1. Wireless LAN Part 1

A wireless local area network (LAN) is a flexible data communications system implemented as an extension to, or as an alternative for, a wired LAN. Using radio frequency (RF) technology, wireless LANs transmit and receive data over the air, minimizing the need for wired connections. Thus, wireless LANs combine data connectivity with user mobility.

Wireless LANs have gained strong popularity in a number of vertical markets, including the health-care, retail, manufacturing, warehousing, and academia. These industries have profited from the productivity gains of using hand-held terminals and notebook computers to transmit real-time information to centralized hosts for processing. Today wireless LANs are becoming more widely recognized as a general-purpose connectivity alternative for a broad range of business customers.

Wireless LANs can be deployed to transmit data, voice and video within individual buildings, across campuses, and over metropolitan areas. Some of the computer and communications industries' leading vendors are introducing personal digital assistants (PDAs), modems, wireless microprocessors and other devices and applications in support of wireless communications.

2. Advantages of Wireless LANs

With wireless LANs, users can access shared information without looking for a place to plug in, and network managers can set up or augment networks without installing or moving wires. Wireless LANs offer the following productivity, convenience, and cost advantages over traditional wired networks:

- Mobility - Wireless LAN systems can provide LAN users with access to real-time information anywhere in their organization. This mobility supports productivity and service opportunities not possible with wired networks
- Installation speed and simplicity: Installing a wireless LAN system can be fast and easy and can eliminate the need to pull cable through walls and ceilings
- Installation flexibility - Wireless technology allows the network to go where wire cannot go
- Reduced cost-of-ownership: While the initial investment required for wireless LAN hardware can be higher than the cost of wired LAN hardware, overall installation expenses and life-cycle costs can be significantly lower. Long-term cost benefits are greatest in dynamic environments requiring frequent moves and changes.
- Scalability - Wireless LAN systems can be configured in a variety of topologies to meet the needs of specific applications and installations. Configurations are easily changed and range from peer-to-peer networks suitable for a small number of users to full infrastructure networks of thousands of users that enable roaming over a broad area

3. *Issues to Highlight*

1999 Business Research Group, a market research firm, predicts a sixfold expansion of the worldwide wireless LAN market by the year 2000, reaching more than $2 billion in revenues. Now, 2003 we know that it expanded more than 100 times since 1999.

4. How Wireless LANs Work

In a typical wireless LAN configuration, a transmitter/receiver (transceiver) device, called an access point, connects to the wired network from a fixed location using standard cabling. At a minimum, the access point receives, buffers, and transmits data between the wireless LAN and the wired network infrastructure. A single access point can support a small group of users and can function within a range of less than one hundred to several hundred feet. The access point (or the antenna attached to the access point) is usually mounted high but may be mounted essentially anywhere that is practical as long as the desired radio coverage is obtained.

End users access the wireless LAN through wireless-LAN adapters, which are implemented as PC cards in notebook or palmtop computers, as cards in desktop computers, or integrated within hand-held computers. wireless LAN adapters provide an interface between the client network operating system (NOS) and the airwaves via an antenna. The nature of the wireless connection is transparent to the NOS.

5. Types of Wireless LANs

Wireless LANs usually have two types of realisation:

- Ad-hoc - Several mobile nodes (for example, notebook computers) may get together in a small area (for example, in a conference room) and establish peer-to-peer communications among themselves without the help of any infrastructure such as wired/wireless backbone. Since a small coverage area does not imply insured communication, there is no real reliability. Despite the possibility of ad-hoc networking, most applications will require communications with services located in a pre-existing infrastructure
- Infrastructured - Such an infrastructure is typically a higher-speed wired (or wireless) backbone. Therefore we can divide typical network traffic into two directions: uplink (into the backbone) and downlink (from the backbone). The contact points to the backbone are called access points. The access points can be either base stations for wired infrastuctures or wireless bridges for wireless infrastructures. Repeaters may also be used for enlarging the coverage area of communication

6. Downlink Traffic

Due to the limited bandwidth of wireless LANs, a common channel is typically used for communication between an access point and mobile nodes. Downlink is achieved by broadcasting on this common channel. More precisely, the access point broadcasts packets to all mobile nodes even if there is only one destination. Downlink activity may constitute up to 75 or 80 percent of the total traffic in wireless LANs because that nodes on modern LANs often operate in a client-server mode. For instance there might be a high performance workstation or PC acting as a file server. A request for file transfer on the uplink may result in a huge file on the downlink.

7. Uplink Traffic

The uplink protocol is the core task for the MAC design of wireless LANs. To recognize and register new mobile nodes that join the network in any time and place, a kind of random access protocol is needed.

Thus uplink traffic needs a multiple access protocol to organize the transmissions from mobile nodes. In the next section 'Expected features of wireless LANs' it is explained why multiple access is more difficult for wireless LANs than for wired LANs.

8. Modulation

This is the process of varying some characteristic of the electrical carrier signal (generated in the modem) as the information to be transmitted (digital signal from the DTE) on that carrier wave varies. The following modulation methods are the most commonly used in modems:

- Frequency shift keying (FSK)
- Differential phase shift keying (DPSK)
- Quadrature amplitude modulation (QAM)
- QAM/Trellis Coding Modulation (QAM/TCM)
- Pulse code modulation (PCM)

For modulations background, look at ADSL modulation section.

9. Features of Wireless LANs

Multiple access is not easy in the wireless environment because of the following reasons:

- Dynamic physical channel characteristics
- Practical implementation
- Mobility and network topology
- Spatial behaviour and handoff

10. Dynamic Physical Channel Characteristics

Wireless LANs typically operate in very strong multipath fading channels, for example, the received signal may suddenly disapear or reappear. Also capture effects may occur when two mobiles transmit at the same frequency it may be that the receiver clearly receives the signal of one of them without detecting any interference). The channel statistics may change significantly within 10 to 20 ms duration or any movement of 1 foot distance!

11. Practical Implementation

Many functions that are trivial to implement in a wired medium cannot be applied to a wireless medium. For example, carrier sensing in cable is easy but carrier sensing in radio takes at least 30 to 50 microseconds. Moreover, a mobile station can't detect collision while transmitting because the difference between the strengths of the signals.

12. Mobility and Network Topology

The network must maintain normal operation while its topology changes with time.

13. Spatial Behaviour and Handoff

Infrastructure LANs are based on some access points that divide the service area of a wireless LANs into different corresponding cells. One of the primary reasons to adopt cellular structure is in order to increase the effective total bandwidth by using different frequencies in

different cells. This concept, known as frequency reuse, is illustrated in the following example. The figure shows a seven-cell structure; suppose a total of 3-B bandwidth is needed to serve users in the seven-cells area. Three different frequency bands can cover this seven-cell region. If frequency reuse were not employed and a single frequency band served all users in the same region, a total of 7-B bandwidth would be needed to support the same quality of service.

As a result of frequency reuse, the total available communication bandwidth for all users is much larger than the transmission speed. Furthermore, frequency reuse not only saves the spectrum but also reduces transmission power by reducing cell size.A function that allows a mobile node to communicate with the access point in a cell and then switch to the access point in another cell is called handoff or handover. The purpose of the handoff is to keep continuous or seamless service to mobile nodes through different cell coverages. Handoff is consequently a special feature to deal with the mobility issue for wireless networks.

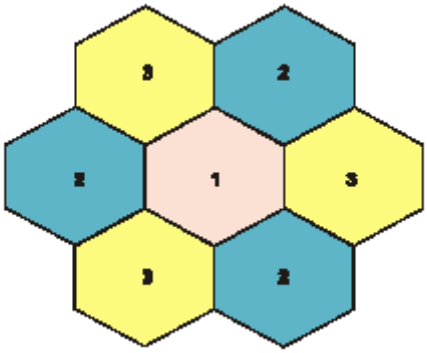

3 Frequencies in classic hexagon formation

14. Three main types of technological implementation:

- Narrowband - A narrowband radio system transmits and receives user information on a specific radio frequency. Narrowband radio keeps the radio signal frequency as narrow as possible just to pass the information. Undesirable crosstalk between communications channels is avoided by carefully coordinating different users on different channel frequencies

- Spread spectrum - Most wireless LAN systems use spread-spectrum technology, a wideband radio frequency technique developed by the military for use in reliable, secure, transmission-critical communications systems. Spread-spectrum is designed to trade off bandwidth efficiency for reliability, integrity, and security. In other words, more bandwidth is consumed than in the case of narrowband transmission, but the tradeoff produces a signal that is, in effect, louder and thus easier to detect, provided that the receiver knows the parameters of the spread-spectrum signal being broadcast. If a receiver is not tuned to the right frequency, a spread-spectrum signal looks like background noise. There are two types of spread spectrum radio: frequency hopping and direct sequence

- Infrared (IR) - IR systems use very high frequencies, just below visible light in the electromagnetic spectrum, to carry data. Like light, IR cannot penetrate opaque objects; it is either directed (line-of-sight) or diffuse technology. Inexpensive directed systems provide very limited range (3 ft) and typically are used for personal area networks but occasionally are used in specific wireless LAN applications. High performance directed IR

is impractical for mobile users and is therefore used only to implement fixed sub-networks. Diffuse (or reflective) IR wireless LAN systems do not require line-of-sight, but cells are limited to individual rooms

The new up comer for higher bandwidth and dataspeeds plus better noise handling than the classical WLAN spread spectrum is OFDM, we will take a look on OFDM little later.

The most popular standards are:

- 802.11 –family
  IEEE 802.11 networks work on license free industrial, science, medicine (ISM) bands
  - IEEE 802.11 (1997) 1 Mbps and 2 Mbps (2.4 GHz band)
  - IEEE 802.11b (1999) 11 Mbps (2.4 GHz band) = Wi-Fi
  - IEEE 802.11a (1999) 6, 9, 12, 18, 24, 36, 48, 54 Mbps (5 GHz band)
  - IEEE 802.11g (2001 ... 2003) up to 54 Mbps (2.4 GHz) backward compatible to 802.11b
- OFDM IEEE 802.11a
- pro.11 -3 Mbit/s (2,4 and 3,5GHz band)
- HiperLAN (5GHz v1=24, v2=54Mbit/s)
- Bluetooth (2,4GHz 740 kbps)
- radioLAN permanent 10Mbit/s (5,8GHz)

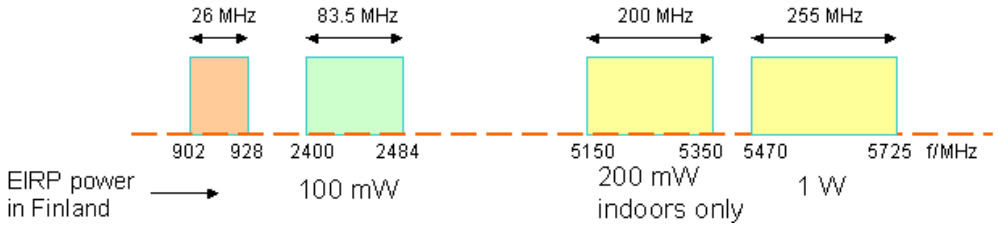

Free frequencies used worldwide.

Frequency	Notes	Standards
2.400-2.483.5 GHz	ISM Band (USA max 4W EIRP, other 100mW)	802.11/11b
902-928 MHz	ISM Band (Used by GSM in most countries)	
5.800-5.925 GHz	ISM Band	
5.15-5.25 GHz	U-NII (Unlicensed - National Information Infrastructure) max. 200 mw EIRP	802.11a
5.25-5.35 GHz	U-NII max. 1w EIRP	802.11a
5.725-5.825 GHz	U-NII USA max. 4w EIRP other 200mW	802.11a

In Japan is EIRP max 10mW ~10 meters isotropic coverage practically
In USA you are allowed to use 4W, all other countries 100mW/200mW and 1W EIRP.

The 3,5GHz band is licensed and is under heavy discussion, some countries are not yet set, Sweden is one of them. In Denmark you have to pay for transported data volume, in Finland it is just to apply for your bands and in Sweden you must show economical solidity and deponate about 100,000SKr. In Germany and other places in Europe it work in similar ways as in Scandinavia. To be sure, you have to ask local authorities, it can even exist local restrictions in country regions.

3,5GHz bands for OFDM:
2 * 21MHz
2 * 28 MHz (The best).
2 * 26,25 MHz
2* 19,25 MHz
2 * 16,75MHz
Classes:

3,5a1	3400MHz-3450MHz common
3,5a	3500-3550MHz (2 *28MHz)
3,5b	3450MHz-3500MHz not common

To overcome this bottleneck you use directional antennas of different kinds. (This is actually juridical a problem in some countries, since directed power from antenna can be very high, imagine a parabolic dish with 38dB gain.) We will look more on network planning.

All those standards act as layer 1 and 2 network component only. They all encapsulate classical LAN protocols such as Ethernet standards and in fact follows CSMA/CD structures more or less. This have some impact on time critical data from higher layers which has to be dealt with before entering the world of wireless, for this QoS is included in the 803.11 family. The systems are more or less transparent to upper layer and lan protocols.

WIRELESS LAN STANDARDS

Standard	Status	What it defines
IEEE 802.11	Completed 1997	Original WLAN standard. Supports 1 Mbps to 2 Mbps.
IEEE 802.11a	Completed 1999	High-speed WLAN standard for 5-GHz band. Supports 54 Mbps.
IEEE 802.11b	Completed 1999	Current dominant WLAN standard. Supports 11 Mbps.
HiperLAN2	Completed 2000	Competing high-speed WLAN standard for 5-GHz band. Supports 54 Mbps.46
IEEE 802.1x	Completed 2001	Comprehensive security framework for all IEEE networks, including Ethernet and wireless.
IEEE 802.11g	Expected 2001	Alternate high-speed WLAN standard for 2.4-GHz band. Supports 20+ Mbps.
IEEE 802.11i	Expected 2001	Wireless-specific security functions that operate in combination with IEEE 802.1x.
IEEE 802.11e	Expected 2001	QoS mechanisms that support all the IEEE WLAN radio interfaces.
IEEE 802.11f	In process	Defines communication between access points.
IEEE 802.11h	In process	Defines spectrum-management techniques for 802.11a.
WISPR	Expected 2001	Wireless ISP Roaming. Recommendations by the Wireless Ethernet Compatibility Alliance on how to support roaming across multiple public WLAN networks.

All existing and proposed standards

The market has also overwhelmingly accepted one wireless standard: IEEE 802.11b. At 11 Mbps, IEEE 802.11b provides sufficient speed for most applications, even though actual throughput is only about 6 Mbps(theoretical 11Mbit/s), and a busy 802.11b network degrades much faster than wired Ethernet because of a less efficient medium-access protocol. IEEE 802.11b is making serious inroads to the home environment as well, so the fate of the home-oriented HomeRF (Home Radio

Frequency) specification has become quite uncertain, especially with one of HomeRF's major initial backers, Intel, defecting to IEEE 802.11b.

However, you should watch standards development most closely. IEEE 802.11b launched the industry, but widespread usage has exposed security flaws that are addressed only by vendor-specific solutions. Keeping track of these developments and designing a network with which you can easily migrate to improved technology is the crux of WLAN deployment today.

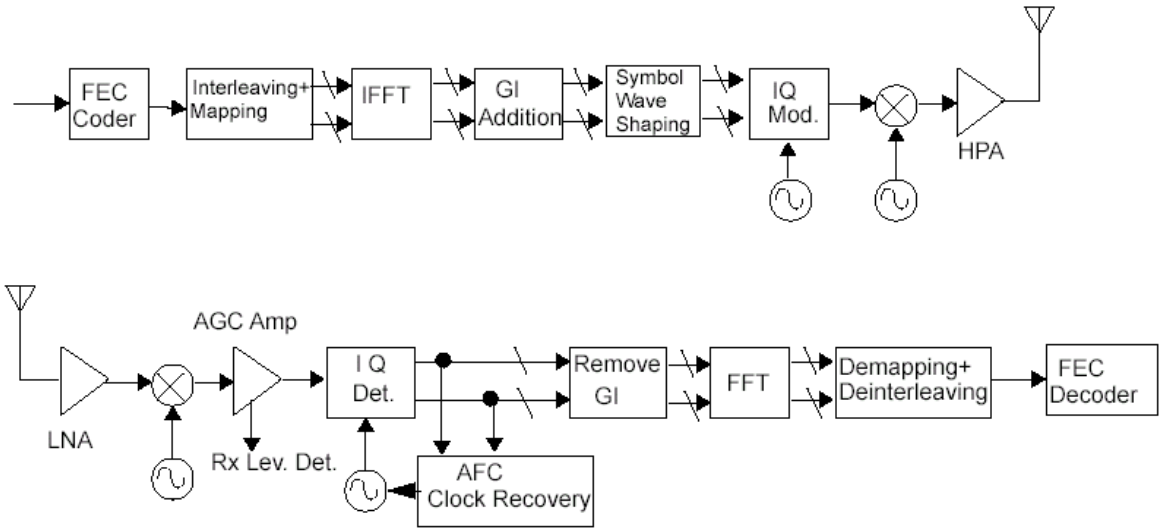

WLAN transiever

15. Technology and Standards Developments

Vendors and standards groups are advancing WLAN technology on three broad fronts: higher speeds, improved security and QoS. In an ideal world, one new standard would encompass these improvements. When a vendor's products support these improvements, you could just upgrade its equipment, and everything would be backward compatible. But this world does not exist, and advancements will occur in stages.

With respect to speed, there are exciting new developments. The IEEE 802.11a standard (which was started before the IEEE 802.11b standard) specifies a new physical layer that runs at a raw data rate of 54 Mbps. Although maximum user throughput is likely to be 25 Mbps to 30 Mbps, this is still a fivefold increase over IEEE 802.11b--almost like going to Fast Ethernet from conventional Ethernet.

IEEE 802.11a uses an advanced radio technique called OFDM (Orthogonal Frequency Division Multiplexing). Instead of sending data bits sequentially at a very high data rate, OFDM sends multiple data streams in parallel over separate radio carrier signals. This results in a more robust radio signal that makes high bandwidth

communications practical. In fact, many next-generation wireless systems, including fixed and mobile wide-area systems, are based on OFDM.

In addition, the radio can dynamically employ different modulation methods based on the quality and strength of the radio signal, resulting in extremely high throughput at shorter ranges and lower but reliable communications at higher ranges. And whereas IEEE 802.11b uses the increasingly congested 2.4-GHz radio band, IEEE 802.11a operates in the less congested 5-GHz unlicensed band, which has more than three times the available spectrum (300 MHz vs. 80 MHz). However, there is no long-term protection against interference in the 5-GHz band either.

Atheros Communications has been aggressively developing and promoting the benefits of 802.11a technology. Atheros shipped chipsets this summer, and we expect a raft of WLAN products using these chips to appear by year's end. With aggressive pricing on these chipsets, building an 802.11a product should cost no more than making an 802.11b device. So why not just wait for 802.11a?

The answer is complex. First, there is the question of range. The laws of physics dictate that the range of free-space radio communications decreases with higher frequencies, but indoor propagation differs from free space because of absorption and reflections. Moreover, power transmit levels and the type of modulation used also affect range. The result is that it is hard to accurately predict in advance the range of any particular radio technology.

According to Mobilian Corp., a manufacturer of both IEEE 802.11b and IEEE 802.11a components, up to four times as many access points are needed to cover an area with 802.11a than an area with 802.11b. However, recent "real-world" testing by Atheros in office environments indicates otherwise.

Atheros claims that, as long as you place access points in close proximity, about 60 to 80 feet from one another, you can readily overlay an 802.11a network on an 802.11b network. For the full 54-Mbps speed of 802.11a, range is restricted to about 50 feet; at 100 feet, throughput drops to 36 Mbps; and at 200 feet, 6 Mbps. Keep in mind that actual user throughput is about half of these link rates.

Although throughput drops off with range, according to Atheros and other vendors, it remains higher with 802.11a than with 802.11b. However, until 802.11a products are available and more testing is done and publicized, laying an 802.11a network over an 802.11b network will remain a complicated issue and will likely not be just a matter of swapping a radio card in a dual-slot access point. Fortunately, being able to power access points using their Ethernet connections does ease the redeployment burden.

There is another issue, though: backward compatibility. While 802.11a and 802.11b employ different radio bands, many initial network cards will support only 802.11a. Dual-mode cards will also become available but will cost more for some time because separate chips are required. With 802.11b so widely entrenched, initial 802.11a deployments will constitute small islands of coverage, making the upgrade hard to justify for many users.

Entrenched 802.11b vendors also are not rushing out with 802.11a products, and many of the initial 802.11a vendors are secondary players looking to gain market footing. Still, higher speeds are inevitable, for the increased bandwidth support not only offers higher throughput but supports a larger number of users, something that will quickly become an issue as the popularity of the technology increases.

IEEE 802.11a is not the only high-speed option, either. The European Telecommunications Standards Institute, or ETSI, has developed a family of high-speed wireless standards, with HiperLAN2 a direct competitor to 802.11a. HiperLAN2 uses the same physical layer as 802.11a, including OFDM and operation in the 5-GHz band, but it differs at upper layers. Whereas 802.11a is based on CSMA (carrier sense multiple access), HiperLAN2 centrally coordinates access, dynamically assigning time slots to individual mobile stations. This deterministic approach (analogous to token ring) is more complicated but provides for QoS--currently missing in 802.11a--and makes HiperLAN2 a more seamless extension of ATM networks.

For IP-based applications, however, the two standards offer comparable capabilities. So will we have to live with two standards? Perhaps, but IEEE 802.11a has greater momentum, with more companies developing components and with end-user products closer on the horizon. And, as we'll see in a moment, QoS is coming to 802.11 networks as well. Another factor is regulations: European regulations governing interference management favor HiperLAN, but standards work under way by the IEEE (802.11h) will address this as well.

To complicate matters further, the IEEE is developing another high-speed standard, 802.11g, which has a peak rate of more than 20 Mbps. This standard will likely use OFDM. Although not directly backward compatible with 802.11b, 802.11g does operate in the same radio band as 802.11b, and vendors will be able to offer cards that support 802.11b and 802.11g, possibly simplifying network upgrades. But if 802.11a products start rolling out, 802.11g could be too little, too late.

What is not yet clear is what vendors will do to facilitate the upgrading of access points to higher speeds. Those with modular radios (such as PC Card format) will be easier to upgrade than those with integrated radios. For example, dual-slot access points from Enterasys Networks and Intermec
Technologies Corp. will support 802.11a and 802.11b simultaneously, though the potential difference in range remains an issue. An alternative approach will be to lay an 802.11a (or 802.11g) network over an 802.11b network and have the two operate independently. This may be simpler but won't be the most efficient tactic in terms of infrastructure. If you expect to consider this approach in the future, make sure you run two Ethernet ports to each access-point location today.

16. IEEE 802.11

The standards define the physical (PHY), logical link (LLC) and media access control (MAC) layers for a wireless local area network.

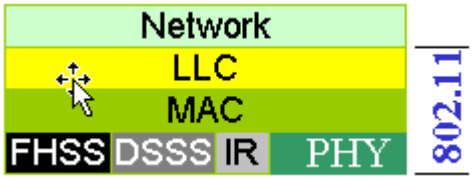

802.11 networks can work as:

- basic service set (BSS)

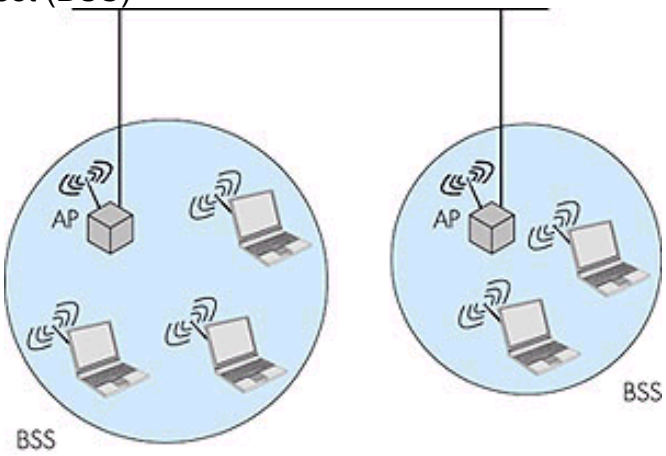

- extended service set (ESS)

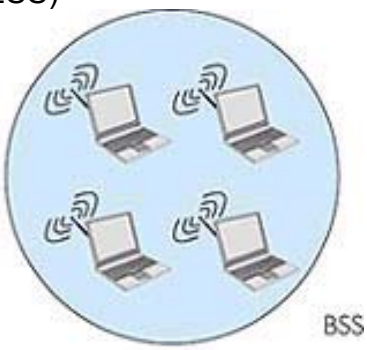

BSS can also be used in ad-hoc networking

The MAC layer provides access to wireless medium with CSMA/CA
Priority based access (802.12) is also supported
Used when joining the network
As well as authentication & privacy with and access point in BSS or ESS

Three physical layers (PHY) variances:
FHSS: Frequency Hopping Spread Spectrum (SS)
DSSS: Direct Sequence SS
IR: Infrared transmission

DSSS transmitter spreads out the original signal over a wider frequency spectra and the receiver recovers the original signal with a correlator, this is done mathematically. This is good for protection against noise which is relative narrowband signals. When you restore the original signal the noise is spread out over the spectra and fall below

acceptable S/N factor. The method is so effective that it can coexist with other narrowband transmissions. The con is that you need much more bandwidth when transmitting data than regular narrowband transmission. One bit is for example “0” is encoded with “chips” wich actually are a binary sequence. All nodes and accesspoint must use the same encoding mechanism within the same “WLAN” Several can overlap eachother with different encoding mechanisms. Modulation used is BPSK and QPSK modulation. This method is very robust against noice.

FHSS transmitter uses several small fixed carrier within the spectra declared. A key is used to jump between those carriers, known as jump key sequence. All nodes and cooperation access points must then use the same key, Several “WLAN’s” can coexist if using different keys, thus they will not collide. A small burstof information is transmitted at each carrier as the transmitter jumps the sequence, the burst time is called dwell time. Modulation used is called GFSK (Gaussian Frequency Shift Keying). FHSS is not so good in multipath, but easy to implement.

OFDM is yet another technology, this makes use of several carriers sending in parallell. Modulation can vary, DMT is popular. OFDM is good for multipath reception.

Multipath means that the signal transmitted has bounced in the terrain, this makes up problems with some technologies FHSS, devasting for the dwell time bursts.

[ADD HERE PART TWO into the deep with 802.11 and WLAN technologies]

17. **The major WLAN threats?**

1. Most current products use spread spectrum technology. Vendors initially claimed it was difficult or impossible to de-spread or demodulate the signals. Wrong, Gemmel says. It's easy.

   All you have to do is steal an SSID (Service Set Identifier), the ID attached to packets sent over WLANs that functions as a password for joining a network. All radios and access points within a network use the same SSID. Packets with other SSIDs are ignored.

2. Vendors also said you couldn't get an SSID unless you were given it. Wrong again. "We now know SSIDs are sent in the clear," Gemmel says. "You can get very simple software, some of it free on the Internet, that easily intercepts somebody's SSID."
3. WLAN signals are prone to being intercepted well outside the facility in which the network resides.

   "A lot of consumers are using wireless LANs now," he points out. "They see on the box that it's 11 Mbps up to 300 feet. They're not educated enough to realize, though, that the signal doesn't necessarily stop at 300 feet. In fact it can go up to 2,000 feet and beyond."
   This makes it easy for eavesdroppers to drive up to an office building - or home - park and infiltrate a network inside without anyone realizing.

4. As everyone who knows anything knows by now - or should do - the 802.11b Wired Equivalent Protocol (WEP) encryption can be compromised by hackers using statistical mathematical analysis tools. Two recent studies, one from AT&T another at Rice University (www.rice.edu) have made this painfully clear, Gemmel says.
5. At the level of what hackers can do once they smash through inadequate WLAN defenses, Gemmel puts "file transposition" at the top of his list. Infiltrators steal an SSID, gain access to a network, hack passwords on the enterprise LAN and then merrily delete or alter files stored on servers - or steal trade secrets contained in files.
6. Or hackers infiltrate the network and leave behind "Easter eggs," hidden and undocumented programs or messages embedded in the code of commercial software residing on the network. Some Easter eggs are harmless, even funny, but they can also be destructive viruses.
7. The last WLAN security threat is really only a perceived threat, he says, because hackers would need a lot of hardware and arcane software to do it. But theoretically, they could intercept WLAN packets, decrypt them if they're encrypted using WEP, change them, re-encrypt them and send them on to the intended recipient - who would never know.

18. Stay Safe

The 802.1X standard is designed to enhance the security of wireless local area networks (WLANs) that follow the IEEE 802.11 standard. 802.1X provides an authentication framework for wireless LANs, allowing a user to be authenticated by a

central authority. The actual algorithm that is used to determine whether a user is authentic is left open and multiple algorithms are possible.
802.1X uses an existing protocol, the Extensible Authentication Protocol (EAP, RFC 2284), that works on Ethernet, token ring, or wireless LANs, for message exchange during the authentication process.

In a wireless LAN with 802.1X, a user (known as the *supplicant*) requests access to an access point (known as the *authenticator*). The access point forces the user (actually, the user's client software) into an unauthorized state that allows the client to send only an EAP start message. The access point returns an EAP message requesting the user's identity. The client returns the identity, which is then forwarded by the access point to the *authentication server*, which uses an algorithm to authenticate the user and then returns an accept or reject message back to the access point. Assuming an accept was received, the access point changes the client's state to authorized and normal traffic can now take place.
The authentication server may use the Remote Authentication Dial-In User Service (RADIUS), although 802.1X does not specify it.

19. WEP
Although speed gets everybody's attention, it is actually new security features that may bring us greater peace of mind. The current IEEE 802.11 security method, called WEP (Wired Equivalent Privacy), employs either 40-bit or 128-bit encryption using the RC4 algorithm. Unfortunately, WEP has serious security holes and relies on manual key distribution.

To address these shortcomings, the IEEE is developing a new security architecture, specified by IEEE 802.1x, that can be applied to all IEEE access networks, including wireless (at any speed) and wired networks. This architecture provides a framework for authentication, encryption, message integrity and key distribution, and is designed to work in conjunction with existing security standards, such as EAP (Extensible Authentication Protocol) and RADIUS (Remote Access Dial-in User Service).
Another new standard, IEEE 802.11i, specifies how security is specifically implemented in wireless networks, including 802.11b and 802.11a. With solid backing by key players, such as Cisco Systems and Microsoft, and standards close to completion, expect products to start supporting these new security standards as early as next year.

Microsoft Windows XP, for example, supports 802.1x and EAP. One result: A single user logon can be used for both the wireless and the infrastructure networks. Taking advantage of these new wireless security features will mean more integration work, but this is far better than the current approach of no security at all. Of course, these security standards are only now approaching completion; it may be some time before vendors support them, and there is the big question of interoperability.

The final major push is QoS, with yet another standard, IEEE 802.11e. This standard provides for both asynchronous data traffic and data traffic that is time controlled, such as voice or video. It also allows each traffic stream to employ different policies. For example, a video stream that is time sensitive could employ forward error correction instead of packet retransmission. IEEE 802.11e--for QoS--in conjunction with IEEE 802.11a--for speed--will match HiperLAN2's capabilities.

QoS is an essential capability for voice and video support, but these mechanisms will need to be integrated with QoS mechanisms in infrastructure networks at large, and this will take some time. So while exciting, it may be years before applications in corporate environments can truly take advantage of this capability. Home use of integrated voice/video/data networks will happen much faster. However, there is no reason to wait for these more exotic features: Today's products offer more than sufficient capabilities for many applications. And as long as you put some hard questions to your vendors about their upgrade paths, you can safely deploy a network that you can enhance as needed over time.

20. Understanding 802.1x authentication

IEEE 802.1x is a draft standard for port-based network access control, which provides authenticated network access to 802.11 wireless networks and to wired Ethernet networks. Port-based network access control uses the physical characteristics of a switched local area network (LAN) infrastructure to authenticate devices that are attached to a LAN port and to prevent access to that port in cases where the authentication process fails.

During a port-based network access control interaction, a LAN port adopts one of two roles: *authenticator* or *supplicant*. In the role of authenticator, a LAN port enforces authentication before it allows user access to the services that can be accessed through that port. In the role of supplicant, a LAN port requests access to the services that can be accessed through the authenticator's port. An *authentication server*, which can either be a separate entity or co-located with the authenticator, checks the supplicant's credentials on behalf of the authenticator. The authentication server then responds to the authenticator, indicating whether the supplicant is authorized to access the authenticator's services.

The authenticator’s port-based network access control defines two logical access points to the LAN, through one physical LAN port. The first logical access point, the *uncontrolled port*, allows data exchange between the authenticator and other computers on the LAN, regardless of the computer's authorization state. The second logical access point, the *controlled port*, allows data exchange between an authenticated LAN user and the authenticator.

IEEE 802.1x uses standard security protocols, such as RADIUS, to provide centralized user identification, authentication, dynamic key management, and accounting.

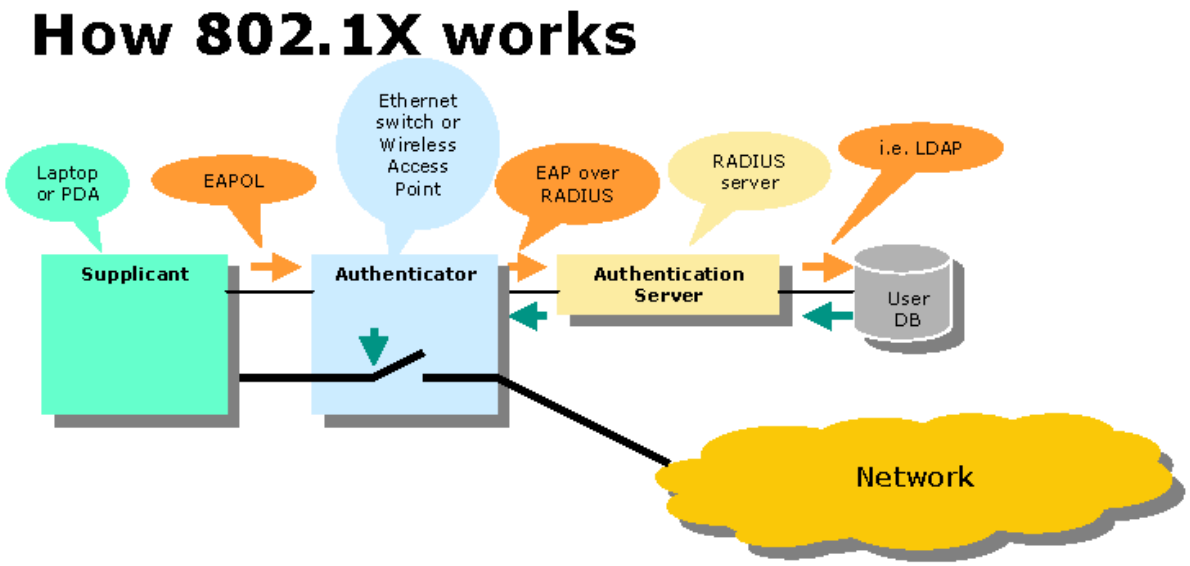

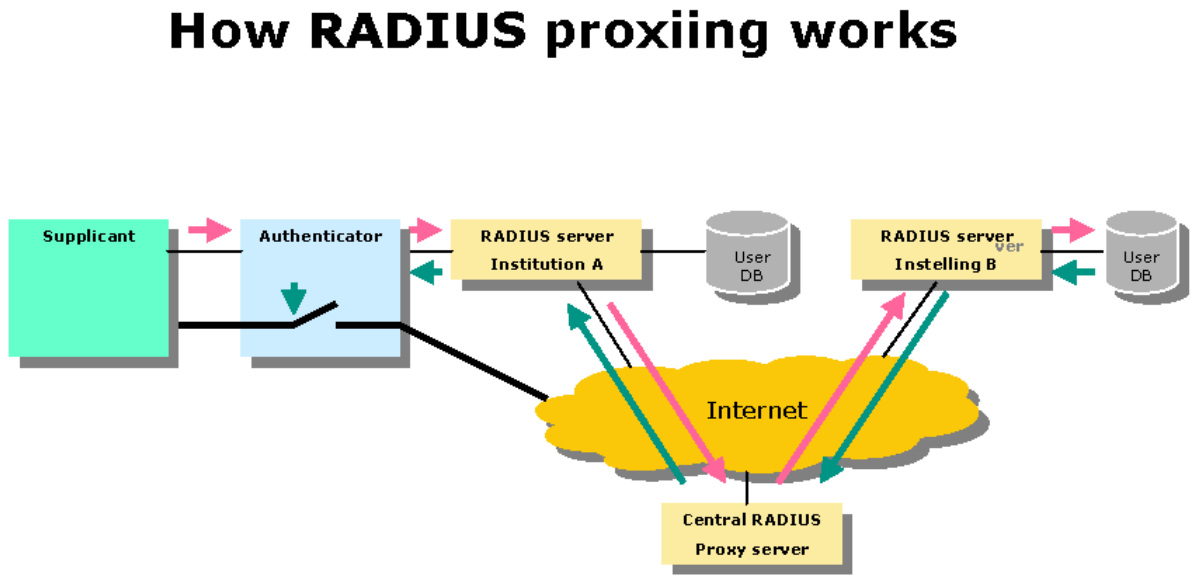

Understanding what the IEEE 802.1x standard is and why you should care means understanding three separate concepts: PPP, EAP and 802.1x itself.
Most people are familiar with PPP - Point-to-Point Protocol. PPP is most commonly used for dial-up Internet access. PPP is also used by some ISPs for DSL and cable modem authentication, in the form of PPP over Ethernet. PPP is part of Layer 2 Tunneling Protocol, a core part of Microsoft's secure remote access solution for Windows 2000 and beyond.

PPP evolved beyond its original use as a dial-up access method and iis now used all over the Internet. One piece of PPP defines an authentication mechanism. With dial-up Internet access, that's the username and password you're used to using. PPP authentication is used to identify the user at the other end of the PPP line before giving them access.

Most enterprises want to do more for security than simply employing usernames and passwords for access, so a new authentication protocol, called the Extensible Authentication Protocol (EAP), was designed. EAP sits inside of PPP's authentication

protocol and provides a generalized framework for several different authentication methods. EAP is supposed to head off proprietary authentication systems and let everything from passwords to challenge-response tokens and public-key infrastructure certificates all work smoothly.

With a standardized EAP, interoperability and compatibility of authentication methods becomes simpler. For example, when you dial a remote-access server and use EAP as part of your PPP connection, the RAS doesn't need to know any of the details about your authentication system. Only you and the authentication server have to be coordinated. By supporting EAP authentication a RAS server gets out of the business of acting as middle man, and just packages and repackages EAP packets to hand off to a RADIUS server that will do the actual authentication.

This brings us to the IEEE 802.1x standard, which is simply a standard for passing EAP over a wired or wireless LAN. With 802.1x, you package EAP messages in Ethernet frames and don't use PPP. It's authentication and nothing more. That's desirable in situations in which the rest of PPP isn't needed, where you're using protocols other than TCP/IP, or where the overhead and complexity of using PPP is undesirable.

802.1x uses three terms that you need to know. The user or client that wants to be authenticated is called a supplicant. The actual server doing the authentication, typically a RADIUS server, is called the authentication server. And the device in between, such as a wireless access point, is called the authenticator. One of the key points of 802.1x is that the authenticator can be simple and dumb - all of the brains have to be in the supplicant and the authentication server. This makes 802.1x ideal for wireless access points, which are typically small and have little memory and processing power.

The protocol in 802.1x is called EAP encapsulation over LANs (EAPOL). It is currently defined for Ethernet-like LANs including 802.11 wireless, as well as token ring LANs such as FDDI. EAPOL is not particularly sophisticated. There are a number of modes of operation, but the most common case would look something like this:

The authenticator sends an "EAP-Request/Identity" packet to the supplicant as soon as it detects that the link is active (e.g., the supplicant system has associated with the access point).

1. The supplicant sends an "EAP-Response/Identity" packet to the authenticator, which is then passed on to the authentication (RADIUS) server.
2. The authentication server sends back a challenge to the authenticator, such as with a token password system. The authenticator unpacks this from IP and repackages it into EAPOL and sends it to the supplicant. Different authentication methods will vary this message and the total number of messages. EAP supports client-only authentication and strong mutual authentication. Only strong mutual authentication is considered appropriate for the wireless case.
3. The supplicant responds to the challenge via the authenticator and passes the response onto the authentication server.

4. If the supplicant provides proper identity, the authentication server responds with a success message, which is then passed onto the supplicant. The authenticator now allows access to the LAN- - possibly restricted based on attributes that came back from the authentication server. For example, the authenticator might switch the supplicant to a particular virtual LAN or install a set of firewall rules.

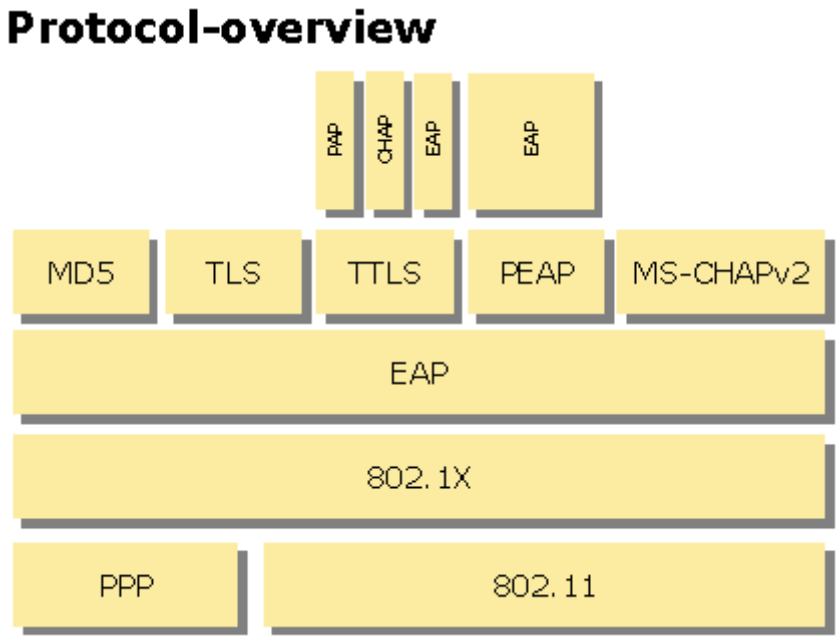

22. Improving WLAN Security
Over the past year, much has been written about the vulnerabilities of 802.11b wireless LANs. Researchers from AT&T Labs, UC Berkeley, Intel [.zip], and University of Maryland have identified holes in Wired Equivalent Privacy (WEP) that let attackers learn the keys used to encrypt 802.11b traffic.
Tools like NetStumbler exploit 802.11b behavior, sniffing the airwaves to discover cards, access points, and the peer-to-peer or infrastructure networks in which they participate. AirSnort and WEPCrack even use captured traffic to recover crypto keys. Today, anyone armed with one of these shareware tools, a wireless card, antenna, and GPS is capable of "war driving".

23. First, acknowledge the problem
802.11b vulnerability assessment products are finding opportunity in WEP's misfortune. One company, Cigital, offers assessment services that survey 802.11b access points, identifying correctable configuration weaknesses that range from default Service Set IDs (SSIDs) to risk factors for ARP cache poisoning [.pdf].
NetStumbler and AirSnort are also handy for self-assessment. By roaming around your building or campus, you may discover underground WLANs that you didn't know about. For more systematic, ongoing introspection, consider commercial products like the ISS Internet Scanner and RealSecure IDS, recently enhanced to spot and monitor 802.11b wireless-borne attacks.

24. Next, make the best of WEP
War drivers report that just 30 to 40 percent of discovered WLANs now use WEP. For heaven's sake, enable WEP and change your keys frequently! Consider using 802.11b products with dynamic key generation, like Agere's ORiNOCO AS-2000 or NextComm's R7210. Configure long, hard-to-guess SSIDs. Apply MAC filters or use VLANs to restrict access to authorized cards. Track inventory to make sure those cards stay in employee hands, and please block MACs that belong to lost or stolen cards. Lock down access point management interfaces, just as you would on any perimeter router or firewall. Use anti-virus and personal firewall software to keep the wireless client clean, preventing back-channels.

By combining firewall defense with IPsec, SSH, or SSL, you can better prevent wireless eavesdropping and block access by unauthenticated clients. For example, many companies have already deployed a SafeNet or Ashley-Laurent VPN client on laptops for secure remote access. The same client can often tunnel IPsec over wireless to a VPN gateway located between the access point and the rest of the corporate network. Alternatively, consider an access point with built-in IPsec, available from vendors like Colubris Networks.

When roaming, wireless cards often use DHCP to obtain a new IP from each access point. This can be a problem for network layer solutions like IPsec. If roaming is essential to your 802.11b deployment, consider wireless "VPN" solutions from companies like NetMotion, Columbitech, or Ecutel. These products use servers that run proprietary, WTLS, or Mobile IP protocols to avoid session interruption when a wireless client changes its address. They also offer user-level authentication, which may or may not be present in your IPsec VPN today.

25. There must be a better WEP
Windows XP and 802.11b gear now using 802.1x authentication and key distribution still use WEP for payload encryption. To "fix" the vulnerabilities inherent in WEP, the IEEE is actually defining a brand new encapsulation protocol. This new protocol is expected to use a stronger cipher the Advanced Encryption Standard (AES) in Offset Codebook (OCB) mode. We can hope that it will replace WEP next year, providing industrial-strength data integrity and privacy for 802.11 wireless

26. For Windows XP, consider using 802.1x
802.11b Open System Authentication is no authentication at all. The alternative, Shared Key Authentication, depends on secrecy of the shared WEP key which can be disclosed or cracked. If your wireless clients happen to run Windows XP, a stronger alternative is available: IEEE 802.1x.

802.1x defines a generic framework for port-based authentication and key distribution. By using the Extensible Authentication Protocol (EAP), an "authenticator" (an Ethernet switch or wireless access point) authenticates a "supplicant" (an Ethernet or wireless NIC) by consulting an authentication server (RADIUS or Kerberos). 802.1x can be implemented with different EAP types, including EAP-MD5 for Ethernet LANs and EAP-TLS for 802.11b WLANs.
802.1x also provides a carrier for secure delivery of session keys used to encrypt traffic between the supplicant and authenticator, addressing another serious omission in the WEP standard. For example, session keys might be created "on the

fly" by the access point or supplied by a RADIUS server. If a war driver with AirSnort recovered keys from WEP session traffic, the keys would be of no value for other sessions.

27. The catch
802.1x products are just now hitting the market. As of this writing, the only operating system with 802.1x support is Microsoft Windows XP (and XP Pro). 802.11b wireless card and access point vendors that support 802.1x today include Agere, Cisco, and Enterasys. Of course, you'll also need an authentication server that supports EAP (Microsoft IAS, Funk Steel-Belted RADIUS, Interlink RAD-P) or LEAP (Cisco ACS).
What if your entire client base isn't running XP? One alternative is deploying an access point with mixed-mode support, like Agere's ORiNOCO AP-2000. For example, the AP-2000 can be configured to authenticate XP clients with 802.1x and other OS clients with Agere's "Closed System Authentication", assigning static IPs to known MAC addresses and denying access to all others.
If you run a Cisco shop, consider using Cisco's proprietary 802.1x derivative, LEAP now available with drivers for other operating systems.

Credentials are another consideration when adopting 802.1x authentication. EAP-TLS requires both the supplicant and authenticator to possess digital certificates, enabling mutual strong authentication. But certificates must be issued by a Certificate Authority more required infrastructure.

Those that prefer (weaker) password authentication may prefer to wait for EAP-SRP (Secure Remote Password), now being defined. Note that 802.1x authenticates MAC addresses, not users. IEEE 802.11i is still working on 802.1x extensions for wireless, including higher level (user) authentication.
Windows XP users seeking wireless public Internet access can give 802.1x a trial run without enterprise rollout. Wayport and Microsoft conducted an 802.1x trial at Seattle-Tacoma airport this summer.

Through the end of January, 2002, Windows XP clients with 802.11b cards get a "free ride" when visiting Wayport-enabled hotel and airport common areas.