Large-Scale Wireless LAN Design

Alex Hills, Carnegie Mellon University

ABSTRACT

The design of a large-scale wireless LAN poses a number of interesting questions. Building the Wireless Andrew network at Carnegie Mellon University has given us the opportunity to deal with these questions. This article describes the approaches we have developed for the design of similar networks.

A large-scale wireless LAN must be designed so that all of the target space has radio coverage (i.e., there are no coverage gaps). It must also be designed so that its capacity is adequate to carry the expected load. These requirements generally can be met by using the proper combination of access point locations, frequency assignments, and receiver threshold settings.

INTRODUCTION

To provide high-speed wireless service on our campus, Carnegie Mellon University has built a large-scale wireless local area network (WLAN). The wireless infrastructure we have installed is called *Wireless Andrew* for Andrew Carnegie and Andrew Mellon. Started in 1994 as a National Science Foundation-funded research network to support Carnegie Mellon's wireless research initiative, it originally provided coverage in seven campus buildings, but it has been expanded to serve all residential, academic, and administrative buildings on the campus. The network now covers 65 buildings, with a total floor area of approximately 3 million ft^2, as well as outside areas [1–4].

Wireless Andrew builds on the university's wired network infrastructure, which currently provides 10 Mb/s, 100 Mb/s, and Gigabit Ethernet service to the desktop. To provide high-speed wireless service to the campus, we have installed WLAN equipment conforming to the IEEE 802.11b standard, which uses 2.4 GHz direct-sequence spread spectrum (DSSS) radio to provide a raw data rate of up to 11 Mb/s [4].

The Wireless Andrew experience has given us insight into techniques for designing large-scale WLANs. Design techniques are discussed here in the context of IEEE 802.11b, but they may be applicable to other kinds of microcellular networks.

WIRELESS LANS

WLANs were originally intended to allow local area network (LAN) connections where premises wiring systems were inadequate to support conventional wired LANs. During the 1990s, because the equipment became available in the PCMCIA form factor, WLANs came to be identified with mobility. They can provide service to mobile computers throughout a building or throughout a campus.

Generally, wireless LANs operate in the unlicensed industrial, scientific, and medical (ISM) bands at 915 MHz, 2.4 GHz, and 5 GHz. The original WLAN standard IEEE 802.11 (with speeds up to 2 Mb/s) allows either direct-sequence or frequency-hopping spread spectrum to be used in the 2.4 GHz band. It also allows operation at infrared frequencies. [5] The high-rate WLAN standard IEEE 802.11b provides operation at speeds up to 11 Mb/s in the 2.4 GHz band and uses a modified version of the IEEE 802.11 direct sequence spread spectrum technique. A newer high-rate standard IEEE 802.11a uses orthogonal frequency-division multiplexing (OFDM) to provide for operation in the 5 GHz UNII band at speeds up to 54 Mb/s. [6] IEEE 802.11b equipment is readily available in the market, and IEEE 802.11a equipment is expected to become available by early 2002.

WLANs typically include both network adapters (NAs) and access points (APs). The NA is available as a PC card that is installed in a mobile computer and gives it access to the AP. The NA includes a transmitter, receiver, antenna, and hardware that provides a data interface to the mobile computer. The AP is a data bridge/radio base station which is mounted in a fixed position and connected to a wired LAN. The AP, which includes transmitter, receiver, antenna, and bridge, allows NA-equipped mobile computers to communicate with the wired LAN. The bridge which is part of the AP routes packets to and from the wired network as appropriate.

Each AP has a radio range, for communication with NAs, from approximately 20 to more than 300 m, depending on the specific product, antennas, and operating environment. The APs can be interfaced to IEEE 802.3 (Ethernet) wired LANs.

Most wireless LANs allow "roaming"; that is, mobile computers can accept a handoff as they move from the coverage area of one AP to the coverage area of another, so service is continuous. In order for this handoff to be successful, it is necessary that the tables of the bridges contained in each AP be updated as mobiles move from one AP coverage area to another. In wireless LANs, direct 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 others two mobiles, even though they are both within range of each other, can communicate only by having their transmissions relayed by an AP.

The use of direct sequence spread spectrum (DSSS) in IEEE 802.11 and 802.11b spreads the signal over a wide bandwidth, allowing transmissions to be robust against various kinds of interference and multipath effects. IEEE 802.11b WLANs operate at raw data rates of up to 11 Mb/s and occupy a transmission bandwidth of approximately 26 MHz. Exact spectrum allocations for 2.4 GHz ISM differ from one country to another. In North America the band is 2.400–2.4835 GHz.

IEEE 802.11 and 802.11b use the carrier sense multiple access (CSMA) with collision avoidance (CA) medium access scheme, which is similar to the CSMA/CD 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 the transmitter signal strength at its own antenna will be so much stronger than the signal received 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.

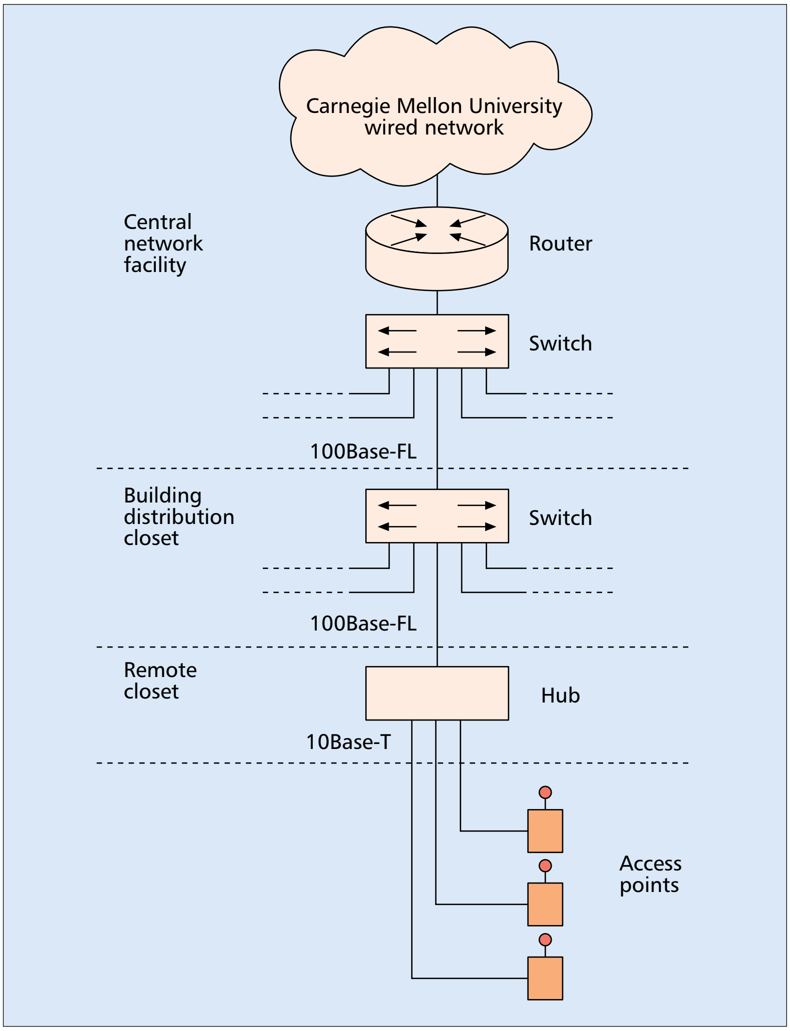

Figure 1. *The structure of Wireless Andrew.*

Structure of Wireless Andrew

Figure 1 shows how our large-scale WLAN is connected to the campus network. We have installed a new IEEE 802.3 backbone network on campus to connect the APs in each building with the rest of the campus wired network through a dedicated switch in the central campus network facility. Each AP is connected to an IEEE 802.3 10Base-T hub in a building's remote wiring closet near the AP. These hubs, in turn, are connected to a switch in the building's main distribution wiring closet, which is in turn connected to a switch in the university's central network facility. This central switch's connection to a router in the central network facility is the link that connects the wireless backbone to the remainder of the campus data network. The APs and hubs are powered by the campus 110 VAC system. (Some recent WLAN products allow APs to be powered by DC through communication cable, e.g., Category 5) [1].

This structure allows us to operate the high-speed wireless network independent of the campus network and to disconnect the two if necessary. The structure also allows us to separate traffic on the wireless LAN backbone network from traffic on the rest of the network. The router connecting the wireless backbone to the rest of the campus network filters packets based on destination address and only passes packets to and from the wireless backbone as needed.

Design Challenges

The challenges in building such a large wireless network are significant. They include designing the network so that coverage blankets the campus and adequate capacity is provided to handle the traffic load generated by the campus community. We define WLAN design as including two components: selection of AP location and assignment of radio frequencies to APs.

In laying out a multiple-AP wireless LAN installation, one must take care to ensure that adequate radio coverage will be provided throughout the service area by carefully locating the APs. Experience shows that the layout must 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 an entire campus [3].

The layout and construction of buildings deter-

In selecting AP locations, one must avoid coverage gaps, areas where no service will be available to users. On the other hand, one would like to space the APs as far apart as possible to minimize the cost of equipment and installation.

mine the coverage area of each AP. Typical transmission ranges go up to 300 m in an open environment, but this range may be reduced to 20–60 m through walls and other partitions in some office environments. Wood, plaster, and glass are not serious barriers to wireless LAN radio transmissions, but brick and concrete walls can be significant ones; the greatest obstacle to radio transmissions commonly found in office environments is metal, such as in desks, filing cabinets, reinforced concrete, and elevator shafts.

Network performance is also an issue. An AP and the mobile computers within its coverage area operate something like the computers on an Ethernet segment. That is, there is only a finite amount of bandwidth available, and it must be shared by the APs and mobile computers. The IEEE 802.11b protocol, using CSMA/CA, provides a mechanism that allows all units to share the same bandwidth resource [5, 7].

The CSMA/CA protocol makes radio interference between APs and NAs operating on the same radio channel a particular challenge. If one AP can hear another AP or a distant NA, it will defer, just as it would defer to a mobile transmitting within its primary coverage area. Thus, interference between adjacent APs degrades performance. Similarly, if a mobile unit can be heard by more than one AP, all of these APs will defer, thus degrading performance [7].

Design Approach

In selecting AP locations, one must avoid coverage gaps, areas where no service will be available to users. On the other hand, one would like to space the APs as far apart as possible to minimize the cost of equipment and installation. Another reason to space the APs far apart is that coverage overlap between APs operating on the same radio channel (co-channel overlap) degrades performance. (This is the reason one should not "overprovision" a wireless LAN by using more APs than necessary.) Minimizing overlap between APs' coverage areas when one is selecting AP locations helps to minimize co-channel overlap.

We have found that rules of thumb are inadequate in doing this type of design. Rather, each building design must be based on careful signal strength measurements. This is particularly challenging because the building is a three-dimensional space, and an AP located on one floor of the building provides signal coverage to adjacent floors of the same building and perhaps to other buildings, as well [3].

After the APs have been located and their coverage areas measured, radio channels are assigned to the APs. Eleven DSSS radio channels are available in the 2.400–2.4835 GHz band used in North America; of these, there are three that have minimal spectral overlap. These are channels 1, 6, and 11 [1]. Thus, in North America, APs can operate on three separate noninterfering channels. Furthermore, some NAs can switch between channels in order to talk with the AP providing the best signal strength or the one with the lightest traffic load. Use of multiple channels can be very helpful in minimizing co-channel overlap, which would otherwise degrade performance.

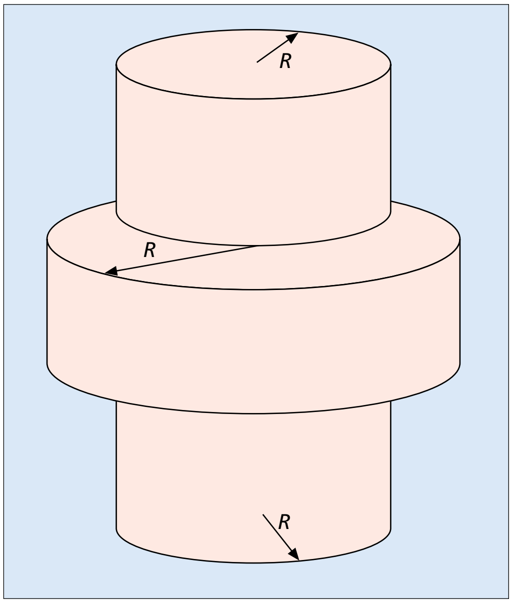

Figure 2. *Idealized access point coverage.*

Our approach is to assign one of these three channels to each of the APs and to do so in a way that provides the smallest possible co-channel coverage overlap. Making these frequency assignments is essentially a map coloring problem, and there are various algorithms that give optimal or near-optimal assignment of the three radio channels, given a particular set of AP placements and coverage areas.

The design must also consider service to areas with high and low densities of users. If many users of mobile computers are located in a small area (a high-density area), it may be necessary to use special design techniques in these areas. We expect that most parts of a campus will be low-density areas. However, there will be some areas, particularly classrooms and lecture halls, that will be high-density areas, with high concentrations of users, mostly students.

Two design layout techniques that are useful in high-density situations are increasing receiver threshold settings and using multiple radio channels. Some wireless LAN products allow one to set receiver threshold, thus controlling the size of the coverage area of the AP. A coverage-oriented design should use the minimum receiver threshold setting, maximizing the size of the coverage area of each AP. When capacity issues are considered, however, one may wish to use higher AP receiver threshold settings in high-density areas, reducing the coverage area of each AP.

The use of multiple radio channels can allow the use of multiple APs to provide coverage in the same physical space. For example, one might use three APs operating on three different channels to cover a large lecture hall with a high density of users. The exact capacity improvement is dependent on the algorithm used by the mobile to select an AP. A load balancing algorithm will provide the greatest capacity increase. An algorithm that selects the strongest AP signal will not provide as great an increase.

Thus, one would like to carry out a design that is coverage-oriented in most (low-density) areas, minimizing the number of APs, but capacity-oriented in some (high-density) areas, assuring adequate capacity to serve all users in these areas. The coverage-oriented design in the low-density areas minimizes the cost of APs, but the use of extra APs with higher receiver thresholds in high-density areas can be used to provide extra capacity.

DESIGN PROCEDURE

Because radio propagation inside a building is frequently anomalous and seldom completely predictable, the design of an in indoor wireless installation must be an iterative one. Our design procedure includes five steps:
- Initial selection of AP locations
- Test and redesign, which is adjusting the access point locations based on signal strength measurements
- Creation of a coverage map
- Assignment of frequencies to APs
- Audit, which is documenting the AP locations and a final set of signal strength measurements at the frequencies selected

In the next section we describe a technique for carrying out the first step, initial selection of access point locations. This initial design is a tentative one and is intended to be modified in the second step of the design process.

After the initial selection of AP locations is complete, APs are temporarily installed at the locations selected. The coverage areas of these APs and the overlaps between coverage areas are measured. Typically, coverage gaps and/or excessive overlaps are found. Based on the measurement results, the AP locations are adjusted as needed, more measurements are done, more adjustments are made, and so on, until an acceptable design is found. The process is an iterative one. It may be necessary to repeat this design-test-redesign cycle several times to find an acceptable solution.

After the final AP locations have been selected, a coverage map of the design area is created. This coverage map may be created using AutoCAD or other computer-based techniques.

After AP locations have been finalized, frequencies are assigned to the APs in a way that minimizes co-channel coverage overlap. Then a complete set of coverage measurements (audit) is made for the entire building with the APs operating at the selected frequencies, and the results of these measurements are documented. At this point the design is considered complete. The coverage map is updated to reflect final AP locations, coverage measurements, and frequency assignments.

DETERMINING THE ACCESS POINTS' INITIAL LOCATIONS

We now describe a procedure for the initial selection of AP locations in a low-density area. In selecting locations for the APs, one should place them so that there are no coverage gaps in the target space, and the coverage overlaps between and among APs are minimized. While the first point is obvious, the second is more important than is immediately apparent. If too many APs are used, the cost of equipment and installation will be higher than necessary, and the performance of the network may also be degraded if the final design involves a great deal of co-channel coverage overlap. The amount of co-channel coverage overlap is determined by both AP placement and AP frequency assignment.

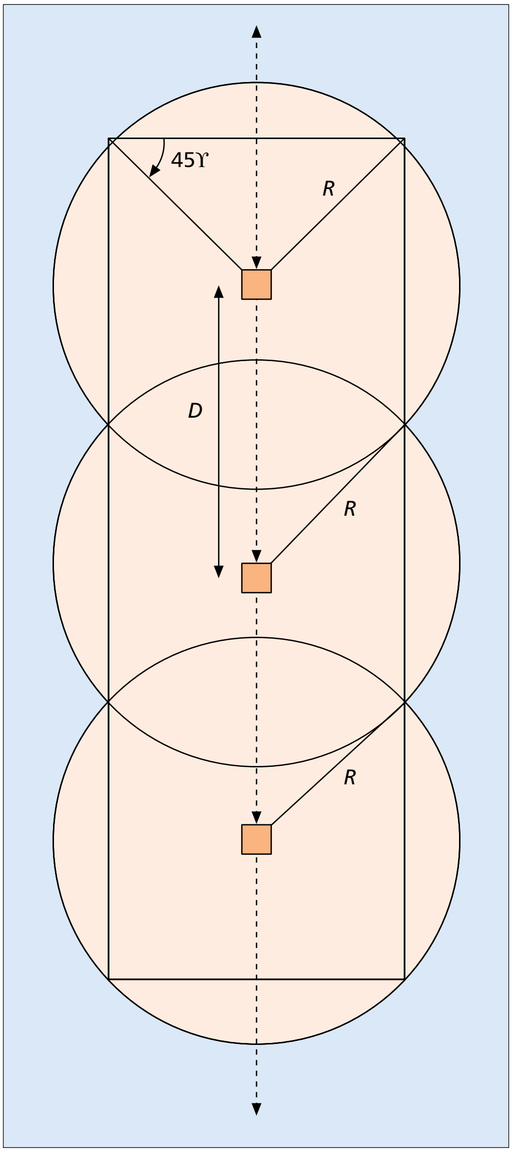

Figure 3. *A linear array of APs in a single-floor building.*

While coverage volumes are not actually perfect cylinders, one can find the average coverage radius inside a building and use this as the radius of an idealized cylindrical coverage volume.

We define coverage area in terms of a specified received signal strength. This threshold level is selected in order to provide an adequate signal-to-noise ratio (S/N) and some additional margin. If, for example, in designing an IEEE 802.11b WLAN, one measures an ambient noise level of –95 dBm and a 10 dB S/N is needed to ensure excellent performance, one might decide to allow an extra 5 dB of margin to allow for noise levels higher than –95 dBm. In this case one would select a threshold of –80 dBm.

When high-density spaces exist, we suggest that the AP placement first be done for these spaces and that the remaining low-density spaces then be designed, filling in the gaps between high-density spaces.

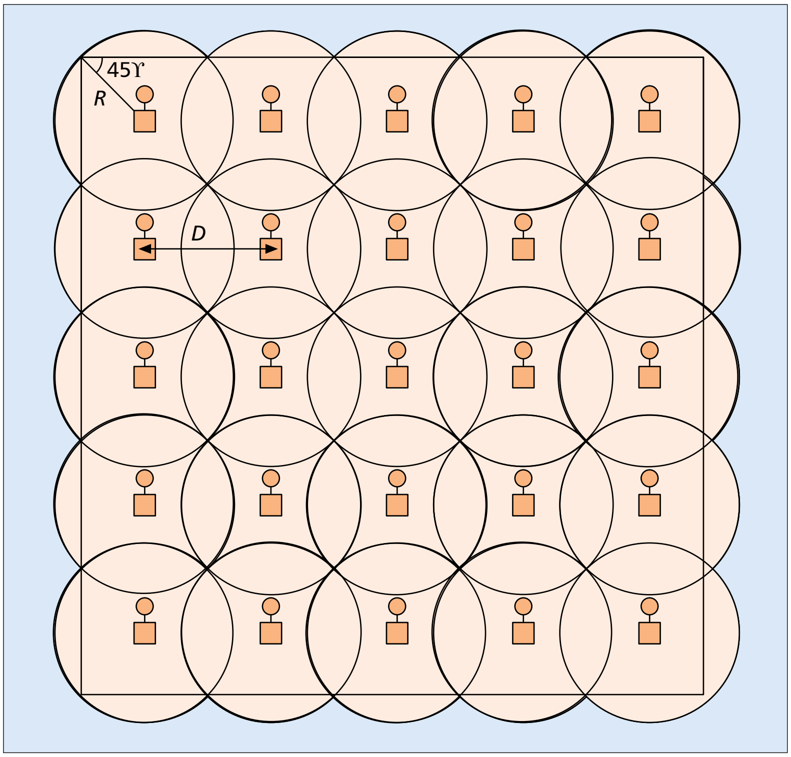

■ **Figure 4.** *A rectangular array of access points in a single-floor building.*

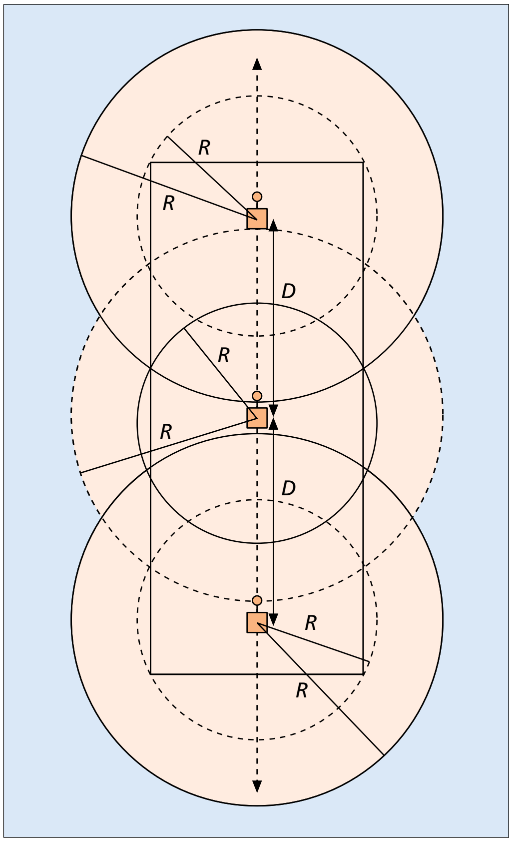

■ **Figure 5.** *A linear array of access points in a multifloor building.*

AP Placement

In this section we introduce an idealized notion of AP coverage. This description is offered only to provide some insight into the layout approaches that can be used in different types of buildings.

We idealize the coverage volume of the AP as three coaxial cylinders, as shown in Fig. 2. The middle cylinder, representing coverage on the floor on which the axis point is located, has radius *R*. The AP is located on the axis of this cylinder. The upper and lower cylinders, representing coverage on the floors above and below the one on which the AP is located, have radius *R*′, which is less than *R*. The height of each of the three cylinders is the height of a floor in the building. These three cylinders can be thought of as a single object, which moves about as the location of the AP moves.

The problem of locating APs within a building can be viewed as a problem of locating these shapes within the building in such a way that all spaces are filled with as little overlap as possible. While coverage volumes are not actually perfect cylinders, one can find the average coverage radius inside a building and use this as the radius of an idealized cylindrical coverage volume. This can be achieved by defining an acceptable signal strength threshold (e.g., –80 dbm) and determining the average distance from the AP at which signals fall below the threshold.

Procedure

The initial selection of AP locations begins with a complete set of signal strength measurements within the building. Signal strength measurements should be made in all areas of the building, with particular attention to the building's construction so that the characteristics within each part of the building are understood. These measurements have two purposes: to divide the building into spaces that are relatively isolated from each other from a signal propagation perspective and to determine the typical coverage radius of an AP. Signal strength measurements should be taken to determine the same floor coverage radius *R*′ and the adjacent floor coverage radius *R*′ of an AP.

Access points can be placed within a building in an array that is either linear or rectangular. An example of a linear array is shown in Fig. 3, and an example of a rectangular array is shown in Fig. 4. Each of these shows how APs can be located in a single-floor building or in a building with only one floor needing WLAN coverage. It is only necessary to locate the APs in a way that provides coverage throughout the floor and also minimizes as far as possible the overlap between and among AP coverage areas. A linear array is used when the building is narrow relative to *R*, and a rectangular array when the building width is large relative to *R*.

On the other hand, in a building that requires coverage on more than one floor, adjacent floor coverage must be considered in locating each AP. Usually, a staggered approach is used. As one moves along the length (or width) of a building, one places APs first on one floor and then on an adjacent floor. In this case the coverage of an AP's adjacent floor coverage must "dovetail" with the coverage of the next AP's same-floor coverage. As in a single-floor building, a linear array is used when the building is narrow relative to R, and a rectangular array when the building width is large relative to R.

We illustrate by using four scenarios one will encounter when designing an indoor wireless network. Each is determined by whether the building is single-story or multistory and by the width of the building relative to R and R'. In each case we give the appropriate layout approach and list the figure that illustrates it. Solid lines show coverage on a floor; dashed lines show adjacent floor coverage.

Scenario 1: Single-floor linear array, illustrated in Fig. 3. A single-story building (or a building that requires wireless coverage on only one floor) whose width (smallest outer dimension) is not large relative to R. D denotes the distance between adjacent APs.

Scenario 2: Single-floor rectangular array, illustrated in Fig. 4. A single story building (or a building that requires wireless coverage on only one floor) whose width (smallest outer dimension) is large relative to R. D denotes the distance between adjacent APs.

Scenario 3: Multifloor linear array, illustrated in Fig. 5. A multistory building whose width (smallest outer dimension) is not large relative to R and R'. D' denotes the distance between adjacent APs on different floors.

Scenario 4: Multifloor rectangular array, illustrated in Fig. 6. A multistory building whose width (smallest outer dimension) is large relative to R and R'. D denotes the distance between adjacent APs on the same floor, and D' denotes the distance between adjacent APs on different floors.

Frequency Assignment

After the AP locations have been finalized and a coverage map has been created, frequencies are assigned to the APs. In the United States and Canada, three nonoverlapping channels (channels 1, 6, and 11) are used. Thus, one can assign one of these three frequencies to each AP, doing so in a way that minimizes co-channel overlap. Assignment of frequencies is essentially a map coloring problem with three colors.

A variety of algorithms can be used to assign AP frequencies when the AP coverages are known. One can do this exhaustively by checking the co-channel overlap for all possible frequency assignments, and this is a reasonable approach if a computer is being used. Other, less time consuming algorithms are also possible, and some of these can give near optimal results. Another approach is to use the building coverage map that has been created to visualize the coverage overlaps and assign frequencies so that co-channel APs have only small coverage overlaps.

We recommend assigning AP frequencies in high-density areas before low-density areas. If, for example, one uses three APs to cover a high-density space, three different channels should be assigned to these APs. These frequency assignments will subsequently need to be considered in assigning frequencies to nearby APs covering low-density areas. This is true because APs covering the high-density space will usually have some coverage overlap with APs covering only low-density areas.

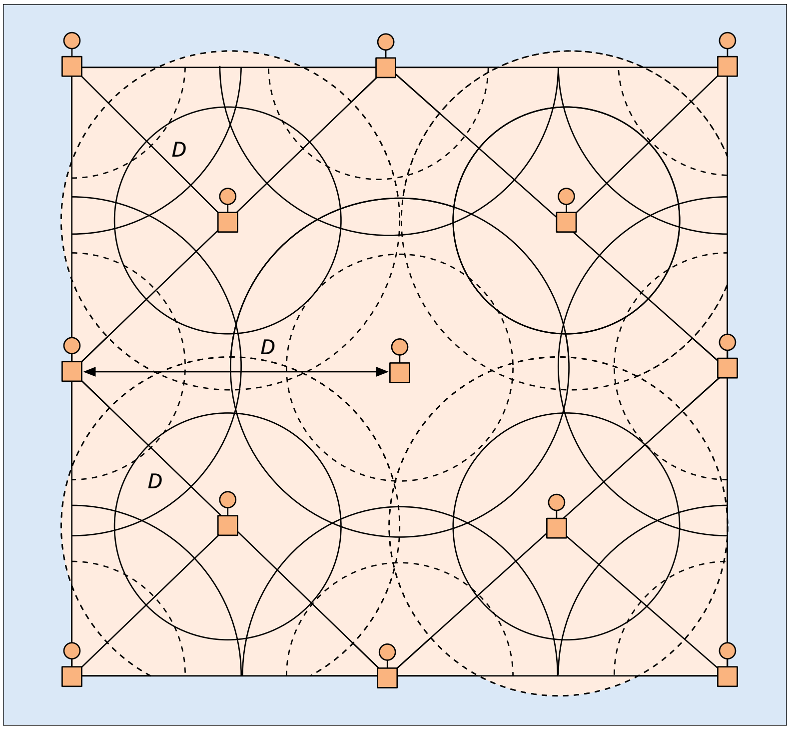

■ **Figure 6.** *Rectangular array of access points in a multi-floor building.*

Design Tool

In order to help with the design process, we have developed a computer-based tool, which assists with AP placement and frequency assignment. The tool allows one to measure AP coverage areas by merely walking around the target space. It generates coverage maps, showing the coverage areas of all APs. It also gives the optimal frequency assignments for this set of coverage patterns, allows the user to use drag and drop to virtually relocate APs, and predicts a new set of coverage patterns for the virtually relocated APs. We have found this tool to be extremely helpful in efficiently designing WLAN installations. A typical display, showing the coverage areas of a few APs on one floor of a building, is shown in Fig. 7.

Conclusion

The design of a large-scale IEEE 802.11b WLAN should be done in a way that ensures complete coverage of the target space and adequate capacity to carry the anticipated traffic load. The design must consider both the selection of AP locations and the assignment of frequencies to the APs.

AP locations should be selected so that all of the target space has radio coverage (i.e., there are no coverage gaps). AP locations should be selected and frequencies assigned in order to minimize co-channel coverage overlap. In high-density areas coverage overlap can be used (with different frequencies) to provide increased capacity. Another technique useful in serving high-density areas is increasing receiver thresholds in order to reduce APs' coverage areas.

The design of a large-scale IEEE 802.11b WLAN should be done in a way that assures complete coverage of the target space and adequate capacity to carry the anticipated traffic load. The design must consider both the selection of AP locations and the assignment of frequencies to the APs.

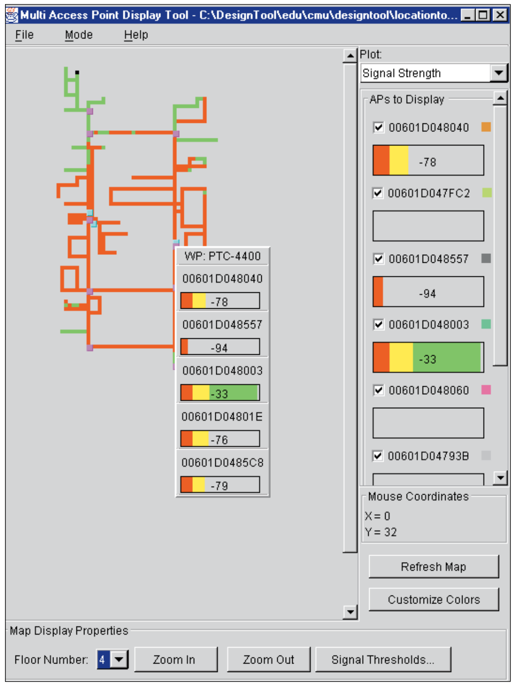

Figure 7. *Design tool display.*

Acknowledgments

There were many contributors to the Wireless Andrew project. Among these contributors were Charles R. Bartel, Bernard J. Bennington, Peter W. Bronder, Lawrence R. Gallagher, Mark Campasano, and Lisa M. Picone. In addition, Jon P. Schlegel, Mitchell A. Rappard, and Benjamin C. Jenkins contributed to the development of the design tool.

References

[1] A. Hills, "Wireless Andrew," *IEEE Spectrum*, vol. 36, no. 6, June 1999, pp. 49–53.

[2] A. Hills and D. B. Johnson, "A Wireless Data Network Infrastructure at Carnegie Mellon University," *IEEE Pers. Commun.*, vol. 3, no. 1, Feb. 1996, pp. 56–63.

[3] B. J. Bennington and C. R. Bartel, "Wireless Andrew: Experience Building a High Speed, Campus-Wide Wireless Data Network," *3rd Annual ACM/IEEE Int'l. Conf. Mobile Comp. and Networking*, Budapest, Hungary, Sept. 1997, pp. 55–65.

[4] http://www.cmu.edu/computing/wireless/index.html

[5] B. P. Crow *et al.*, "IEEE 802.11 Wireless Local Area Networks," *IEEE Commun. Mag.*, vol. 35, no. 9, Sept. 1997, pp. 116–26.

[6] R. van Nee *et al.*, "New High-Rate Wireless LAN Standards," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 82–88.

[7] H. S. Chhaya and S. Gupta, "Performance of Asynchronous Data Transfer Methods of IEEE 802.11 MAC Protocol," *IEEE Pers. Commun.*, vol. 3, no. 5, Oct. 1996, pp. 8–15.

Biography

ALEX HILLS [SM] (ahills@cmu.edu) is Distinguished Service Professor of Engineering & Public Policy and Computer Science at Carnegie Mellon University, where he also served until 1999 as vice provost and chief information officer. He is the founder of Carnegie Mellon's Wireless Andrew project. His teaching and research focus on wireless technology and telecommunications policy. Over the course of his career he has been a university professor and executive, a state government official, a broadcast station manager, a U.S. Army Signal Corps officer, a computer designer, a radio announcer, and a commercial fisherman. He is an avid runner, backpacker, and cross country skier