
Radio Resource Management in Wireless LANs



ALEX HILLS, CARNEGIE MELLON UNIVERSITY
BOB FRIDAY, AIRESpace, INC.

Abstract

Some chief information officers and information technology managers are reluctant to deploy wireless LANs. Among their concerns are reliability, availability, performance, and deployment. Each of these concerns can be directly addressed through the radio resource management techniques used in a new generation of wireless LAN equipment. The new capabilities include dynamic channel assignment, dynamic power control, and load sharing. Changing from the relatively static radio resource management techniques generally in use today to dynamic methods like those highlighted in this article will help to increase capacities and improve performance of large-scale wireless LANs.

Introduction

It is well known that the IEEE 802.11/wireless LAN market is expanding rapidly, but there is more to the story. While first-generation IEEE 802.11 technology is adequate for residential and small office/home office (SOHO) customers, the same is not always true for enterprise customers. In fact, some chief information officers (CIOs) and information technology managers are reluctant to deploy wireless LANs. Among their concerns are security, reliability, availability, performance under load, deployment, mobility, and network management.

While security is often mentioned as CIOs' greatest worry about wireless, some of their other concerns (e.g., reliability, availability, performance, and deployment) can be directly addressed through radio resource management techniques. The use of such techniques can make possible the rapid deployment of wireless infrastructure with much greater flexibility than has previously been available.

Wireless LAN technology evolved gradually during the 1990s, and the IEEE 802.11 standard was adopted in 1997 [1, 2]. As the technology has evolved, current products do not scale as well as they might in large-scale enterprise networks. Furthermore, as IEEE 802.11 wireless networks have become increasingly popular and more widely deployed, the need to expand the functionality of wireless LAN equipment has become obvious. In fact, IEEE 802.11 task groups and study groups are working to improve the standard, and some of these improvements will help with the issues that have been mentioned.

This article discusses how radio resource management is beginning to be used to mitigate some of the problems in enterprise wireless LANs. First, the problems of reliability, availability, performance, and deployment are described. Then a new architecture and its use with three radio resource management techniques are explained.

Enterprise Wireless LAN Problems

Enterprise CIOs want wireless networks that have several important qualities. High security (which is not addressed by radio resource management) is one, but in addition they want

networks that are highly reliable and available with very little downtime. The networks should also perform well (i.e., be capable of high throughput and low latency). The ideal is to have reliability, availability, and performance that are comparable to those of wired enterprise networks.

In addition, it should be possible to deploy wireless networks very quickly and without the need for extensive and time-consuming site surveys. Furthermore, the networks should have the flexibility needed to adapt to shifts in traffic loads and changes in the radio environment.

Reliability and Availability

The user experience in a wireless LAN is dependent on the radio propagation environment in which the wireless LAN operates. The radio propagation environment may change from time to time, affecting connection speeds and error rates. In a manufacturing environment, for example, where the multipath environment changes as equipment is moved about, it is quite possible for a link to fail completely even if the mobile is stationary.

To make matters worse, network management personnel in information technology departments are often unable to manage the network all the way to the user's mobile computer. The network operations center may be completely unaware of the ways in which radio propagation conditions are affecting the quality of the service received by the user. In fact, network personnel may be dependent on users to notify them of some types of problems. (This contrasts with many CIOs' increasing interest in network management solutions that extend all the way to users' laptops and other handheld devices.)

Furthermore, as users move around within the wireless LAN, the number of users being served by an access point (AP) may vary dramatically from a very low number to a very high one. When an AP is called on to serve a high number of users, it is likely to become overloaded, and the resulting congestion may significantly degrade the service received by users dependent on that AP.

These and other factors inevitably lead to problems of reliability and availability. Users have become accustomed to high levels of reliability and availability in wired data networks. The ideal is for wireless LANs to provide similar reliability and availability, and, to the extent possible, wireless LAN designers strive to provide the kind of service to which wired network users have become accustomed.

Performance

Wireless LANs often provide worse performance than wired data networks. Wireless users are likely to experience slower service, which is the result of the often lower transmission speeds and higher error rates on wireless links.

Raw data rates on wireless LANs are typically lower than those on Ethernet networks. For example, at one time most Ethernet networks provided 10 Mb/s service to users. At that time, wireless LAN technology provided no more than 1 or 2 Mb/s raw data rate to the user. Today the numbers have

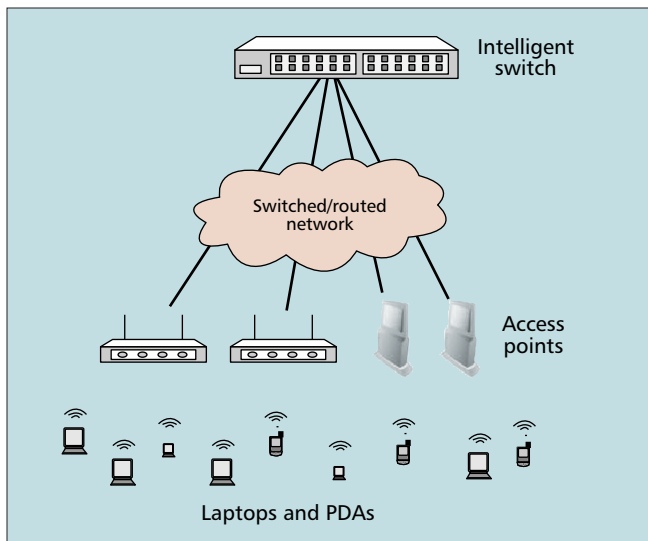


FIGURE 1. *Wireless LAN architecture with access points controlled by an intelligent switch.*

increased considerably, but there is still a gap. Ethernet networks can provide 100 Mb/s or even 1000 Mb/s service to users. While IEEE 802.11n promises higher speeds, wireless LANs currently in use operate at speeds up to 11 Mb/s in the case of IEEE 802.11b, and up to 54 Mb/s in the case of IEEE 802.11a/g.

But these numbers make wireless LANs sound much faster than they really are. Like a stepdown modem, a station in a wireless LAN will reduce its data rate from the maximum (11 or 54 Mb/s) to a lower rate, one that is sustainable by the quality of the radio path that exists. So, for example, an IEEE 802.11b link may step down from the data rate of 11 Mb/s to 5.5, 2, or even 1 Mb/s. Furthermore, overhead bits dramatically reduce the effective data rate available [1, 3–5].

Congestion further reduces the throughput experienced by a user. An AP and the mobile computers it serves share a single radio channel. As with Ethernet, when the traffic level or number of active computers is high, congestion occurs, and poor performance is the result. With switched Ethernet, however, each station can have a segment of its own. On the other hand, all stations using an IEEE 802.11 AP share the same bandwidth resource, and congestion is likely to be particularly severe in areas of high user density [3, 6]. It is highly desirable for wireless LAN equipment to include provisions to mitigate this problem.

Deployment

Although it may appear that one attraction of a wireless LAN is ease of deployment, in fact, the deployment of a well-designed, large-scale wireless LAN requires a careful site survey and design, which can be difficult and time consuming.

There are normally two parts of a careful wireless design: selection of AP locations and assignment of radio channels to the APs. The design is usually based on signal strength measurements and on consideration of radio propagation issues. This can be challenging because the building is a three-dimensional space, and an AP located on one floor of the building may provide signal coverage to adjacent floors of the same building and perhaps to other buildings, as well [3, 7].

The first part of the process, selection of AP locations, should be done to provide complete coverage of the target space without undue coverage overlap. Consideration of the characteristics of the radio propagation environment in which the wireless LAN is being deployed can be difficult but is important in a wireless LAN design [7].

In a coverage-oriented design one would like to space the

APs as far apart as possible while still providing complete coverage of the target space. This will minimize equipment and installation costs, and it will also allow the minimization of coverage overlap between APs operating on the same radio channel. Such “co-channel overlap” degrades performance [7].

Channel assignment, the second part of the design process, is normally done in a way that minimizes co-channel overlap. This is because, with carrier sense multiple access with collision avoidance (CSMA/CA), the IEEE 802.11 multiple access scheme, co-channel overlap causes interaction between stations in different cells, degrading performance [7].

A good site survey and design for a large-scale enterprise wireless LAN requires radio expertise. Since most data communications personnel lack this kind of expertise, enterprise organizations often provide personnel with the necessary training or hire an outside company to handle deployment.

Careful site survey and design are time consuming but are important to the successful deployment of first-generation wireless LAN networks. Although many customers attempt to shorten the process, they may experience performance problems resulting from a less than adequate design.

Architecture

In a first-generation IEEE 802.11 wireless LAN, the network’s intelligence is distributed among the APs. But radio resource management techniques require access to information that must be gathered across a number of APs, and the techniques involve control decisions that apply to a number of APs, not just one.

Thus, some centralized decision making is appropriate. An architecture that makes this possible is shown in Fig. 1. Illustrated is a wireless LAN that includes APs deployed throughout the target space but connected to and controlled by an intelligent switch. The software running on the intelligent switch is capable of collecting information from the APs and also sending control signals back to the APs, thereby making radio resource management and other features possible. Depending on the number of APs in an installation, it may be necessary to use more than one intelligent switch in the same network.

In this architecture, lightweight APs are little more than radio transceivers. Their operating parameters (e.g., radio channel, transmit power) can be controlled by the intelligent switch.

The intelligent switch controls a number of APs, and the intelligence that makes radio resource management and other features possible resides in the intelligent switch’s software. This software can instruct the APs on which channels they should operate, which transmit powers they should use, whether or not to accept association requests from specific clients, and so on. These actions, among others, are needed to make dynamic channel allocation, dynamic transmit power control, and load sharing possible, which can be expected to significantly improve the performance of the wireless LAN. They also will make the site survey and design process easier because these techniques can, to some degree, compensate for errors in designing the network.

Dynamic Channel Assignment

There are 14 2.4 GHz radio channels available for use in IEEE 802.11b/g networks. Table 1 lists these channels, and which of them can be used in North America, most of Europe, Spain, France, and Japan. Each of these channels has substantial bandwidth, and therefore they have significant spectral overlap with each other. In North America channels 1–11 are available for use. Of these channels, channels 1, 6, and 11 have minimal spectral overlap with each other, so these three

Channel number	Frequency (GHz)	North America	Most of Europe	Spain	France	Japan
1	2.412	x	x			x
2	2.417	x	x			x
3	2.422	x	x			x
4	2.427	x	x			x
5	2.432	x	x			x
6	2.437	x	x			x
7	2.442	x	x			x
8	2.447	x	x			x
9	2.452	x	x			x
10	2.457	x	x	x	x	x
11	2.462	x	x	x	x	x
12	2.467		x		x	x
13	2.472		x		x	x
14	2.483					x

TABLE 1. IEEE 802.11b/g channels in various parts of the world [3].

channels are frequently used in wireless LANs in North America [1, 3].

On the other hand, the IEEE 802.11a channels are laid out somewhat differently. The 5 GHz unlicensed band, called the unlicensed national information infrastructure (U-NII) band, has 12 available channels, exclusive of the new 5 GHz spectrum recently allocated for unlicensed operation by the U.S. Federal Communications Commission (FCC) [5, 8]. These fall into three contiguous blocks, the U-NII lower, middle, and upper bands (Table 2). Unlike IEEE 802.11 b/g, there is no significant spectral overlap, even between adjacent channels, in these bands, and a transmitter sending on one channel will cause little interference on an adjacent channel [5].

A recent order of the FCC opened a new band, 5.470–5.725 GHz, to unlicensed operation [8]. It is likely that this new spectrum will add approximately 11 new channels to those that can be used in IEEE 802.11a networks.

As described above, the performance of a network depends, in part, on the assignment of radio channels to APs. This assignment is often done using a manual process in which the designer attempts to assign the channels in a way that minimizes co-channel overlap. The coverage areas, and therefore the channel assignments, are dependent on, among other things, the radio propagation environment. But the radio propagation environment changes, so one cannot be sure that the channel assignments valid at the time the network was designed will continue to be valid.

Using radio resource management techniques, a second-generation wireless LAN can sense the radio environment from time to time and then dynamically adjust channel assignments accordingly. Such a process eliminates the necessity of doing the channel assignment as part of the design, and also helps to ensure that the current set of channel assignments reflect the current radio environment.

As described above, the optimal channel assignment for a given wireless LAN should minimize the overlap between coverage areas of co-channel APs. This will enhance the performance of the network by reducing interaction between co-channel APs and between clients associated

with different APs but operating on the same channel. The optimal set of channel assignments is the one that minimizes co-channel coverage overlap.

The problem of assigning channels to APs can be characterized as a graph-coloring problem. For example, Fig. 2 shows the graph associated with a small-scale IEEE 802.11b wireless LAN. The nodes represent APs whose channels are to be assigned. The edges represent coverage overlaps between APs, and the weights associated with the edges represent the amount of overlap measured in square meters.

One would like to color the nodes (representing APs) of a graph in a way that minimizes the sum of the weights of the edges (representing coverage overlaps) connecting nodes of the same color (representing channel). This is equivalent to assigning channels to APs in a way that minimizes the total co-channel coverage overlap. Converting the channel assignment problem to a graph-coloring problem allows the known techniques used in graph coloring to be employed directly to solve the problem of channel assignment [9, 10].

Since it is not feasible to directly measure coverage overlaps while the network is in operation, inter-AP signal strengths (signal strengths received by APs from other APs) can be used as a proxy for coverage overlap. Thus, one can reasonably assume that two APs that receive high signal strengths from each other have a high degree of coverage overlap.

But there are also other factors that influence the selection of the optimal set of channel assignments. These other factors include the noise and non-802.11 interference being received by each AP on each of the channels available for use, and the interference being received from other 802.11 networks. If such interference and noise are considered, it is possible that some APs will receive more interference and noise on some channels than on others. If this is the case, one would like to avoid assigning these APs to the channels on which they are receiving higher levels of interference and noise.

Finally, since an AP, along with its associated clients, might transmit only occasionally (low utilization or duty cycle), it is less likely to cause co-channel interference to other APs using the same channel than an AP that, along with its associated clients, is transmitting frequently (i.e., having a high utilization or duty cycle). Thus, one might, if necessary, assign this AP to a channel that causes some co-channel interference because little damage will be done.

Band (GHz)	Operating channel numbers	Channel center frequencies (MHz)
U-NII lower band (5.15–5.25)	36	5180
	40	5200
	44	5220
	48	5240
U-NII middle band (5.25–5.35)	52	5260
	56	5280
	60	5300
	64	5320
U-NII upper band (5.725–5.825)	149	5745
	153	5765
	157	5785
	161	5805

TABLE 2. IEEE 802.11a channels in the United States [5].




FIGURE 2. The graph associated with a small-scale wireless LAN. Nodes represent APs whose channels are to be assigned. Edges represent coverage overlaps between APs.

Thus, inter-AP signal strengths should be considered in combination with AP utilization levels.

Assignment of channels to APs can be made systematically according to a set of decision criteria that consider interference among APs, AP utilization, noise and non-802.11 interference, and interference from other 802.11 networks. Channel assignments can be dynamically updated as the operating environment changes in order to maintain optimal network performance.

To illustrate the performance improvement of dynamic channel assignment, we tested an IEEE 802.11b wireless LAN with four APs operating on channels 1, 6, and 11. With 10 fixed clients, each streaming traffic at 1 Mb/s, we moved another client to each of 41 locations in the space covered by the network. Without dynamic channel assignment operating, excessive co-channel overlap caused degradation in performance, but, as shown in Fig. 3, the situation was improved with dynamic channel assignment in operation. Figure 3 is a histogram that shows the throughput distribution across the 41 locations both with and without dynamic channel assignment.

Dynamic Transmit Power Control

AP locations and antenna types are normally selected by the wireless LAN designer in order to provide complete coverage of the target space. In most cases the designer will necessarily include some coverage overlap between APs but attempt to keep such overlaps to a minimum because the coverage overlaps make it difficult to assign radio channels in a way that minimizes co-channel coverage overlap.

But the radio environment may change from time to time. Thus, the conditions that existed when the APs' locations and antenna types were selected may no longer exist. If one could adjust the transmit power of each AP to reflect the current radio environment, one could perhaps maintain continuous coverage throughout the space without undue overlap. One

difficulty with this approach, however, is that AP transmit power control only affects the coverage area in the downlink direction. If the transmit power of the client is unchanged, the effective coverage area of the AP in the uplink direction is also unchanged.

Still, it is clear that dynamic transmit power control has the potential to reduce the effort involved in the site survey and design of a wireless LAN. It may be possible to carry out an abbreviated site survey and design process, placing APs in good, if not the best, locations and allowing the dynamic transmit power control capability to make the necessary adjustments.

After the wireless LAN is operational, there may be changes in the propagation environment (e.g., when equipment is moved in a factory) that can cause gaps or "holes" in coverage. Dynamic transmit power control can compensate for such changes, filling in the coverage holes. Thus, the technique can help to ensure that coverage will remain continuous throughout the target space.

Furthermore, APs fail from time to time. Depending on the exact AP locations and antenna types, dynamic transmit power control can be helpful in temporarily filling in coverage holes caused by AP failures.

To facilitate dynamic channel assignment, one would like to set the transmit powers of the APs in a wireless LAN in a way that will provide complete coverage without excessive coverage overlap. Dynamic transmit power control can improve the results of the channel assignment process, accommodate changes in the propagation environment, and also compensate for lost coverage due to failed APs.

As with dynamic channel assignment, one can use inter-AP received signal strengths as a proxy for coverage overlap. If the system is well calibrated, this technique can be effective. APs listen to each other's signals, and each AP's transmit power is set in a way that will achieve the desired signal strengths at (and coverage overlaps with) other APs. The situation is complicated somewhat by the fact that path loss as well as inter-AP received signal strength between APs depend on the physical environment. Thus, such an approach depends on calibration of the target received signal strengths to the radio environment.

Dynamic transmit power control can be very helpful in reducing coverage overlap, as shown in Fig. 4, which is based on radio frequency (RF) propagation modeling results and presented here for illustration purposes. In Fig. 4a, with no transmit power control the APs have significant coverage overlap, even when one considers the -50 dBm coverage contours. In Fig. 4b, with transmit power control

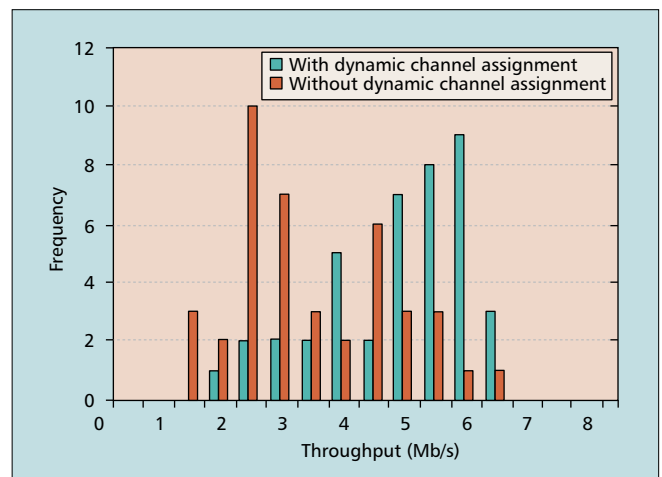


FIGURE 3. Throughput distribution across 41 locations in an IEEE 802.11b network with and without dynamic channel assignment in operation.

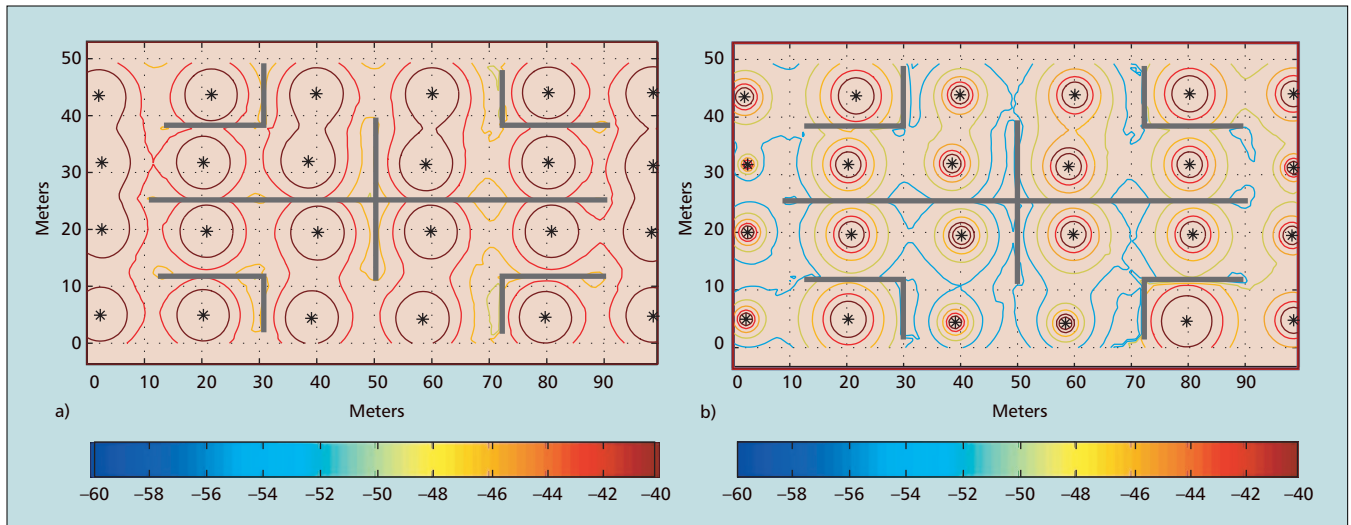


FIGURE 4. Propagation modeling results showing received signal strength contours (in dBm) without transmit power control (4a) and with transmit power control (4b).

operating the coverage overlap has been significantly reduced. There is virtually no coverage overlap at the -50 dBm level, an even at -60 dBm the coverage overlap has been reduced considerably.

As noted, the transmit power control technique described here affects only AP transmit power and therefore only the downlink coverage area. At the present time the IEEE 802.11 standard does not provide a way for the wireless LAN infrastructure (e.g., an intelligent switch/AP combination) to control clients' transmit powers, and therefore uplink coverage areas. It would be very helpful if an AP were able, via the air interface, to request that a client increase or decrease its transmit power. Fortunately, a study group within the IEEE 802.11 Working Group is currently considering the transmission of client power control and other control signals to be sent over the air interface.

Load Sharing

Since an AP and its associated clients share a limited bandwidth resource, APs can become overloaded, leading to congestion and poor network performance. On the other hand, a client may be able to communicate quite successfully with two or more APs. Thus, one would like to have a wireless LAN that is capable of distributing client associations among APs more or less uniformly so that no one AP is unduly overloaded. Wireless LAN equipment with this capability can enhance network performance considerably.

Association between a client and an AP begins with an association request that is initiated by the client. This association request is normally preceded by the client's transmitting one or more probe requests on channels it selects. In each of these probe requests, the client asks for a response from all APs operating on that channel and able to receive the client. This tells the client which APs are within radio range, and the signal strengths received from the APs give an indication of which APs will be able to provide higher-quality service. Before sending an association request, a client should also have previously sent an authentication request that has been granted.

The method by which a client decides with which AP to request association is not specified in the IEEE 802.11 standard. Thus, client cards produced by different manufacturers use different algorithms for requesting association.

When an AP receives an association request, it can

either accept or deny the request. Although the IEEE 802.11 standard does not specify an algorithm for making this decision, APs (or the intelligent switch that controls them) can profitably consider the load currently carried by the AP and also the loads being carried by nearby APs. For example, an AP that is heavily loaded might not be the best one to accept a new association request. If such a request is received and the radio resource manager running on the intelligent switch knows that a lightly loaded AP is also within radio range of the requesting client, it may decide that it is best for the requested AP to deny the association request.

If, of course, the AP receiving the association request is lightly loaded and receiving a good signal from the client, there may be no reason not to accept the request. But denial of an association request by a heavily loaded AP may be best for the overall performance of the network.

In a wireless LAN controlled by intelligent switches running such a load sharing algorithm, we have a situation in which decision making is distributed between clients and APs. Clients look after their interests, and APs look after their interests. The client is likely to request association with the AP that from the client's point of view is most likely to provide the best service. On the other hand, an AP will accept association requests only if they are unlikely to degrade the quality of service provided to other clients.

As with transmit power control, load sharing techniques can be effective with the current IEEE 802.11 standard, but such techniques can be even more effective if there are appropriate changes to the standard. In this case, changes to facilitate communication regarding load sharing decisions between APs and clients would be helpful.

For example, network resources might be allocated using a method that relies on information received by the intelligent switch from a set of APs. This information would help an AP determine what service level it can provide to a client requesting association. If the AP can meet the level of service requested by the client and is the best AP for association with the client, the request is accepted. Otherwise, the AP can deny the request and advise the client with which AP it should request association.

Although we have not yet run definitive performance tests, we believe that optimized load sharing software, running on both clients and switches, can in some situations double or triple a network's throughput.

Conclusion

The architecture that has been described has the potential to improve reliability, availability, performance, and deployment effectiveness in enterprise and other large-scale wireless LANs. These improvements arise from the radio resource management algorithms contained in the software running on the intelligent switches that control APs.

In the approach that has been described, the software controlling the APs attempts to optimize performance without having any direct control over client behavior, and this limits the effectiveness of the approach. Efforts currently underway in the IEEE 802.11 Working Group promise to allow for communication between APs and clients in ways that will allow better radio resource management techniques than those described here.

Acknowledgments

The authors wish to thank Airespace, Inc. (<http://www.airespace.com>) for allowing information about their product to be included in this article.

References

- [1] IEEE 802.11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 1997.
- [2] B. P. Crow *et al.*, "IEEE 802.11 Wireless Local Area Networks," *IEEE Commun. Mag.*, vol. 35, no. 9, Sept. 1997, pp. 116–26.
- [3] B. O'Hara and A. Petrick, *The IEEE 802.11 Handbook: A Designer's Companion*, IEEE Press, 1999.
- [4] R. van Nee *et al.*, "New High-Rate Wireless LAN Standards," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 82–88.

- [5] IEEE 802.11a, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High Speed Layer in the 5 GHz Band," 1999.
- [6] 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.
- [7] A. Hills, "Large-Scale Wireless LAN Design," *IEEE Commun. Mag.*, vol. 39, no. 11, Nov. 2001, pp. 98–104.
- [8] Report and Order (FCC 03-287), In the Matter of Revision of Parts 2 and 15 of the Commission's Rules to Permit Unlicensed Infrastructure (U-NII) Devices in the 5 GHz Band, ET Docket 03-122, Washington, D.C.: Federal Communications Commission, Nov. 12, 2003.
- [9] D. Brelaz, "New Methods to Color the Vertices of a Graph," *Commun. ACM*, vol. 22, no. 4, Apr. 1979, pp. 251–56.
- [10] J. Randall-Brown, "Chromatic Scheduling and the Chromatic Number Problems," *Mgmt. Sci.*, vol. 19, no. 4, Dec. 1972, pp. 456–63.

Biographies

ALEX HILLS [SM] (ahills@cmu.edu) is Distinguished Service Professor of Engineering & Public Policy and Electrical & Computer Engineering 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.

BOB FRIDAY (bfriday@airespace.com) is chief scientist at Airespace, Inc. In this position he has led the company's radio resource management design effort. He is also a co-founder of Airespace. Previously, he was director of advanced technology and chief scientist at Metricom, Inc., where he was a key architect of Metricom's Ricochet network. Before joining Metricom, he was a member of technical staff at Watkins Johnson, where he developed software for automated test systems. He holds a B.S. in electrical engineering from Georgia Tech and an M.S. in electrical engineering from San Jose State University.