the BANMS signals. This confirmed random experiences of a number of users that there are intermittent sources of significant interference on campus. For example in Hamburg Hall, pulsating interference is seen within the ISM band, manifested in noise readings seen by the WaveLAN tools and erratic performance. However, although spectrum analyzer readings have been taken at various sample intervals over times of several hours, this source of interference has so far proved elusive. Work is continuing on this.

There are still other sources of interference localized to certain rooms which have not yet been investigated. These also will be examined in the coming year, interference analysis representing a fruitful area for further work.

3.2 Throughput

Although throughput had been measured in a simple comparative manner prior to implementation, this was more thoroughly investigated as installation proceeded. The concern was the performance in a classroom situation where a large number of users, concentrating on the same application, could all hit the enter key demanding simultaneous service. Multiple copies of a 25 Mbyte test file were created on a Hewlett Packard server, and this was networked to be accessible from a single WaveLAN access point, isolated in both network and RF terms from any other access point in the network. Thirty laptop computers were configured for both wired Ethernet and wireless WaveLAN access. These included eight Mac Powerbook 5300cs, five TI TravelMate 5000, six DEC PC 425SL, seven IBM Thinkpads including 750C, 701C and 750CS, three Compaq Aero, and a NEC Versa. Operating systems included Mac 7.5.3, DOS 6.22, Windows 95, NetBSD Unix, and Mach 2.6. This formed a fairly representative set of what might be found in a classroom. A total of 12 tests were run of simultaneous FTP transfers for 1, 5, 10, 15, 20, and 29 PCs[2]. This was done first over wired Ethernet then wireless WaveLAN. The result is shown in Figure 2. For the WaveLAN network this shows, in essence, an effective 1.7 Mbit/sec link shared among the contending applications.

Although the performance is not that of wired Ethernet, it meets our expectations that there be a data throughput of 2 Mbit per second reduced by overheads but shared reasonably equally among contending stations. This, however, becomes a problem where there is overlap of cell coverage. The nature of the CSMA/CA protocol is that an access point will defer if it detects a nearby transmission in the spectrum -- irrespective of whether it is within the same cell or not. This includes both data traffic being transmitted by the other access point and also the regular beacons from that access point. Substantial overlap of cell coverage designed to fix the pager interference problem referred to in Section 3.1 would result in the possibility of reduction of throughput, essentially placing more stations contending for single capacity. This is a tricky issue to wrestle with. It is one of the factors that contributes to the complexity of the design methodology discussed below.

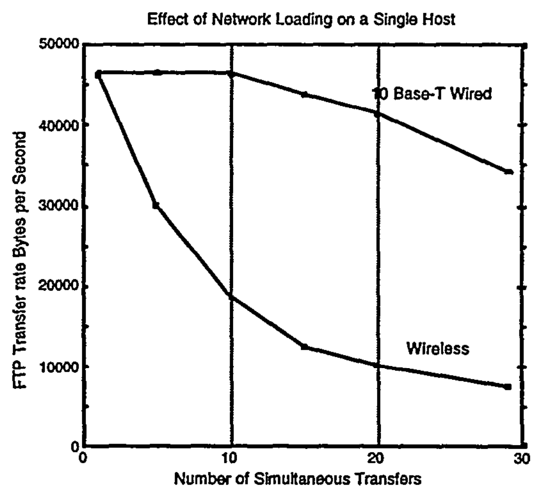

Figure 2: Effect of Network Loading on a Single Host

3.3 Design Issues

Crucial to establishing a network with good performance and economics is the design methodology and the ability to judge trade offs. Three issues encountered are:

- Holes in coverage,
- Marginal coverage areas, and
- Positioning for future upgrading from 900 MHz to 2.4 GHz.

Regarding holes in coverage: the optimum layout of cells would ensure minimal overlap to give maximum coverage per unit cost. However, there must be sufficient overlap to allow the roaming algorithm time to begin to search for another cell, fix on it and transfer registration, as a user moves from cell to cell. This needed overlap is also a function of speed of the

[2] The 30th PC failed several times and its results were discarded.

mobile user and the time they take to transgress the boundaries of the cell. This can include the obvious case of a human wandering from location to location with a laptop, but on our campus could also be a walker with a wearable computer, a mobile robot, or even a moving vehicle. To overlap the cells unnecessarily is inefficient with regard to throughput, as discussed above. The complexity of space-shapes in large multistory buildings, the inevitable overlap between floors as propagation is in three dimensions, and the variability of penetration due to different materials (e.g., office partitions vs. elevators) makes design something of an art form. Tools and procedures suitable for the current typical market of a small open retail area do not scale to large complex buildings. Propagation models prove ineffective. This complexity unfortunately was something neither we or our partner understood prior to the initial design. The only previous experience had been with small groups of cells in relatively simple, mainly two dimensional, spaces. Also, these typically did not have the demands of our users, i.e., large numbers of users in a classroom and large file transfers competing for bandwidth.

The challenge of scaling up to our environment had several facets to it:

- some tools were not designed for continuous use i.e. hour-by-hour logging large amounts of data;
- human factors that were appropriate for small efforts (like having to restart between measurement tools) were onerous in long continuous design sessions, leading to skimping the readings or making errors;
- standard signal strength values, necessary to start the roaming algorithm, were initially selected too low to ensure continuity of connection in some cases, albeit providing apparent additional coverage per cell. This resulted in unexpected disconnects;
- there was a general tendency on the part of the original design team to fix coverage problems by adding a cell, understandable in a simple environment but disastrous in a complicated building; and finally,
- the design of a network this size turned out to be more complex than expected. Some areas had virtually no coverage, prohibiting connection; some areas had multiple coverage, with the reduction of throughput mentioned above.

However, on the positive side, as we faced each problem, we learned to confidently design in almost any environment. We made changes in strategy, process, and in design tools and aids. Because experience with large installations did not exist at the start of the project in 1995, the strategy was very basic: assume that alternate floors would not need access points as they would be filled from above and below, and assume coverage per access point is fairly constant. The process at that time was to examine the drawings of the building, estimate coverage, place the access points, make spot measurements of actual signal to noise ratio, and then fill coverage holes with additional access points. The tools at that time were DOS-based. Separate tools, such as the overall monitor of all signals at a given point (wmonitor) and the tool used to diagnose signal, noise, and packet loss in a link (ptpdiag) had to be restarted between usage. Tools were calibrated in coarse intervals as a way of overcoming signal fluctuation, and only spot readings were taken, as opposed to logging data over time to determine the mean value.

We changed the strategy to one of designing each floor separately, staggering access points on alternate floors to avoid undue overlap, and iteratively moving access points to avoid having undue duplication of coverage. The process become one of undertaking three planned successive iterations. First, after estimating a coverage plan from the drawings, a single access point was taken to each of the proposed locations and sufficient readings taken to learn that specific environment's propagation and attenuation patterns, e.g. in the presence of elevators, banks of filing cabinets, sources of noise, or other peculiarities. Based on this, a first design was made for each floor and then groups of floors were tested with all access points temporarily in place. The tools were modified to allow finer readings and a Windows version gave the ability to click across different tools to move easily from broad to narrow pictures of signal strength and link quality. The new tools also allowed logging of data, marking of location, and automatic processing of data into Excel spreadsheets for later display prior to the final design iteration. With these more human-friendly tools it was easier to more accurately collect, store, and process the huge amounts of data necessary for a large building such as Wean Hall. Before assessing the final iteration, color maps were drawn of overlapping cells, to begin to visualize the performance of the whole building and see where minor changes in location of access point were necessary. Usually this third round produced a perfect building design.

A second complicating problem experienced in the design process was that the designer sees fluctuating field strengths, yet there is a need for precision in the balance between coverage, throughput, and cost. Every human being is effectively a six-foot column of water, and taking design readings to establish the correct threshold at the edge of cells while people walk by can make precision impossible. To do the readings at night is meaningless as the people will be there during operational hours. This is made more complex in areas where there is marginal field strength due to shape of the space, separation walls, or sources of interference. Once again this is not critical in small installations but becomes so in large three-dimensional spaces. To overcome this, we logged data for an extended period of time and used the mean value.

Finally, regarding use of 900 MHz vs. 2.4 GHz. All things being equal, 900 MHz provides substantially more coverage (of the order of twice as much linear coverage) and with the largest potential interference source, cordless telephones, proving no problem in our tests, 900 MHz was very attractive for the early deployment of the network in CMU. However, from the beginning, Lucent made it clear that their development path was towards a common IEEE 802.11 standard product across the Asian, European, and United States regions. Consequently, we anticipated from the beginning that at some point downstream, moving to 2.4 GHz or higher will be an attractive option with future improvements in coverage, cost, and features. As was discussed earlier, 25 percent of the cost of installation can be attributed to installing the necessary data and power wiring. Consequently, the ability to reuse the physical infrastructure is important. Siting access points today to ensure reusability of the wiring implies the ability to predict RF coverage of future devices. Erring too conservatively can cause cell overlap with the consequent throughput problem described earlier. Not being conservative enough can mean unnecessarily having to add cells in the future if the reach of the new device is only fractionally short of the old ones.

Lucent informed us that their strategy would be to ensure as far as possible that future 2.4 GHz products would overlay the coverage patterns of current 900 MHz products by suitable improvement of circuitry and changes in emitted power, but even so, it would be naive to feel that the propagation patterns would overlay exactly. As a consequence, accommodation of future migrations took subordinate place to other factors.

3.3 Network Management Issues

Management of wireless networks is not a problem with a small number of access points but, as the wireless network is scaled up to that of a large campus, the challenges found in managing any large network become relevant. In the Andrew wireline network, as with most large networks, the goals of current network management efforts are twofold :

- know that a problem exists in the network before our customers do; and
- solve and repair problems centrally rather than having to dispatch a technician.

The level of service that this creates for our end-users presents de facto service goals of the new wireless access network. However, there are several issues that challenge this:

- Lack of tools suitable for management of a large dispersed wireless network;
- The dispersed nature of devices compared to wireline units;
- Difficulty in diagnosing problems in the link between the access point and the end user;
- The mobile nature of possible problem sources; and,
- Peculiar problems, such as routing anomalies caused by the overlap of cells.

Because of the length of time that large wireline networks have been with us, there are now sophisticated network management tools available at both the network and individual component level. For example, in the wireline Andrew system, the concept of an inverted backbone architecture (built as a star topology from hubs at the central Computer Services building) has been used extensively. This approach allows all bridges, routers and servers to be monitored and isolated for diagnosis from the central site.

Andrew uses a number of sophisticated network management tools [14]: some home grown such as domain name server (DNS) utilities and SNMPCON (SNMP Console), a tool which monitors most of the critical devices connected to the campus network and immediately notifies Operations of any problems with network connectivity. The manufacturers of network components, such as Cisco, have developed over the years extended implementations of SNMP to allow very detailed diagnosis and testing of network devices. Wireless devices are very early in their development cycle and, as we have said, typical networks up to now are very simple, rarely more than a few access points. Hence the current implementations of SNMP are primitive for these devices.

How this is different is seen, for example, in the wireline network being remotely checked with DNSMON and DNSTEST (the domain name server monitor and testing tools) to see if any server is the right machine, is working, and what process it is running. All of this can be seen centrally. It also means being able to use SNMPMON (the SNMP monitors) in the wireline network to even see the temperature, voltage, use of memory, and routing tables of wireline devices. Ultimately, in the case of our hub management system, we are able to shut down and repower wireline network items or isolate a piece of equipment centrally. This is not possible currently with the equivalent wireless devices. The only alternative is sending out a technician. The technician's role is further complicated by the fact that access points are fairly uniformly distributed throughout buildings, whereas wireline components can be deliberately organized and isolated to a few telecommunications closets. In Wean Hall, for example, this is the difference between visiting a maximum of seven wiring closets for *all* wireline devices vs. visiting 32 access points just for the wireless network, and some of these may be located in inaccessible areas behind walls or above ceilings for aesthetic purposes.

There are other differences in managing the wireless network due to the nature of the RF link between the access point and the end machine differing from that of a wireline connection. In the wireline network, pinging to the bridge and then to an end machine, known to be operational, can help to isolate the pinpointing of any trouble to the backbone or to the access link. Without permanently locating fixed wireless machines in each cell to monitor the RF connection (a relatively costly proposition) there is no permanent end point to ping to, making it difficult to isolate failure points in the wireless network. In fact, the RF link may be down and this will not be discovered until a user, failing to establish connection, reports a fault.

Another example of where mobility causes a difference occurs when an end-user machine is, in fact, the source of the network problem. In a wireline network, it can be isolated by the hub to remove the source of the problem. In the case of a mobile platform, the option is to isolate all connections in a given cell containing the faulty laptop by shutting off the access point. However, this does not stop the faulty laptop from roaming into the next cell, making the isolation meaningless.

Finally, there are routing anomalies due to the fact that RF coverage of different cells can overlap. This has no topological equivalent in wireline other than a "short circuit". Routing anomalies can take place, for example, as shown in Figure 3, where a laptop is connected to access point B, but the coverage of access points A and B overlap. Although access points A and B have different domain IDs, and hence are different networks; and although the correct routing path between Net 1 and Net 2 is through the router shown, the laptop also being in the signal strength area of B will promiscuously listen and see the beacons from both A and B and also see ARP requests on both networks. This in turn can corrupt the routing tables.

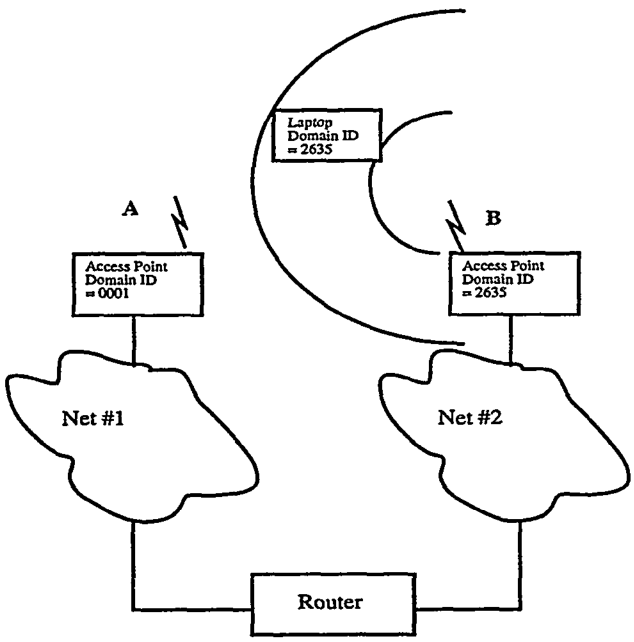

Figure 3: Routing Anomalies Due to Cell Overlap

4. Releasing the Network

In the last quarter of 1996, the design team felt all of the problems discussed above were reasonably solved. We then carefully audited the signal strength and network performance of all installed buildings, and so were in a position to understand what needed to be done to release the network to a broader user base. The decision was made to re-design and reinstall all six buildings and add a seventh, Baker Hall, the extension of Porter Hall. The rationale for this was that there are few visible cues in moving from Porter Hall to Baker Hall and that a mobile user would have reasonable expectations of continuity of service. In addition we added coverage of the two main common outside areas between the buildings, areas known as The Mall and The Cut. A combination of 6 dB and 12 dB directional antennas was utilized for this. With this final network of seven buildings accomplished by Spring 1997, we were now ready to notify the campus community and release the network. The released network covers 59 percent of the technical community of the University (i.e. excluding Fine Arts, Architecture and Warner Hall, a purely administrative building). Table 3 shows the details of cells for each of the buildings and may be contrasted with the original design placement before it was understood how to properly design. Note that there is no significant addition of access points (other than in the previously uncovered Baker Hall) but the placement has been significantly changed. The Table does not capture the fact that almost all access points were moved even if the total number of access points on a particular floor did not change significantly. The large cost of rewiring was the price of being first.

All of these seven buildings were officially released at the beginning of this year and an active website (http://www.ini.cmu.edu/WIRELESS/andrew.html) provides information to users on topics such as: availability of services and hardware, registration details, and where and how WaveLAN and Digital Ocean devices can be purchased. The website also handles user questions and trouble calls. Access devices were made available for purchase via the website or by telephone at that time. And so the community is beginning to build from the first, non-naive, selected group of 100 users.

The decision was also made to hold campus coverage at the level of seven buildings for the first public year while our attention turns to usage. Though we feel confident that we can now accurately design and install networks, we do not yet have answers to some of the challenges of expanding the network to the whole campus and challenges in taking a whole network operational. These include:

- better understanding of usage so that we can better design the help desk and other user support functions, as well as ensure consistent service levels;
- better network management to allow hand-over of the system to the technician support group of the main Andrew system; and
- implementation of transparent migration from the network to other networks on and off campus to allow users to use a single IP address.

4.1 Usage Patterns

An essential part of developing the wireless infrastructure further is to measure, analyze and better understand usage. This includes human factors that may present barriers not only to usage of the CMU network but also factors that effect market growth for similar systems. Two measurement systems have been put in place: a direct, per packet, measurement of usage and a human survey. With the permission of our user community, we log the traffic generated on our wireless network. This traffic logging is performed on the backbone of the network using a PC running UNIX. Header information from the packets necessary to distinguish the type of transaction is stored and the contents are discarded. This large amount of logging data is accumulated on recordable compact disks at the end of each day. Custom programs have been developed to access the data, performing preliminary analysis, and storing the results in an accumulating database. The database can then be queried to retrieve information such as:

- aggregate and average usage,
- usage by time of day,
- usage while roaming,
- amount of time spent on different applications,
- the amount of traffic generated by each application,

Cyert Hall (University Computer Center)

Floor	Old	New
Basement B	1	2
Basement A	6	2
1st	0	3
2nd	3	2
Total	10	9

Hamburg Hall (Information Networking Institute)

Floor	Old	New
A Level	3	8
1st Floor	7	4
2nd Floor	0	6
3rd Floor	3	1
Total	13	19

GSIA - Old Building (business school)

Floor	Old	New
Basement	3	3
1st Floor	1	0
2nd Floor	2	1
3rd Floor	0	1
Total	6	5

GSIA - Posner Hall (business school)

Floor	Old	New
Basement	2	2
1st Floor	1	1
2nd Floor	3	3
Total	6	6

Porter Hall (includes engineering)

Floor	Old	New
Basement B	1	2
Basement A	5	4
1st Floor	0	2
2nd Floor	5	1
3rd Floor	0	1
Total	11	10

Wean Hall (computer sciences)

Floor	Old	New
1st Floor	0	2
2nd Floor	4	2
3rd Floor	0	8
4th Floor	14	3
5th Floor	0	5
6th Floor	0	2
7th Floor	9	6
8th Floor	0	4
Total	27	32

Baker Hall (includes engineering)

Floor	Old	New
Basement	0	1
1st Floor	0	5
2nd Floor	0	2
3/4th Floor	0	6
Total	0	14

Outside Areas	**0**	**6**
Grand Total	**73**	**101**

Table 3: Access Points by Floor, Original and New Design

- locations of high usage, and
- roaming profiles.

We are currently capturing traffic [15] between a mobile station and:

- any machine in another wireless cell,
- any machine on the campus wireline network, or
- any machine external to CMU.

We are unable to capture peer-to-peer traffic within a cell as it would involve installing a monitoring station in each cell, which would be prohibitively expensive. However, based on sample measurements, we are reasonably certain that most traffic falls into the three categories above that we can measure. There are a number of problems however. For example, the WavePOINT bridging tables do not expire and hence we do not know if a mobile unit is switched off unless it tries to reregister elsewhere. An earlier problem where we could not track roaming was solved by Lucent providing a change to the WavePOINT software so that we can monitor sign-on and hand-over requests between a mobile unit and the access point of the cell that it is in. However, as mentioned, there is no way currently that we can locate the position of a mobile unit within a cell. Without location data being currently available, we are concentrating on analysis of applications that are being used by mobile users and comparing them to applications use over the wireline network. To do this we are monitoring: e-mail, http (web usage), Telnet, FTP, printing, IRC (chat room), use of the Finger utility, and unknown traffic which is none-of-the-above. During the coming year we will also establish average session durations.

So far, the monitor has been running for approximately three months. These have coincided with the summer semester -- an atypical semester. Usage is hence not currently high with the exception of http and unknown traffic. The http usage corresponds to typical work day experience monitored on the wired network and we currently suspect the unknown traffic to be related to isolated research projects. One item worth mentioning is that early results show little mobility: users seem to prefer to use wireless machines when stationary, but this again can be the result of most usage being for isolated research projects at this time. We are not rushing to draw conclusions but will explore more thoroughly the nature of the unknown traffic and await a full year's data collection and analysis as the user community grows and encompasses other than the immediate research community.

In addition to direct monitoring of actual usage, the usage team also undertook a paper survey of registered WaveLAN card owners to establish and measure how they perceive their usage in order to compare this with actual data. Users participating in the human survey completed questionnaires at the beginning of the trial and will again complete surveys at the end of the coming year. The surveys will uncover information that cannot be measured by logging traffic such as the percentage of time spent on various applications on the wired network, work hours, mobility, technical background and personal profile information.

Much of the structure and techniques of this measurement system have been based on experience gained in another CMU project, a wireline communication project called HomeNet, an urban laboratory of some 120 households connected to and monitored by CMU to determine usage behavior [16].

With the measurement system in place and data being collected and archived for detailed examination over the coming year, as the user community grows, the user monitoring system is expected to be a rich source of further research. Traffic studies will allow comparisons to be made between wireless and wired network usage. In addition, the logged information will allow certain key applications to be optimized for wireless usage, will check the significance of total traffic, and will allow tracking of roaming patterns, allowing the verification of various mobility models.

4.2 Operation

The network for the time being remains a research network and is not being handed off to Computer Services as part of the normal campus operation, hence, there is little to say about operation. Lucent's focus is on the 2.4 GHz product and there is little likelihood that SNMP improvements will be made to our network. Consequently, we are using the rudimentary system currently available and monitoring network performance and troubles. We have also started parallel development work into issues such as how to monitor RF in cells in a cost effective manner and how to better configure and fix problems remotely. We also have begun looking into the possibility of some automatic self-reconfiguration of the nodes, using configurable antennas and variable RF power, to accommodate changing demand and environment.

4.3 Mobile IP

Finally, concerning transparent migration, we have implemented Mobile IP to allow use of a single IP address for a user whether on the wireless system or various wireline systems. The IETF Mobile IP standard [17] allows mobile hosts to move about within the Internet, from one IP subnet to another, without the burden of rebooting or reconfiguration of either the mobile host or other hosts communicating with it. The mobile host's movement is completely transparent, and the host continues to be addressed by its home IP address, independent of its actual point of connection to the Internet in another subnet away from home. The base Mobile IP protocol is now at the first level of official Internet standards status (Proposed Standard), but little or no actual experience yet exists with the protocol in a real-world, operational setting. In order to further understand and refine the protocol, such experience is vital, and in order to further advance along the Internet standards track, such experience is required by the IETF.

To provide this experience, and to provide a valuable service to the CMU user community, four Pentium-based computers have been installed to deploy Mobile IP service on campus at CMU. Each of these four computers serves one of the four main subnets on campus, acting as a Mobile IP "home agent" (for mobile hosts for which this is their home subnet) and as a Mobile IP "foreign agent" (for mobile hosts visiting this subnet while away from their home subnet). The four subnets to be covered are the WaveLAN network, the School of Computer Science network, the Department of Electrical and Computer Engineering network, and the main campus Andrew wired network. These computers are currently being installed and will all run the FreeBSD version of the Unix operating system, using an implementation of Mobile IP written at CMU.

In addition, users will be able to transparently migrate to and from the CDPD network using their CDPD assigned IP address as a "co-located Care-of address" using Mobile IP.

CMU has been an active participant in the design and standardization of the Mobile IP protocol [18], and involvement in the IETF in further Mobile IP work will be continued, including the Route Optimization extensions to Mobile IP and the modifications to the new IPv6 protocol for supporting Mobile IP. As this new work reaches a sufficient level of maturity, these services will also be deployed on the CMU campus, utilizing the Wireless Andrew network and these home agent and foreign agent computers as an infrastructure for continued research and evaluation in mobile networking protocols.

5. Summary

The Wireless Andrew project is showing the challenges of designing and managing large scale wireless networks. The examples show the differences experienced in wireless vs. wireline networks due to the nature of mobility vs. static components, the nature of RF propagation vs. wireline connection, and the difference in sophistication of tools because of the different place in the product development cycle. However, the appeal of mobility has been aptly demonstrated by the incredible growth of mobile telephony over the last ten years. Add to this the ability of mobile computer systems to reach a class of users so far left out of the information revolution, namely people who don't sit at desks, and we see an inevitable drive towards networks such as we are creating on the campus of CMU. Answering the questions needed to get there, makes the return worth the investment.

At this point in time, we are looking into the next phase of three years. This includes CMU and Lucent seeking additional partners, in a consortium of common interested organizations, to use the network as a research platform, to extend the network throughout the campus, and to extend the network off campus to some of the homes of the campus community.

Further Information

Website http://www.ini.cmu.edu/WIRELESS/ contains additional material on the CMU Wireless Initiative, which Wireless Andrew supports. The Wireless Andrew project site includes reports and status. It can be seen at: http://www.ini.cmu.edu/WIRELESS/Wireless_Infrastructure.html

Acknowledgments

The Information Networking Institute would like to acknowledge and thank the National Science Foundation for its key support in this project under Grant CDA-9413346, Lucent Technologies for their support, and Bell Atlantic NYNEX Mobile Systems for their support by means of a service grant.

This paper represents the work of many more people than the authors hence special acknowledgment is given to the CMU community:

- Computing Services (especially Frank Skelley for early WaveLAN test data, Mark Campasano for WaveLAN test data and cellular interference test data as well as considerable general input, Jenny Ladd for Figure 2, Alex Hills for coverage ratio analysis, Erikas Napjus for input regarding network management, and Pete Bronder for general input);
- Telecommunications (especially Bruce Taylor for input regarding cellular and other interference testing);
- Ratish Punnoose of the INI, Harinderpal Grewal and Jon Peha of Electrical and Computer Engineering, and Sara Kiesler of Social and Decision Sciences, for input regarding the usage system;
- Dave Johnson, of Computer Science, for input regarding use of Mobile IP in the network; and
- the user community, for actively taking part and making the project work.

Thanks is also given to Lucent Wireless Communications Networks in Utrecht, Holland (especially Leo Montelban and Albert Eikelenboom for coverage test results and general input, and Nedim Erkocevic for cellular test specification).

References

[1] A. Smailagic and D. P. Siewiorek, "Modalities of Interaction with CMU Wearable Computers", *IEEE Personal Communications*, February, 1996, pp. 14-25.
[2] M. Satyanarayanan, " Mobile Information Access", *IEEE Personal Communications*, February, 1996, pp. 26-33.
[3] J. M. F. Moura, R. S. Jasinschi, H. Shiojiri, and J. C. Lin, "Video Over Wireless", *IEEE Personal Communications*, February, 1996, pp. 44-54.
[4] B. Bruegge and B. J. Bennington, "Applications of Mobile Computing and Communication", *IEEE Personal Communications*, February, 1996, pp. 64-71.
[5] J. H. Morris, M. Satyanarayanan, M. H. Conner, J. H. Howard, D. S. H. Rosenthal, and F. Donelson Smith, "ANDREW: A Distributed Personal Computing Environment", *Communications of the ACM*, Vol. 29, No. 3, March, 1986., pp. 184-201.
[6] W. Y. Arms, "Reflections on ANDREW", *EDUCOM Review*, Vol. 25, no. 3, Fall 1990, pp. 33-39,42-43.
[7] P. Clegg, "PCMCIA Radio Network Adapters", LAN Times October 3 1994
[8] B. Tuch, "Development of WaveLAN, an ISM Band Wireless LAN", *Lucent Technical Journal*, July/August, 1993, pp. 27-37.
[9] M. Schoffstall, M. Fedor, J. Davin, and J. Case, "A Simple Network Management Protocol (SNMP)", *Internet RFC 1157*, May 1990.
[10] A. Hills and D. B. Johnson, "A Wireless Data Network Infrastructure at Carnegie Mellon University", *IEEE Personal Communications*, February 1996, pp. 56-63.
[11] B. Kobb, "Spectrum Guide: Radio Frequency Allocations in the United States, 30 MHz-300 GHz", New Signals Press, 1996.
[12] D. Eckhardt and P. Steenkiste, "Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network", *Proceedings of SIGCOMM '96*, ACM, August 1996, Stanford, CA.
[13] S. C. Harlan, "CMU WaveLAN Testing", Bell Atlantic NYNEX Mobile, June 1996.
[14] P. Schnaidt and S. Reddy, "Case Study - Speedy and Smart", *LAN Magazine*, Vol. 6, No. 3, March 1991.
[15] H. S. Grewal, R. J. Punnoose, S. Kiesler, J. M Peha, "User Behavior and Mobility on Wireless Networks: A Preliminary Empirical Analysis of the WaveLAN Network at Carnegie Mellon University", Internal Report, Carnegie Mellon, May, 1997.
[16] R. Kraut, W. Scherlis, T. Mukhopadhyay, J. Manning, and S. Kiesler, "The HomeNet Field Trial of Residential Internet Services", *Communications of the ACM*, December, 1996, 39, pp. 55-63.
[17] C. Perkins, editor, "IP Mobility Support", *Internet RFC 2002*, October 1996.
[18] D. B. Johnson and D. A. Maltz, "Protocols for Adaptive Wireless and Mobile Networking", *IEEE Personal Communications*, February 1996, pp. 34-42.