The Technology Guide Series
techguide.com

Wireless Enterprise Networking

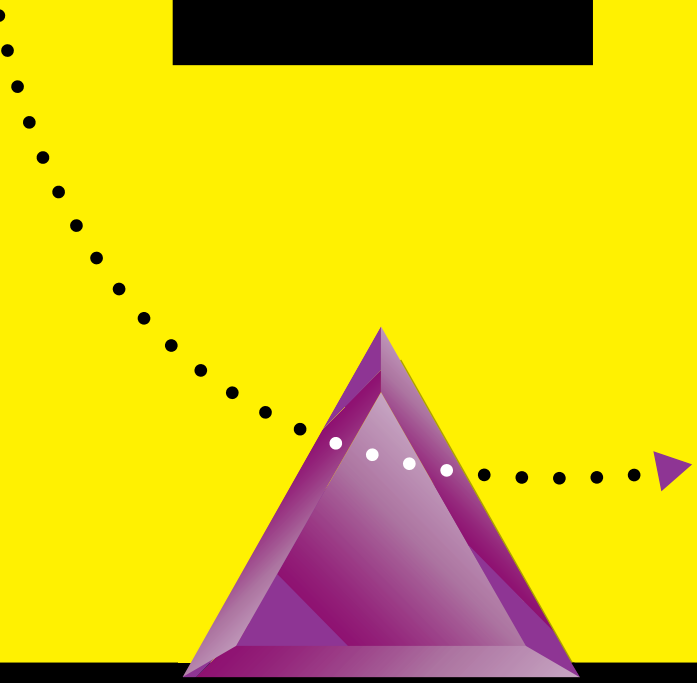

This guide has been sponsored by

BREEZECOM®
Wireless Access Solutions

Founded in 1992, *BreezeCOM, Inc.* is a leading wireless technology company that develops, manufactures and markets products for telecommunications, data communications and wireless local-area network (LAN) applications. *BreezeCOM's* high-performance products use frequency-hopping, spread-spectrum radio technology and operate in the license-free 2.4 GHz ISM (industrial, scientific and medical) band. They are designed for compliance with international regulatory, transmission and communication standards.

The company's *BreezeNET® PRO.11* line of plug-and-play wireless Ethernet products offers ease of installation, optimum performance and hardware compatibility. *BreezeNET* products have achieved worldwide recognition for high performance and ease of installation.

"Wireless network connections just don't get any easier than with BreezeNET PRO.11"

BreezeNET PRO.11 is an IEEE 802.11 interoperable, wireless network product line that is ideal for users of standard platform computers who need a flexible, plug-and-play network connection that can maintain mobility and accommodate temporary, ad hoc, building-to-building and Internet applications. Common *BreezeNET PRO.11* installations include manufacturing and warehouse facilities, hospitals, schools, construction and excavation sites, shopping malls, airports, and many others.

Table of Contents

About the Editor...

Jerry Ryan is the Vice President of Editorial Development for the Technology Guides on Communications and Networking. Mr. Ryan is also a principle at ATG. Mr. Ryan has developed and taught many courses in network analysis and design for carriers, government agencies and private industry. He has provided consulting support in the area of WAN and LAN network design, negotiation with carriers for contract pricing and services, technology acquisition, customized software development for network administration, billing and auditing of telecommunication expenses, project management, and RFP generation. He was the president and founder of Connections Telecommunications, Inc., a Massachusetts based company specializing in consulting, education and software tools which address network design and billing issues. Mr. Ryan is a member of the Networld+Interop Program Committee. He holds a B.S. degree in electrical engineering.

Introduction

In this last decade of the twentieth century, people have become increasingly dependent on and comfortable with, mobile wireless communications. The ability to carry cell phones and pagers that allow people to communicate within the local area as well as across long distances has also whetted peoples' appetites for mobile wireless data communications. This is evident in the rapidly growing use of wireless LANs to support contemporary business models. Wireless LANs (WLANs) have many important advantages over fixed wired systems. They allow rapid deployment, support for mobile personnel, use in temporary or industrial locations, and are remarkably cost effective. Wireless LANs are a natural complement to wired LANs and, in some cases, a natural replacement - particularly when deployed in enterprise-wide solutions using wireless hubs and switches that can be deployed more easily than wired solutions. Additionally, mobility lends a layer of management ease in that stations can be deployed where the data is generated. In the past, however, there has been some resistance to deploying wide-spread wireless LANs because of the lack of an industry-wide standard. But this barrier has been dramatically removed by the recent adoption of the new 802.11 wireless LAN standard. With this standard in place, network managers can now consider wireless LANs as an integral option for provisioning the corporate network enterprise. This Technology guide examines the issues, benefits, and limitations of wireless LANs within the context of the 802.11 standard.

Wireless LANs in Today's Enterprise Environment

In March of 1971, Gordon Moore and Robert Noyce, founders of a fledgling Silicon Valley start-up, announced that their Intel 4004 microprocessor was fully functional and ready for general release. That invention, the direct descendant of Shockley's transistor a decade earlier, launched a revolution that is propelling the business world forward at an ever-increasing accelerated rate. In 1996, an estimated 3.6 billion microprocessor parts were sold worldwide, almost one for every human on the face of the earth. Concomitant with this rapid development, computer networks and telephony, using microprocessors, have also developed at an accelerated rate. With the addition of video and the placing of all four technologies on the same binary infrastructure, we are now seeing the convergence of these technologies into a single, unified force, a force that is revolutionizing business and society. The general deployment of distributed computing in corporate America has changed the way we do business. The static hierarchies and massive headquarters of brick and steel have been replaced with a virtual environment of networked employees. The new business model calls for economy, shortened delivery cycles, faster access to customers, and greater emphasis on quality. All of these pressures are forcing corporations to go to where the data is generated; at the customer's site. Consider the way that the widespread deployment of laptops has radically changed the way leading-edge applications process remote data. By going to the source of the data and capturing the data in situ, rather than manually entering it on a form and scanning it into a mainframe, the new business model goals are achieved. These same advantages are gained by WLANs.

Wireless communications are the logical end point of Moore's and Noyce's invention. The shift from mainframes, to client/servers, to tetherless, laptop connections, pushes the computer closest to where the data is generated, at the customer site. In addition, the economic forces driving the corporations of the 90's are forcing constant change upon all aspects of our business. We are constantly being relocated. World-class companies are characterized by their ability to form small workgroups quickly and by their ability to rearrange those groups

quickly for new initiatives. Organizations that are tied down to physical, static, wired LANs will compete with wireless LANs (WLANs) about as well as mechanical cash registers did with microprocessor driven ones in the 1970s. In addition, there are environments where traditional cabling is not possible or is cost prohibitive. Some typical examples include hospital emergency rooms, medical offices, classrooms, conference rooms, small retail operations, open office areas, outside environments, and employees' homes. In fact, the wired environment is the anomaly here; it forces the user to go to the wire, as we were forced to go to the mainframe before PCs.

If WLANs and WWANs were as cost-effective and equivalent in performance to traditional LAN/WANs, there would be no reason to run cables. As we will see later, those conditions are becoming very close to being met. The ability to cost-effectively outfit users with tetherless laptops confers a huge strategic advantage to a company in the highly competitive marketplace of this decade. The user, armed with a laptop, can capture the data at the source. In the office environment, the user can connect to the WLAN anywhere in the building. Finally, new devices can be connected at any time or place without the costly delay of laying new cabling.

Why has this technology not been deployed until now? One reason is that older, wireless technologies were too slow. Furthermore, without uniform industry standards, older technologies were expensive, non-interoperable, and often unreliable. However, these barriers are being breached.

This guide examines the history of wireless technology and what has changed. It then examines the new wireless paradigm and shows how it is an excellent candidate to give the contemporary organization a competitive advantage over its rivals. Finally, the guide will conclude with an examination of a new 802.11 standard and with predictions of the future evolution of WLANs.

History of Wireless LANs

Wireless LANs have been around for a long time. Indeed, they predate wired LANs. ALOHA, arguably the first LAN and the basis for Ethernet, was radio-based. However, wireless evolution has lagged far behind wire-based networks. The latter evolved from 10-16 megabit speeds up to gigabit speeds for Ethernet-based LANs and ATM SONET WANs. In contrast, the most common form of wireless communications today, modem connections over analog cellular connections, provides a transmission rate of only a few Kbps. Infrared LANs, one of the alternative technologies, have also been limited by dramatic interference problems caused by sunlight and artificial light. Moreover, it is primarily a point-to-point paradigm, which means, to use infrared to construct a WLAN, a very large number of access points would need to be installed every thirty feet or so. Finally, any form of radio communication requires an allocation of the frequency spectrum, a crowded arena with many competing technologies demanding space. Likewise, many product sets have been based on proprietary non-standard technologies. Making a business case based upon a communications technology owned by a single vendor is risky in the extreme.

Emergence of the IEEE 802.11 Standard

Similarly, the WLAN industry has not provided a unified strategic product set that could interoperate between vendors. In addition, wired vendors, instead of designing their wired LAN infrastructure products to be backbone products, based their components (bridges, routers, etc.) and product sets on adapter technology. Thus, users wanting to deploy WLANs were placed in the difficult situation of either having to be dependent on a

single vendor or solving the infrastructure problems on their own. Convincing a LAN administrator used to the plethora of integrated standardized wire-based protocols from multiple vendors to consider a fragmented, non-open wireless solution, had been a difficult task. Recognizing this, the IEEE struck a new subcommittee of the 802 standardization process, the 802.11 Working Committee for Wireless LANs, to introduce a standard for the 2.4 GHz range of WLANs. The slowness of that group to produce a timely standard delayed widespread acceptance of WLANs in the early part of this decade. However, as of June 26, 1997, the final 802.11 specification was ratified and the wireless LAN industry began working on conformance and interoperability immediately.

This open standard is expected to be the catalyst for the marketplace to unify their offerings and present to their customers, a mature, and complete product set that will usher in a new paradigm shift in networking, the wireless LAN.

Goals of Wireless LANs

What are the goals of a wireless LAN? The immediate goals are to provide the customer flexibility, mobility, ease of deployment, and cost-effectiveness.

Clearly, the WLAN must use radio waves to interconnect users in a radius of several hundred meters and transmit reliable data at rates comparable to wire LANs. It must also provide kilometer-wide connectivity for campus-wide networks. Finally, it must interoperate with standard networking technologies such as Ethernet and ATM. To be successful in a general sense, WLANs must be ubiquitous, able to seamlessly integrate with wired technologies, and must be based on an open standard. In addition, they must be robust, able to transfer data at the megabit range (initially), scalable, secure, and easily maintained, as well as being constructed from an integrated product set. And, of course, they must be priced competitively with wired LANs.

In addition to all of these goals, WLANs have to compete with a LAN technology that is mature, one in which users have come to expect near-instantaneous response times while meeting most of the above goals. In the past, WLANs have not been able to meet many of these goals and have thus been rarely deployed. The only cases in which one would find them was a unit which made a local, tactical decision to acquire one for specific reasons. There have been pockets of acceptance. In the laboratory, ISM (Industrial, Scientific, and Medical) networks operate in the unlicensed 902-928 MHz band at speeds up to 100 Kbps. Satellites have provided similar rates for years. Three new satellite initiatives will shortly provide PCS (Personal Communications Services), but again, the data rates will be low (64 Kbps). Additionally, as cellular digital services become more widespread, we can expect wireless services from them as well, although, they too will be in the kilobit range initially. However, until this year, no wireless technology has provided a service set which would satisfy all of the above conditions. This is about to change.

The New Wireless Paradigm

With the acceptance of the IEEE 802.11 standard, an important necessary condition has been met; interoperability. Vendors are now producing products that are mature and affordable. Fully-featured product sets allow network managers to build customized WLANs and to integrate them into their legacy wired LANs. The technology is mature and cost-effective. Because of the openness of the standard, competition is driving prices lower. The enterprise manager can now consider wireless as a strategic option both in complementing wired LANs and as an alternative in appropriate circumstances.

Why Choose Wireless LANs?

When would one choose a WLAN and why? There are at least four situations in which wireless is an appropriate option; scenarios in which mobility, short-term usage needs, speed of deployment, and the need to overcome difficult wire installation situations are of importance.

- Mobility. The ability to access real time information while dealing with customers is enhanced with a WLAN. In hospitals, for example, health care providers can improve the quality of patient care. With a WLAN, bedside inputting of data and immediate decision-making can reduce cycle times for patient care. Likewise, the reduction of errors by handling the data once is significant.

  In an office situation, the ability to roam around the building while processing information is an advantage. Similarly, point-of-sale employees can circulate freely while serving customers. Insurance agents can input data directly in the customer's premises and receive real time on-line analytical processing. If there is business advantage in going to the customer rather than forcing the customer to come to you, the case for wireless can be compelling.

  Finally, WLANs permit mobile applications to be launched. Consider the WLAN-enabled student that can take her WLAN-connected laptop from lecture to lecture, and remain connected at all times to her files and applications. Indeed, such an environment is now in place at Carnegie-Mellon University, which has installed a campus-wide WLAN with over 100 interconnected APs.

- Short-Term Usage. Similar to the issue of mobility, short-term connectivity allows users to deploy capabilities on an as-needed basis without concern for the cost justification for wired solutions. Financial auditors, for example, can just connect for the time necessary to conduct the audit. This allows significant operational flexibility and facilitates the formation and support of ad-hoc working groups. Being able to connect to the network for a short period of time in this manner can provide a competitive advantage.

- Speed of Deployment. WLANs permit quick connectivity to the network. Forming and disbanding work groups can be done easily with WLANs. The complexity and long cycle time of moving new nodes into and out of wired LANs introduces massive on-going operational costs compared with the flexibility of wireless attachment, where the operational costs are almost zero.

- Difficult Wiring Environment. Many situations do not permit the easy installation of wires. Historic buildings or older buildings make the installation of LANs either impossible or very expensive. Trying to establish LANs in the out-of-doors is virtually impossible with legacy LANs. Consider situations in parks or athletic arenas where one wants a temporary WLAN established and removed. There are other situations where it is vital to be WLAN-enabled. Disaster recovery for example can make immediate and effective use of WLANs in the field to gather data and co-ordinate relief efforts. The use of WLANs in the battlefield is obvious. Finally, there are situations where wires cannot be laid, for example, across busy streets. Likewise, building to building connections can be facilitated where no existing underground cabling is present. Using wireless bridges to connect physically separated LANs or internet connections can be very effective.

The Business Case for Wireless LANs

The cost for wireless LAN NICs (network interface cards) are rapidly dropping in comparison with wired Ethernet. In addition, the data transfer rate, although currently lower than Fast and Gigabit Ethernet, is comparable to standard Ethernet speeds. When considering cost issues and the need for fast deployment, one must gather all of the costs associated with wire-based LAN systems, not just the cost of the node adapter (It has been estimated that the initial purchase price of a wire-based LAN is only 20% of the total cost of operating that LAN.). In many cases, the costs of managing the changes and the cycle time delays make the case for WLAN access compelling.

Some wireless LAN products, especially those with proprietary architectures, are more complex than the average wired LAN and usually, more expensive. However, there are several WLAN product lines on the market that have been deployed by end users because of lower total installation and setup costs in comparison with a completely wired scheme. The higher total installation cost for a wire system is especially true when cabling is difficult to reach or even non-existent. Additional benefits of the WLAN include maximum utilization of mobile client devices. In general, companies pay a premium for mobile or portable devices and then proceed to lose the primary benefit of such a portable device by tethering it to the network, forcing the user to become completely stationary. Because of their flexibility and potential to become even more cost-effective in the future, the WLAN market is expected to grow by 40% per year, leading to a billion dollar market in two years. As is common for IT, this, along with the 802.11 standard, will encourage competition and drive prices down dramatically.

Wireless LANs—A Technology Overview

Radio Frequency Technology Overview

A WLAN is just like a wired LAN with the cable replaced by radio spread spectrum signal. The fundamental components of a WLAN architecture are straight forward. First, we need a radio path for communications. Radio waves occupy a subset of the electromagnetic spectrum that also include radio, television, air traffic control information, and so on. In fact, the radio frequency spectrum is not only a busy place but a scarce resource. Thus, regulatory agencies have apportioned out certain frequency bands for specific purposes to prevent unnecessary collisions since the spectrum is a shared medium. These bands may either be licensed or unlicensed. Access to licensed bands is restricted to those holding licenses but access to unlicensed bands is open. WLAN bands are unlicensed. Two decades ago, the 900 MHz band was reserved for ISM use. Then, regulatory agencies allocated the 2.4-2.483 GHz band for WLAN traffic in North America. The FCC, in particular, has restricted the power output of devices operating in the 2.4 GHz band to 4 watts. However, the story is more complicated than just considering the radius limitations imposed by power restrictions. Radio waves at that frequency bounce around a lot, reflecting off hard objects. Reflections or multiple transmissions of the same signal are common and must be accounted for. There are other electronic devices that also generate signals in that range. Such signals are noise to the WLAN cell. Therefore, the designers must find a way to reduce the interference from these random noise sources and compensate for multiple reflections. Inherently, the problem of engineering a WLAN product is more complicated than designing a simple cable-based system.

Spread Spectrum Technologies

The obvious way to begin the design of a WLAN architecture would be to fix a signal at a certain frequency and use that as the "wire" of communication. However, the noise problems are so severe that an alternate method must be chosen; the so-called spread spectrum solution. Noise has its own frequency and would destroy any signal being sent through that noise cloud. However, spread spectrum, as the name implies, uses multiple frequencies in the band to increase the immunity to noise at any specific frequency. Today, two approaches are used to implement spread spectrum for WLAN transmissions, Frequency Hopping Spread Spectrum (FHSS) and Direct Sequence Spread Spectrum (DSSS). FHSS uses a large number of frequency channels, 80 with the transmitter sending a burst over one and then "hopping" to another channel.

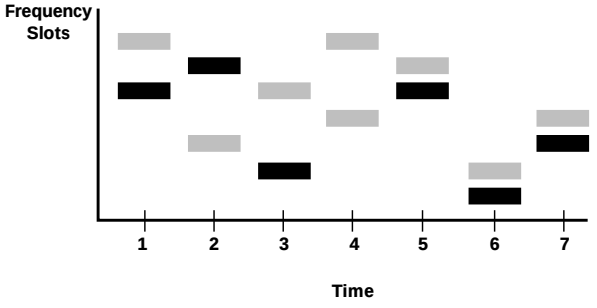

As one can see from the above Frequency Hopping Spread Spectrum, two stations are transmitting at the same time. Each picks a frequency and transmits for a specified time slot. At the end of that time slice, they both shift or hop to another frequency slot. The precise hop sequence must be known to both the sender and receiver of each channel. The typical bandwidth of the information signal is 1 MHz and the time slot interval is a tenth of a second. This diagram is an extreme simplification since the actual sequence used in the 802.11 standard would have 79 slots. Both the number of hops and the maximum "dwell" time are dictated by a regulatory agency, such as the FCC in the US. There must be at least 75 hops and the dwell time must be no more than 400 milliseconds. The hop sequence is quasi-random but is known by both the sender and receiver. There is a number of fixed sequences (26) that can be selected by an installer. If noise occurs, the worst case scenario is that a packet might be slightly corrupted on one frequency necessitating retransmission on the next hop.

An alternative spread spectrum technology is the Direct Sequence Spread Spectrum (DSSS) approach. DSSS takes a baseband signal and replaces the message with calculated blocks of fixed length codes, spreading the bandwidth by a large factor. The receiver knows the decoding sequence and retrieves the original message. DSSS is the older technology, but has gradually been replaced by FHSS. It is important to note that the two are not interoperable, although some vendors do make both product sets.

Some of the advantages of FHSS compared with DSSS are that FHSS does not require a contiguous band of frequency allocations, is simpler to implement, is cheaper to implement, is more secure, and permits multiple simultaneous transmissions. On the other hand, DSSS is easier to handoff from cell to cell and may provide a higher point-to-point data transfer rate in certain circumstances.

Today, almost all vendors have chosen the FHSS method, which is superior to DSSS in many ways. For instance,

- FHSS is more immune to signal interference. DSSS networks can be crippled by outside interference within the same frequency range because DSSS is not frequency-agile. The frequency is pre-selected and cannot avoid interference on the pre-selected band. FHSS on the other hand, hops around the noise source.

- FHSS also has a higher total aggregate capacity. The maximum number of non-overlapping Mbps DSSS channels is 3 for a total of 6 Mbps capacity. Typically FHSS systems can provide up to 15 non-overlapping 1 Mbps channels for a capacity of 15 Mbps.
- FHSS is more scalable. If a FHSS needs to handle additional cell activity, one only needs to add an additional AP in the cell, thus doubling the capacity. Since the co-located APs are naturally non-frequency overlapping, they interfere with each other very little.
- FHSS systems have some physical advantages. FHSS units are typically lighter than DSSS units. DSSS units also require more power to operate.

CSMA/CA

As we shall see, the 802.11 standard defines a MAC layer interface that is compatible with wired Ethernet. However, instead of using Ethernet's CSMA/CD (Carrier Sense Multiple Access/Collision Detection), it uses a variant called CSMA/CA (Collision Avoidance). The CD protocol would require that the wireless radios be able to send and receive at the same time, which would increase the product price and complexity. Also, with wireless stations, it is not always the case that all stations can be in a position to hear all of the other stations. To minimize the possibility of stations not being able to hear each other, 802.11 defines the notion of a Virtual Carrier Sense. The transmitting station first sends a very short packet called the Request to Send (RTS) packet which contains the source and destination station addresses plus an indication of the duration of the intended message. If the medium is free, the receiver will reply with a Clear-to-Send (CTS) packet. On receipt of that, the planned transmission commences. The use of the CSMA/CA protocol reduces the chances of collisions. If the medium is busy, the transmitters will perform an exponential backoff similar to their wired counterparts. Thus, 802.11 uses a protocol called positive acknowledge. When a station wants to transmit, it first checks the medium to see if it is free, as with a wired Ethernet. If it is free, the station transmits. The receiving station, after receiving the message and making sure that the message has not been corrupted, sends back an acknowledgment. If the sending station does not receive an acknowledgment, it assumes that the original message did not make it and retransmits it.

WLAN Architectures
Ethernet Connections

About 98% of the time, the WLAN interfaces with wired Ethernet LANs. In this environment, the Access Point (AP) acts as a bridge into the LAN, permitting WLAN devices to share the same LAN resources as the wired LAN stations. Small or home office (SOHO) environments, on the other hand, are usually supported by a wireless system only , but can still use an access point for peer-to-peer networks. A client-to-client wireless LAN can be configured but has very limited use.

Cells and Access Points

The area covered by a single wireless LAN is called the cell. Stations comprising the WLAN itself are located within the cell. All communications inside and outside the cell must be coordinated by a single unit called the Access Point (AP). The AP connects the cell with other cells and with wired LANs. The AP must also synchronize all of the stations within the cell so that they perform the frequency "hopping" at the proper time and frequency.

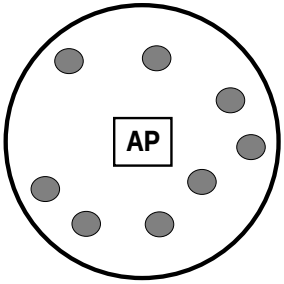

The above diagram shows a basic cell. Within the cell, each station can "hear" all of the others. In a more practical situation, not all of the nodes can "hear" all of the others. In this case, the cell must have a master controller, called the Access Point. The AP must be able to hear all of the nodes and coordinate all intra- and inter-cell traffic.

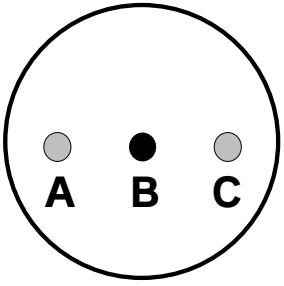

Hidden Nodes

Suppose that A wants to send a message to B. A listens to be sure there is no carrier, hears nothing and, assuming that all is clear, transmits. At the same time, C, which cannot hear A, also wants to send to B. It makes the same decision as A but both packets collide at B. This is the hidden node problem in WLANs. If this happens, retransmissions will lower the effective throughput rate. Rather than do that, the IEEE 802.11 standard inserts a slight complication in the protocol. It forces A to send a Request packet first, indicating when and how much data it would like to send. Node B, if it wants to connect, responds with a Clear to Send, which station C also hears, causing C to defer its transmission according to the information it hears about A. Thus, collisions are reduced significantly. The addition makes the protocol more complicated and introduces extra delay. However, most designers chose it for the better overall performance and as we shall see, it is an option in the 802.11 protocol.

Cells can be connected in four fundamental ways: as stand-alone cells, as linked cells within multi-cell configurations, as an element in a wired Ethernet environment, or to a remote LAN by a wireless bridge.

Stand-alone Cells

A stand-alone cell consists of the AP and all associated wireless stations. The maximum number of stations depends on the nature of the data traffic. If there were a lot of data moving around, 15 might be an upper limit while a less-intensive environment might support up to 50. In such a scenario, no wires need be laid at all. The cell diameter depends on many factors, but one can expect up to 200 meters inside buildings and about a kilometer in outside environments.

Linked Cells

When the maximum number of cell members is exceeded or the cell diameter is exceeded, more cells can (and must) be created. These cells must overlap to permit seamless cross-cell communications. When a user walks from one cell coverage to another, the AP's must "hand-off" the signal from one cell to the next without the user being aware that she has crossed cell boundaries. Such cells are called "linked" cells but the hand-off is so smooth that the user is unaware that she has crossed into a new cell. Combined with a backbone connection, a complete campus could be "wired," as Carnegie-Mellon is today.

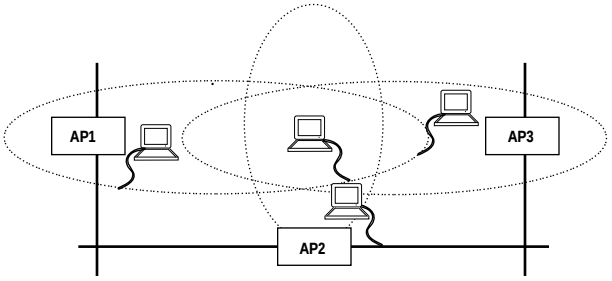

Multi-Cells

If several cells overlap the same physical area, the configuration is termed a "multi-cell". With clever algorithms, the APs can decide which of them is best positioned to orchestrate the WLAN communication from the sender to the receiver. When heavy traffic is expected, this is a distinct advantage. Such coverage, called multi-cell coverage, not only load-levels the inter-cell traffic but also provides redundancy to ensure reliable fail-safe operation.

Roaming and Handoff

Roaming is the capability of portable stations to move freely between overlapping cells, either in a linked topology or a multi-cell one. Roaming is "seamless" in that the user experiences no noticeable interruption when moving from cell to cell. APs must "hand-off" the signal and synchronize appropriately. Some key differences between vendors may be seen in how they handle roaming and AP synchronization. These have not been standardized and the user would be wise to ask prospective vendors how this is done.

Roaming and fast handoff is important. As an acid test, one might ask the vendor how fast a vehicle can move through the multi-cell environment with an attached laptop. Clearly, if the WLAN is functioning well, one might expect the WLAN to handle speeds in the range of many kilometers per hour. Similarly, one might ask how the stations communicate with different APs in order to assess which AP they should currently belong to. This is a dynamic decision depending on the location of the station and the traffic patterns at that instant.

A related issue is the question of how load leveling is accomplished in adjacent WLAN cells. If there is significant traffic within overlapping cells, it would be better for the WLAN to reallocate some stations to less loaded APs and, to initiate a backoff procedure to ensure that the station is not "bounced" from AP to AP.

Remote Bridges

Finally, cells can be connected to remote WLANs by means of wireless bridges. Although not part of the 802.11 standard, such bridges can span kilometers at high data rates and can turn WLANs into WMANs if necessary.

PC Connectivity

Laptop and desktop PCs are connected to the WLAN through a NIC card, such as a PCMCIA card, or an ISA card. In this implementation, the Ethernet NIC is replaced with a wireless NIC which is installed into these same I/O slots inside the client device. In addition, all the necessary drivers must be installed and hardware interrupts configured, just like a wired NIC. Another approach to building a wireless client device is to use an external wireless transceiver which simply plugs into the wired Ethernet NIC or port on the client device and provides instant wireless access, without the need for additional drivers or other software. These devices can have a single or multiple RJ-45 connections. This approach provides the flexibility to connect virtually any Ethernet device to the WLAN.

Sources of Interference

Designers also try to minimize the sources of interference which, through signal corruption cause retransmissions and, as a result, reduce throughput and performance. Three major sources are multipath propagation, microwave devices such as cooking ovens, and ISM network interference.

As has been noted, radio waves bounce around a lot causing reflections or multiple instances of the same signal. Specially designed antennae are used to reduce the effect of these sources. The modem, if properly built, will select the strongest signal on a frame-by-frame basis. The FHSS technique also spreads the signal over 79 hops using bands of 1 MHz each. The hops are changed 8 to 30 times a second in a pre-defined order. The sending and receiving stations must use the same hopping sequence with the same synchronization of timed shifts. Thus, in the case of interference, only a single channel will suffer jamming and for a very short period of time. Another advantage of this technique is that it provides excellent security from hostile, naive listeners. This essentially makes WLANs as secure as wired-based Ethernets. Microwave ovens are a minor nuisance at close distances in that they generate interference in the 2.4 GHz range. Finally, scientific and medical equipment may be using periodic low power ISM transmissions which may occasionally be interpreted as noise by collocated WLANS.

Other Considerations

The WLAN architecture, in addition to being 802.11 standard, should permit open connections to standard LAN interfaces, such as 802.3 Ethernet. The WLAN should transparently interwork with common network protocols such as IP, IPX, AppleTalk, Netbuei, DECnet, and so on. Standard management protocols should be followed, such as SNMP. A good vendor will support standard MIBS like MIBII and bridge MIB and also provide a private MIB for their units.

Finally, most wireless LANs will have laptops, hand-helds, or palm-tops as their user stations, as well as desktop PCs, high powered 68000 based workstations, and assorted peripheral devices. The portable devices such as laptops and hand-helds, are, by definition, battery-driven. Thus, a key component of the IEEE 802.11 standard was to define a stand-by mode in which stations could render themselves comatose and hibernate, awaiting a wake-up call from the AP. Such strategies will significantly extend the operational life of the batteries and the usefulness of the station to users.

IEEE 802.11: The New Wireless LAN Standard

As has been noted, the agreement of the 6th draft version of the IEEE 802.11 standard was a huge leap forward for the WLAN industry. Growth was stunted in the early nineties because vendors could not decide on an open standard. Customers were very leery, and rightly so, of embracing proprietary standards that might be obsolete in months. Staying away in droves, the market penetration stagnated until this year. However, the adoption of the new standard has changed all that. Let us take a look at the pertinent parts of the standard.

IEEE 802.11 Basic Architectural Model

The 802.11 defines a cellular network. The basic cell is called the BSS (Basic Service Set). Each BSS will contain many stations (STAs) with a head station called the Access Point (AP). If the AP connects with another 802 network (such as an Ethernet), it is called a portal. The connecting 802 LAN is called a Distribution

Service (DS). Collections of BSSs and DSs form a whole called an Extended Service System (ESS). Although the standard does not require it, the AP and the portal are normally on the same physical device.

IEEE 802.11 Layer Description.

The 802.11 MAC Layer specification fits in under the 802.2 upper part of the Data Link Layer of the standard IEEE model. Underneath that, three physical layer specifications are defined:

- FH Frequency Hopping Spread Spectrum DS Direct Sequence Spread Spectrum and IR infrared. The FHSS and DSSS are assumed to run at 1 or 2 Mbps nominal. In addition to the normal MAC layer functionality, the 802.11 MAC layer also provides for fragmentation, packet reassembly and acknowledgments. It also defines two different access methods: Distributed Coordination Function, and Point Coordination Function.

- Basic Access Method. The basic access method, called the Distributed Coordination Function, is similar to Ethernet. Using a modified CSMA/CD (Carrier Sense Multiple Access/Collision Detect) technique, it performs the CSMA, listening to sense if the medium is free. If it is, it transmits, if not, it defers. However, there is no CD (collision detect), and instead we use Collision-Avoidance protocol and thus the MAC layer protocol is called a CSMA/CA protocol. There are two reasons why a CSMA/CD approach was not used: it would require that the station sender/receivers be full-duplex which would drive up the cost significantly, and, it cannot be assumed that all stations can hear all of the other stations at all times (as compared with a wire-based Ethernet). Just because a station is ready to transmit and senses that the medium is clear around itself, does not mean that the medium is clear around the receiver.

In addition, the standard defines a term called the Virtual Carrier Sense. In order to reduce the chance that a station has become "hidden" (and thus cannot hear), the sender will begin a transmission by sending off a very short message, Request To Send (RTS) to the receiver, indicating the address and an approximate duration time. The receiver will respond with a Clear To Send (CTS) if appropriate. All stations seeing this will set their NAV (Network Allocation Vector) indicating that they have seen the "use" of the virtual carrier. Such negotiation will reduce the chances of collisions and the probability of a station not hearing properly. The standard defines a RTS Threshold parameter so that small messages will not experience extra overhead. Finally, if the sender sees that the medium is busy, it backs off, as described later on in this document. In terms of the timing sequences, the basic protocol looks like this:

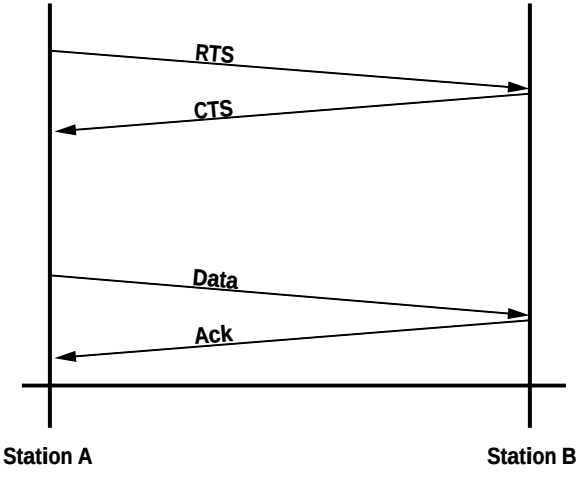

Fragmentation and Reassembly

Ethernet packets can vary in size up to 1518 bytes and indeed, to maximize its effective transfer rate, one ought to use packets that are as large as possible. This is not always the case with 802.11. There are good reasons for making WLAN packet sizes smaller. They include the inherently higher Bit Error Rate of radio transmission, the smaller costs of retransmissions of smaller packets, and, in FHSS, the frequency is typically hopped every 100 milliseconds. Also, microwave ovens have a 4ms noise and 4ms clear duty cycle which is about the size of the full Ethernet packet, making it highly likely that collisions will occur, forcing a retransmission. Thus, the protocol enforces a fragment frame, send fragment, and wait approach. Therefore, part of the MAC protocol is to split the frame (the MSDU) into several fragments (the MPDUs) and send the fragments until the entire message has been acknowledged.

Inter Framing Spacing

The basic frame formats defined by the standard look like this:

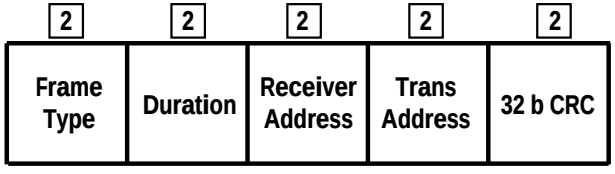

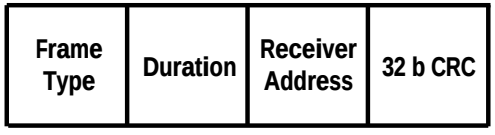

The precise details of the field sizes and contents are explained in the standard. We briefly recall the important elements to give the reader a flavor of the MAC layer. The frame elements include:

- last fragment,
- retry (frame is a retransmission),
- elements present; frame not empty,
- duration in microseconds (for hidden nodes),
- address fields (all 6 bytes),
- source address,
- destination address,
- AP address,
- transmitting station address,
- receiving station address,
- sequence control,
- dialog control,
- fragment number,
- frame body (0 to 2304 bytes long), and
- CRC error control.

There are four inter-frame times specified to separate out the message interactions. They are:

1. SIFP (Short Inter Frame Spacing). This is the maximum time that the sender has to turn itself around when expecting a reply (for example, from frag to ack or from RTS to CTS). For the FS Physical layer, the value is 28 microseconds.
2. PIFS (Point Coordination Inter Frame Spacing). This is the time used by the Access Point to gain access to the medium before any other station. It is defined to be a SIFP plus a slot time or about 78 microseconds.
3. DIFS (Distributed Inter Frame Spacing). This is the time a station waits when it wants to initiate a conversation. It is a PIFS plus a slot time or 128 microseconds.
4. EIFS (Extended Inter Frame Spacing). This is the time a station must wait if it has not understood a message defining a time before sending something out. Otherwise, the station would likely collide with incoming packets.

Exponential Back Off

The Backoff method used in 802.11 is similar to the Ethernet's. Using the definition of a slot time (the time in which a station can tell if the medium is busy or about 64 microseconds), each station picks a random number of slots from one to a maximum, and waits that long to retry. If the medium is still busy, it will increase the maximum and retry (hence the use of the word "exponential"). In 802.11, the exponential back off is executed when:

- the station has sensed that the medium is busy as it prepares to transmit,
- after each retransmission, and
- after each successful transmission.

The only time that a station will not exponentially back off is when the medium has been free for more than a DIFS and it has more to send.

How a Station Connects

A station needs to join a BSS when first powering up. It can do this in two ways:

- Passive scanning. The station waits for a Beacon Frame message from the AP. This is sent out periodically to check for new and included synchronizing information.
- Active scanning. The station sends a Probe Request Frame which is a request for the AP to acknowledge its existence. It then awaits a Probe Response Frame. The station then hops to the next channel and tries again.

Synchronization Process

Synchronizing for hopping is crucial as are other timing signals that are needed. The AP does this by sending out a Beacon Frame which contains a time signal from the AP. Note that this is the actual time that the frame is transmitted and NOT the time when it is queued, which could be significantly different because of the collision algorithm. The receiving stations then reset their clocks according to the new value.

The Authentication Process

Either of the two above methods is acceptable in joining the cell. After the initial connection, the AP and the station enter into the Authentication Process where they exchange passwords, accounting information, and the like.

The Association Process

When the station has been authenticated, the AP begins the Association Process. The basic capabilities of each are exchanged and registered. The characteristics of the station are recorded and exchanged with the AP. Only after this process has been completed can the station transmit and receive data.

Roaming

Roaming is the process of moving from one BSS (cell) to another and having the handoff accomplished smoothly. It is similar to roaming in cell phones, except that roaming on a WLAN is done on packet boundaries. How this is to be done is not defined in the standard but left up to the vendors to decide in their proprietary offerings. Consequently, care must be taken with multiple vendor solutions due to the possibility of interoperability problems. Another issue related to the different implementations is the speed of roaming handoff. Some implementations are faster than others, which also contributes to problems of interoperability.

- Security because the WLAN is wireless. Security is of extreme importance. The committee defined the term WEP for Wired Equivalent Protection to stress that a WLAN has to be as secure as a LAN. Two types of security are considered.

- Restricting Access. The idea here is to provide something equivalent to the physical key to the LAN. It is assumed that access areas will be protected with a similar mechanism.

Eavesdropping

Each message is encrypted with a standard pseudo-random number generator (PRNG) based algorithm that uses the RSA RC4 algorithm. That method is considered reasonably strong as it would take a very determined effort to crack it. The fact that each message contains a new Initializing Vector that generates a new PRNG makes it all the more difficult to crack. It would be much easier to copy the key to the lab.

Power Saving

WLANs stress the use of portable devices. Portables use batteries and battery power is a very scarce resource. Therefore, the standard has gone to considerable lengths to define a Power Save mode. Basically, it permits a station to “go to sleep”. The AP keeps track of any messages, buffering them until the station requests a wake up.

Summary

The IEEE 802.11 standard thus defines an 802 compatible MAC layer which can interoperate with the other 802 technologies. Three physical standards are defined; the FHSS, DSSS and Infrared. At the moment, the focus of the standard is on the 2.4 GHz band but this will likely shift to the higher bands as they are released for open use. Essentially, the 802.11 standard provides open, asynchronous networking that requires a distributed control function. The supported data rates are 1 or 2 Mbps; however, wireless LAN vendors who are creative and have extensive radio and digital signal processing expertise can produce products with a data rate as high as 3 Mbps, and by slowing down to talk to 802.11 devices, can still be 802.11 compliant.

Summary and Conclusions

With the agreement of the IEEE 802.11 standard, there is a new standard in place to drive the implementation of WLANs. We can expect many vendors to bring 802.11 compliant products to market in increasing numbers and dropping costs. Moreover, vendors will be supplying complete architectural solutions that will permit the building of wireless LANs and also integrate existing wired LANs with the new components. Using standard management protocols such as SNMP, users will be able to seamlessly integrate their wireless components with their legacy systems thus protecting their old LANs while leveraging advantage from the new capabilities of the emerging WLAN products. As we noted in the beginning of this pamphlet, the original transformation of computing power has finally completed its evolution from mainframe, customer-unfriendly environments to customer-friendly situations which put our networking and computing power where it should be; with the customer.

In conclusion, when one considers the total cost of ownership, including the on-going cost of maintenance and the costs of missed opportunities, the overall business case for WLANs becomes not just compelling but insistent. The network enterprise manager needs to understand both the Technology and the business case for wireless LANs.

Case Study

BreezeNET at the Super Bowl

It was a groundbreaking day both on and off the field in San Diego for Super Bowl XXXII, but it was what took place off the field that figures to have the most significant effect on how news is gathered and distributed to the public in the future.

For years, media giants such as the Associated Press (AP) have used runners, and more recently fiber optics to transport photographs and other important information from the field to its subscribers around the world during major events like the Super Bowl. This method proved to be extremely costly and required permanent installation of fiber optic cable at various sites.

The Associated Press, the world's pre-eminent news gathering organization, identified the problem and began looking for a way to improve the old method of transmitting photographs from an event's location to the world's newspapers. During the World Series, the AP decided to experiment with wireless networking as a solution. They purchased wireless products from five different wireless LAN manufacturers and tested them for range, speed and ease of use during the World Series. BreezeCOM®, a San Diego based manufacturer, proved to be the winner in all three categories. As a result of these tests the AP decided to go completely wireless for Super Bowl XXXII in San Diego. For the Super Bowl, the Associated Press decided to use BreezeNET PRO® wireless network products from BreezeCOM to efficiently transmit photographs of the game worldwide for substantially less money than conventional methods, such as installing fiber optic cable throughout the stadium.

AP technicians used BreezeNET to wirelessly network six fixed digital cameras and several roving photographers with digital cameras to a series of Macintosh Powerbooks. As the cameras captured the

action from the game, the digital photographs were stored on PCMCIA hard drives. Once the cameras' hard drives were full, they were inserted into one of six field-level Macintosh G-3 PowerBooks. Each PowerBook was outfitted with a BreezeNET SA-10 Station Adapter connected to its Ethernet port.

The captured digital images stored in the Powerbooks were remotely accessed from AP's mobile newsroom located in a trailer outside the stadium approximately 1,500 ft. away. These images were transmitted to the mobile newsroom via six BreezeNET model "D" access points featuring detachable antennas. Access points' antennas were pointed up to six repeaters in the stadium's light standards, and then down to the six PowerBooks in the field. From the AP's newsroom, the images were processed and file transferred using FTP to an AIX server in New York via a router located at AP's Technical Center in Cranbury, NJ. The thousands of photographs captured during the Super Bowl were archived in one central location allowing AP's subscribers to access the vast amount of images worldwide.

"This technology allows us to cover a one-day event with tight photo deadlines without the need for fiber optic installation and maintenance," said Howard Gros, Associated Press Project Manager. "The BreezeCOM system is also re-useable which allows us to bring it into similar events in the future with no additional cost."

Incorporating the BreezeCOM wireless method of transmitting photographs proved to be a major success for the Associated Press. Increasing overall efficiency, the media giant was able to transmit photographs from the cameras to the trailer outside the stadium in 25 seconds. AP also transmitted the photograph of the first play of the game to subscribers via satellite within five minutes of the kickoff. The BreezeCOM wireless method of data transfer was so successful that the Associated Press immediately transported the system to Nagano, Japan to facilitate the organization's coverage of the Winter Olympics.

Glossary of Terms

2.4 GHz range—The frequency spectrum assigned by the FCC to wireless LAN systems.

802.11—An IEEE standard for the 2.4 GHz range of WLANs.

Access Point—Central point in a wireless cell which acts as a bridge for traffic to and from wireless stations in the cell.

Active scanning—The station sends a Probe Request Frame which is a request for the AP to acknowledge its existence.

Antenna Diversity—An antenna arrangement allows multiple antennas to transmit and receive amplitude and phase weighted signals.

Association—The service used to establish access points/station mapping and enable STA invocation of the distribution system services.

Authentication—The service used to establish the identity of one station as a member of the set of stations authorized to associate with another station.

Basic Access Method—The method used by wireless stations to access the AP.

Basic Service Set (BSS)—A set of stations controlled by a signal coordination function.

Bridge—A device that connects and passes packets between two network segments. Bridges operate at Layer 2 of the OSI reference model (the data-link layer) and are insensitive to upper-layer protocols. A bridge will examine all frames arriving on its ports and will filter, forward, or flood a frame depending on the frame's Layer 2 destination address.

Cell—In WLANs, the local area in which a particular transmitter/receiver operates.

Carrier Sense, Multiple Access/Collision Avoidance (CSMA/CA)—A protocol to avoid traffic conflicts in a shared radio system.

Distributed Inter Frame Spacing (DIFS)—The time a station waits when it wants to initiate a message.

Direct Sequence Spread Spectrum (DSSS)—An alternative spread spectrum approach in which the transmitter replaces a baseband signal with calculated blocks of fixed length codes.

Distributed Coordination Function (DCF)—A class of coordination functions where the same coordination function logic is active in every station in the basic service set whenever the network is in operation.

Distribution Service (DS)—The connecting 802 LAN is called a Distribution Service (DS).

Extended Inter Frame Spacing (EIFS)—This is the time a station must wait if it has not understood a message defining a time before sending something out via Point Coordination Inter Frame Spacing (PIFS).

Electromagnetic Spectrum—The full frequency range of electromagnetic emissions from visible light through radio waves.

Encryption—Applying a specific algorithm to data so as to alter the data's appearance and prevent other devices from reading the information. Decryption applies the algorithm in reverse to restore the data to its original form.

Enterprise Network—A complete business network consisting of functions, divisions, or other components used to accomplish specific objectives and defined goals.

Exponential Back Off—When a collision occurs, the method of calculation of the time before retrial.

Extended Service Area—The AP's overlapping radio coverage area of contiguous cells.

FCC—Federal Communications Commission.

Fragmentation—The process of breaking a packet into smaller units when transmitting over a network medium that cannot support the original size of the packet.

Frequency Hopping Spread Spectrum (FHSS)—A radio system operating over a large number of frequency channels in which the transmitter sending a burst over one and then "hopping" to another channel.

Infrared LANs—Wireless LANs based on directed or diffused single-beam or multi-beam infrared transmissions and diverse signal antennae.

Interoperability—A term that implies that different vendor products of the same technology can successfully operate with each other.

ISM Networks—Industrial, Scientific, and Medical networks which operate locally in the unlicensed bands.

Linked Cells—Contiguous cells interconnected to provide seamless interoperability.

Media Access Control (MAC Layer)—The second layer of the protocol stack.

Media Access Control Layer Address (MAC Layer Address)—Also called hardware address or physical address. A data-link layer address associated with a particular network device. Contrasts with network or protocol address which is a network layer address.

Multi-cells—A set of wireless cells that overlap the same area.

Multipath Propagation—When the transmitted signal arrives at the receiver from different directions with different path lengths, the signals are attenuated and delayed differently.

Network Allocation Vector (NAV)—A part of the WLAN protocol indicating that the station has seen the use of a virtual carrier.

Passive scanning—The station waits for a periodic Beacon Frame message from the AP.

Point Coordination Inter Frame Spacing (PIFS)—The time used by the Access Point to gain access to the medium before any other station.

Point-to-Point Multipoint Configuration—A point-to-multipoint wireless bridge configuration is used when connecting network nodes or remote networks back to a central network hub.

Point-to-Point Configuration—A point-to-point wireless bridge configuration is used to connect two remote nodes or networks to each other.

Portal—The AP that connects with another 802 network (such as an Ethernet).

Radio Frequency—The range in which radio signals are transmitted.

Roaming—(1) The capability of portable stations to move freely between overlapping radio cells. (2) The function of moving the wireless station from cell to cell.

Scalable—A term implying that a particular, referenced technology can be expanded in terms of capacity or performance.

Short Inter Frame Spacing (SIFP)—The maximum time that the sender has to turn itself around when expecting a reply.

Small or Home Office (SOHO)—A term for small offices or home based operations.

Spread Spectrum—A radio transmission system that uses multiple frequencies within the assigned band to increase the immunity to noise at any specific frequency.

Stand-alone Cells—A wireless cell with a single AP.

Wired Equivalent Privacy (WEP)—The optional cryptographic confidentiality algorithm specified by IEEE 802.11 used to provide data confidentiality that is subjectively equivalent to the confidentiality of a wired Local Area Network medium that does not employ cryptographic techniques to enhance privacy.

Wireless LANs (WLANs)—Wireless Local Area Networks.

WMANs—Wireless Metropolitan Area Networks.

WWAN—Wireless Wide Area Networks, for example, cellular telephones.

NOTES

NOTES

NOTES

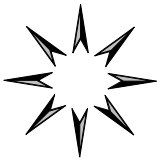

BreezeNET PRO.II Advantages

- ✔ **Data-Rate:** Industry's fastest @ 3 Mbps, auto fall-back to 2 or 1 Mbps
- ✔ **Range:** 2000' open space, 200' - 600' within offices
- ✔ **Seamless Roaming:** Fastest @ 60 Mph, no lost or duplicated packets
- ✔ **Easiest Installation:** Attach to ANY Ethernet device, no drivers or software
  - ✻ *Supports IP, IPX, NFS, NetBEUI, AppleTalk, EtherTalk, LAT, Etc.*
- ✔ **Security:** Hopping, authentication, can use "off-the-shelf" encryption
- ✔ **Management:** SNMP, SLIP, plus TFTP download of upgrades
- ✔ **Standards:** IEEE 802.11, 802.3, 802.1d: Secure investment, interoperable
- ✔ **22 Channels:** Overlapping cells provide for 15 Mbps ++ aggregate throughput

BREEZECOM
Wireless Access Solutions

2195 Faraday Ave., Suite A, Carlsbad, CA 92008
(760) 431-9880 • www.breezecom.com

This Technology Guide is one in a comprehensive series of Guides that provide objective information and practical guidance on technologies related to Communications & Networking, the Internet, Document Management, Data Warehousing, and Enterprise Solutions. Our team of technical editors writes each Technology Guide to assist IT and business professionals in making informed decisions about all aspects of technology application development and strategic deployment.

techguide.com is supported by a consortium of leading technology providers. BreezeCOM has lent its support to produce this Guide.

Visit our Web Site at www.techguide.com to view and print this Guide, as well as all of our other Technology Guides. This is available as a free service.