WaveLAN Wireless LAN
Technology and Market Backgrounder

Enterprises worldwide are finding that wireless data- communication schemes can be leveraged to increase work-force efficiency, improve customer service, and simplify the installation and connection of information systems. More succinctly, wireless LAN and point-to-point WAN installations can increase income and lower cost. The benefits these technologies bring to the enterprise promise to make wireless LANs a booming market over the next decade.

Wireless LANs can be used both as an extension of a wired enterprise LAN and in place of a wired LAN where warranted. The wireless alternative can match the performance of standard wired networks, and prices are dropping. The same wireless technology can be enhanced with specialized antennas to provide point-to-point WAN links over distances of 5 to 25 miles.

This backgrounder will reveal the opportunities, capabilities, and limitations of wireless data-communication schemes. The subject at hand is examined from two perspectives. First, the backgrounder offers a system-level examination of wireless LANs and WANs including basic specifications and a look at already proven installations and applications. The second section looks underneath the system level and examines the enabling RF communications technology. The section explains the spread spectrum communication schemes used in WaveLAN and other systems. Moreover, competing DSSS (direct-sequence spread spectrum) and FHSS (frequency-hopping spread spectrum) schemes are compared based on capacity, performance, range, and robustness to noise and cost.

Section 1
Wireless LANs -- How, Why, and Where

Wireless telephony technology has been widely accepted as an extension to wired infrastructure for mobile workers. Now, information managers can extend wired data networks with reasonably priced wireless technology for mobile workers requiring seamless data links. Moreover, the wireless LANs offer secure data transmission that rivals desktop schemes such as Ethernet and Token Ring in terms of latency and data availability.

Wireless LANs are a close cousin of the wired Ethernet or Token Ring networks that today, are ubiquitous in business computing environments. The wireless alternative simply uses RF data transmission techniques in place of the wire used in Ethernet or Token Ring environments.

In either the wired or wireless case, the LAN allows users connected to the network to seamlessly access disk drives, printers, and other peripherals on other connected systems. In fact, wired and wireless nodes are regularly combined on the same network with no logical distinction from a users perspective.

The same technology underlies wireless LANs that is widely used in a variety of digital communication applications ranging from digital cellular to in-building wireless phone systems. For example, Lucent Technologies WaveLAN family of products employee a spread-spectrum technology for wireless data transfers that shares a heritage with digital cellular phone networks based on CDMA (code division multiple access). For more information on the technology, see Section 2.

The synergy between the wireless LAN technology and other digital communication schemes is important for two reasons. One, the underlying communication technology, originally developed for military communications, has been proven reliable and secure. Two, the widespread use of spread-spectrum communications will help lower cost across all applications over time.

Wireless performance

Performance is perhaps the most misunderstood characteristic of wireless LANs. In practice, wireless LAN performance compares favorably to a wired LAN such as Ethernet both in terms of latency and data rate. For example, a computer user running Microsoft Windows/Office on a local disk and loading data files over the network would notice no difference whether connected to Ethernet or Lucent's WaveLAN. In fact, any unsatisfactory performance in either case would result from a network saturated by too-many users rather than from the baseline performance offered by the network.

Conversely, a worker connected to Ethernet or WaveLAN might realize unsatisfactory performance when running the Windows operating system and the Office application across a network link. Neither Ethernet or WaveLAN can approach local hard disk performance and the Windows/Office environment accesses a disk almost constantly. Multimedia data transfers such as live video can also quickly saturate both wired and wireless networks.

The qualitative comparison of a wireless LAN and Ethernet is quite valid despite the fact that a wireless LAN such as WaveLAN tops out today at a 2-Mbps peak data rate. The WaveLAN implementation, however, includes a low-latency medium-access scheme. And in either the wired or wireless case, performance centers more on data availability than data rate. The typical Windows/Office user or a user running an on-line transaction processing task such as a retail point-of sale (POS) application depends more on fast access to small bursts of data than on continuous streams of data at maximum rate.

Wireless capacity and range

A modern wireless LAN can connect anywhere from a handful to more than 100 users depending on the application. For example, a single WaveLAN access point can comfortably support 30 to 40 typical Windows nodes that store Windows and applications locally and use the network for file and printer access. In POS applications, a single WaveLAN access point has been successfully used to connect 120 users or more. WaveLAN has no hard physical limit on the number of nodes connected to a single wireless channel.

The users of a wireless LAN can be spread throughout an organization. A single WaveLAN cell can cover an area of 5000 square meters. Moreover, multiple cell implementations can expand both user capacity and range.

Benefits of wireless LANs

Workers in jobs ranging from retail organizations to shipping companies to the office executive can derive benefits from wireless LAN technology. The most basic benefits are exactly analogous to using wireless telephones.

Across industries today, for example, management has accepted that the company bottom line benefits when some employees carry a cordless phone. Even within a facility, many workers carry a cordless handset that acts as an extension to the wired phone plant and PBX. The fact is that many employees spend several hours per day away from their desks.

These workers that need cordless phones are often knowledge workers that also use a computer constantly. They regularly carry a notebook computer from meeting to meeting away from their office. To be most effective, the workers try and connect their notebooks to the corporate LAN in places such as a conference room or colleagues' office. Without wireless LAN capabilities, the notebook user must constantly search for a wired network connection and manage the logistics of the wire itself.

For example, Carnegie Mellon University has developed campus- and citywide schemes to allow professors and students with portable computers to maintain constant links to the campus-wide computer network. The university, located in Pittsburgh, is using a combination of CDPD (cellular digital packet data) technology and WaveLAN wireless LAN technology to provide anytime anywhere connections for portable systems.

The computer services and computer sciences department found ways to add both CDPD and Lucent WaveLAN radios to portable systems. The CDPD interface connects through the citywide cellular phone network and provides packet-data services. Most users only send and receive e-mail through the CDPD link which the university has found can provide throughput approaching that of a 9600-kbps modem.

When users are within the bounds of campus, they can connect to the campus network at much higher speed via the WaveLAN interface. The university ultimately plans to spread approximately 200 wireless LAN access points around the campus providing coverage throughout 12 major buildings on campus and adjacent outside areas. Already, access points have been activated in five buildings. When in range of a wireless LAN node, users can perform tasks ranging from internet access to database searches, to software uploads/downloads.

University life demands that the users spend many hours away from an office each day. These knowledge workers still need computer access however, and wireless LAN technology provides that seamless access.

Side benefits of wireless LANs include the elimination of the need for wiring and the elimination of the need to constantly reconfigure the LAN due to moves, adds, and rearranges. The elimination of wiring can be a major benefit when installing a LAN in a building with no existing LAN infrastructure. In some cases, such as in historic buildings, drilling of holes in walls may be prohibited. In other cases, marble or granite walls and floors can make the installation of wiring problematic. Even hazardous materials such as asbestos embedded in walls and ceilings can block the installation of wire.

The moves, adds, and rearranges problem is common today especially in the dynamic high-tech world. Companies constantly shift the responsibilities of members of the workforce and regularly move personnel around a facility. A wireless LAN can substantially simplify such moves. In fact, the system administrator can accommodate such changes from a central control console without the need to visit a wiring closet or modify a LAN switch.

Proven applications for wireless LANs

Specific applications reveal the most compelling benefits of wireless LANs. Retail POS applications, for example, have clearly benefited from the installation of wireless LANs. One good example of a point-of-sale application is restaurants.

Applebees International, for example, operates a chain of casual restaurants. The company is using WaveLAN to connect touch-screen monitors spread throughout a restaurant with a local compute server. Food servers take orders from a customer, enter the order through a touchscreen, and the order is forwarded to the bar or kitchen. The processor also allows the local compute server in each restaurant to automatically track sales and upload the information to headquarters.

WaveLAN allowed Applebees to install the POS systems without retrofitting all of its locations with new wiring. Moreover, the company avoided wiring problems that are common in high-traffic areas such as a restaurant and areas with heavy equipment such as refrigeration units and open grills.

Down the road, however, the restaurant chain also plans to take further advantage of the portability afforded by WaveLAN. The company plans to eventually equip all food servers with portable handheld terminals that will further speed the flow of orders to the kitchen, and allow the kitchen to signal servers when orders are ready for delivery to a customer.

Successful retail applications also include warehouse stores such as GORO Yudhistra Utama that sells toys, tools, food, and other items in locations throughout Indonesia. GORO's warehouse business model is based on order stations that customers use to place orders. Forklift trucks are then dispatched in the product warehouse to fill the order.

GORO has installed a WaveLAN node, notebook computer, and bar-code reader on each forklift truck, which operates in an 11,000 square meter store and warehouse space. The trucks can receive and process customer orders immediately. The order stations are also connected to a central server with WaveLAN and allow customers to check stock and prices in addition to placing orders.

Medical applications

Other applications include medical and manufacturing. Manufacturing operations can use wireless LANs for inventory control and even to control factory equipment. In Grandview Hospital and Medical Center in Dayton, OH, for example, nurses use WaveLAN-equipped DoCuCarts to eliminate constant trips back to a central nursing station. The carts carry the medical equipment and supplies that a nurse needs during rounds. The nurse can use the WaveLAN terminal to immediately add patient data to the hospital database and significantly increase the amount of time the nurses can spend at bedside with patients.

Even government organizations have begun to make use of wireless technology. For example, the Ministry of Posts and Telecommunications in Japan has used the technology to simplify its job of managing government purchases. The Ministry had developed a wired infrastructure based on ISDN and Ethernet to connect severs at the main office in Tokyo with eleven regional offices. The Ministry sought a way to connect individual workstation clusters with the servers in each office without temporarily moving the staff and installing cables. The wireless LAN alternative allowed the Ministry to avoid disrupting the staff and to bring the network on line immediately without the typical installation delays.

WAN opportunities

A number of opportunities are also available for using wireless LAN technology in WAN applications. For example, the technology can be used to provide a point-to-point link that can bridge two wired LANs. Using line-of-sight, directional antennas, Lucent specifies that WaveLAN can cover distances as great as five miles. Third-party partners have developed even more elaborate antenna schemes that can stretch the distance to 25 miles.

The point-to-point capability is especially valuable in areas that have little installed telecommunication or data communication infrastructure, and in areas where barriers prevent the installation of additional wiring.

For example, the Powerhouse Museum in Sydney Australia consists of two buildings that are separated by a few blocks. The museum had a 48-kbps leased line that had been installed in the days of terminals and simple ASCII characters. The line had long been overwhelmed by the more than 200 computers spread throughout the museum for the purposes of classifying more than 300,000 objects in the museum.

Adding a higher-speed landline between the buildings was not feasible due to a major thoroughfare separating the two. A WaveLAN point-to-point link was installed boosting the building-to-building connection from 48 Kbps to 2 Mbps.

LAN in a box

One final application for wireless LANs has been termed "LAN in a box". A wireless LAN offers an organization the unique ability to fully preconfigure a network and ship the network to a site for immediate installation. The central organization can install the wireless LAN node in each computer, set user names and addresses, and preconfigure the server. On site, installation simply requires that the server and workstations be powered on. The LAN in a box concept is proving popular with organizations ranging from financial institutions to travel agents.

The future of wireless LANs is bright throughout the world. Lucent has developed a strong customer base in the USA, with Japan and Korea following closely as the second and third largest markets for the technology. Meanwhile, VLSI advancements and increased volumes continue to make the technology more affordable. Interface cards for portable or desktop PCs have dropped to $695 and server access points cost less than $2000. A compelling price, ease of use, and wired-LAN performance promise to further boost the burgeoning market.

Section 2
Spread Spectrum Data Communications
Direct Sequence vs. Frequency Hopping technology

Why should anyone care exactly how any specific company's wireless LAN is implemented? As stable a foundation as one can find anywhere underlies the burgeoning wireless LAN and WAN products such as WaveLAN. The RF-based products that dominate the market virtually all rely on one form or another of spread-spectrum technology originally developed for military applications. The techniques have been proven reliable and secure.

Still, the implementation details are important for several reasons. For starters, no standards exist today for wireless LANs although the IEEE 802.11 committee is working to establish standards. The lack of standards means that one company's product will not interoperate with another company's offerings the way Ethernet cards from any vendor will interoperate. Of even greater importance today and for the foreseeable future, the different spread-spectrum implementations available vary widely in terms of performance, capacity, range, and immunity to noise and cost.

The type of spread spectrum implementation stands out as the most significant characteristic of any wireless LAN product. Generally speaking, the available products employ either DSSS (direct-sequence spread spectrum) or FHSS (frequency-hopping spread spectrum) techniques. Both are viable choices that have been field proven. An examination of the two technologies, however, can reveal how well either fits in specific applications.

Commonalties in DSSS and FHSS

Before addressing differences in wireless LAN implementations, first consider the commonalties. First, all of the products must work in the same frequency bands. In the United States, vendors can use the 915-MHz, 2.4-GHz, and 5.8-GHz bands that have been designated by the FCC (Federal Communications Commission) for spread-spectrum-based ISM (industrial, science, and medical) applications and require no license. Similarly, Japan's MPT (Ministry of Post and Telecommunications) and Europe's CEPT (Conference of European Postal and Telecommunications) have made the 2.4-GHz band available for unlicensed use. Most of the remainder of the world follow the policies set in the US or Europe.

Because the 2.4-GHz band is available virtually worldwide, most vendors are concentrating new wireless LAN efforts on that band. Some products, such as Lucent's WaveLAN family, do support both the 915-MHz and 2.4-GHz bands thereby allowing users maximum flexibility in finding free spectrum. Implementations that will be based on the IEEE 802.11 standard are focused on the 2.4-Ghz band.

DSSS and FHSS will also share the IEEE 802.11 standard. Currently under development, the standard will include provisions for both DSSS and FHSS physical layers (PHYs), as well as an infrared physical layer. Completion of the standard is not expected until mid 1997. Down the road, compliance with the standard will enable products from different vendors to interoperate and virtually all vendors are expected to migrate their wireless LAN implementations to comply with the standard.

The two RF PHY options resulted from IEEE 802.11 working group members that cited the need for a choice between complexity of implementation and cost relative to performance and reliability. The wireless PHY is analogous to the choice of 10BaseT, 10Base5, 100BaseT, and other PHYs in the hugely successfully Ethernet arena.

Anyone contemplating the subject of wireless LANs should fully understand the present and future ramifications of IEEE 802.11 despite its work-in-progress status. First, some portions of the standard, such as modulation schemes and rates, fundamentally limit the throughput that any compliant LAN implementation can realize. Together, parts of the IEEE standard and governmental regulations are key to any comparison of DSSS and FHSS technologies and this backgrounder will in all cases make these limitations known.

Many suppliers of wireless LAN products, including Lucent, are already attempting to comply with any known elements of the standard. Moving toward the standard where possible today will minimize the differences in present and future LAN configurations for customers. Everyone should realize, however, that any claims of complete compliance today are inherently false and that the draft standard could include mistakes and is still likely to change.

DSSS and FHSS technologies also share the name spread spectrum. Generally, the spread spectrum moniker implies that a RF transmitter spreads a relatively-narrow-band signal over a wide-band spectrum. By using a wider-than-necessary band to transmit a signal, the scheme ensures that the receiver can correctly interpret the transmission even in the presence of noise. Moreover, the techniques used to spread the signal inherently add security to the system.

Despite sharing the spread spectrum name, DSSS and FHSS could hardly be less similar. DSSS systems broaden the signaling band by artificially increasing the modulation rate using a spreading code. FHSS systems, meanwhile, hop from narrow band to narrow band within a wide band, using each narrow band for a specific time period. The frequency hops appear random although they actually occur in a pseudo-random sequence tracked by sender and receiver.

DSSS basics and operation

To understand the differences in DSSS and FHSS systems, consider each in more detail. A DSSS transmitter operates on an incoming data stream of a certain bit rate (bps). The incoming bit stream is typically converted into a symbol stream where each symbol represents a group of 1, 2, or more bits. Modulation techniques similar to those used in voice-band data modems are used to operate on the symbol stream and generate the transmitted signal. The DSSS transmitter modulates or multiplies each data bit or symbol with a pseudorandom noise (PN) sequence that is also called a "chip" sequence. This multiplication provides the spreading phenomena.

To understand the DSSS modulation process consider an example. Lucent's WaveLAN products use DQPSK (differential quadrature phase shift keying) modulation and the draft IEEE 802.11 standard also specifies DQPSK. In a QPSK or DQSP modulator, symbols are differentiated by the phase of a sinusoidal wave. Each symbol can be represented by a point in a symbol constellation plotted on a standard Cartesian graph. The so-called I (in-phase) and Q (quadrature phase) axes of the graph actually represent the cosine and sine of the phase angle respectively.

At the receiver, the implementor has the choice of using DQPSK or QPSK detection. In a DQPSK (also called a non-coherent QPSK) design, the receiver uses the previous phase as a reference to detect the current symbol. Coherent QPSK receivers, meanwhile, include PLL-based, carrier-recovery circuits. A QPSK receiver design is slightly more complex than a DQPSK design and offers a 3-dB advantage in detection efficiency.

A typical QPSK implementation such as Lucent's codes two bits in each symbol. The constellation for this implementation consists of four symbols each separated by 90 degrees. The constellation points are actually at 45, 135, 225, 315 degrees -- each at constant amplitude -- and each represents a two-bit symbol (00, 01, 10, 11). The phase spacing also means that each symbol's I and Q Cartesian coordinates are either +1 or -1.

A standard QPSK modulator would simply transmit the symbol stream. In this example, the transmitted sinusoid would undergo a phase shift a maximum of once during each symbol period and the symbol rate is half the incoming data rate.

To spread the signal, the DSSS modulator then multiplies each symbol by the chip sequence. The chip sequence is actually a series of +1 and -1 values. You can also think of each symbol as being represented by its I and Q values of +1 or -1. Consider a symbol, represented by the I,Q coordinate (1,1) -- polar or phase coordinate 45 degrees. That symbol is multiplied by each member of the chip sequence. In the case of a +1 chip, the symbol is unchanged. In the case of a -1 chip, the symbol shifts 180 degrees in phase or to coordinates (-1,-1).

Where standard QPSK results in at most one phase change per symbol period, DSSS spreading can result in as many as "n" phase changes per period where n is the number of entries in the chip sequence.

DSSS modulation effectively combines a relatively low-rate bitstream with a higher-rate chip sequence of n chips per symbol interval. The parameter n is called the spreading ratio. In the case of the DSSS modulated signal, the apparent symbol rate and required bandwidth for transmission are increased by a factor of n.

The WaveLAN system developed by Lucent modulates each symbol with an 11-chip PN sequence. The resulting 11-Mchip/sec transmitted stream requires 11 MHz of bandwidth to yield an actual raw bit rate of 2 Mbps, and Lucent has plans to increase this raw bit rate in the future. To provide a clear, non-interfering channel, WaveLAN actually reserves 22 MHz for each channel.

The DSSS receiver uses the same PN chip sequence as the transmitter to extract data from the spread signal. The incoming signal is first processed by a pass-band RF filter to reject out-of-band noise and interference. This operation isolates the entire band of interest -- for example the entire 83.5-MHz ISM band located at 2.4 GHz. A down converter (or selective filter) then isolates the 22-MHz channel. A decorrelator function then removes the chip sequence and restores the original signal.

A measure of the results of the decorrelation function is referred to as the processing gain of the receiver. The processing gain actually represents an increase in signal-to-noise ratio with respect to the signal in and out of the decorrelator. This gain allows more reliable data communications.

Think of the robustness to noise in terms of the I and Q coordinates of each chip. Because the original symbol is present in each chip, several chips can be corrupted yet the decorrelator will still recover good data.

Security and robustness to noise in a DSSS system results from the complexity of the DSSS modulation. A higher spreading ratio adds security and immunity to noise to the system but also limits the amount of data that can be handled in a given band. Engineers designing a DSSS system must choose a spreading ratio that allows both robust and cost-effective data transmissions. For wireless LANs, these parameters will ultimately be determined by the IEEE and governmental bodies.

FHSS basics and operation

FHSS systems spread the data stream quite differently than DSSS systems. An FHSS transmitter sends a short burst at one carrier frequency hops to another frequency and sends another short burst and continues this hop-transmit sequence. The rate at which the transmitter moves from one carrier frequency to the next is referred to as the hopping rate. The hopping pattern or sequence is determined by a periodic PN sequence. FHSS systems can be susceptible to noise during any one hop or center frequency but typically can achieve error-free transmissions during other hops. During any given hop, an FHSS transmitter uses a standard non-spreading modulation technique. The draft IEEE 802.11 standard prescribes the use of GFSK (Gaussian frequency shift keying) modulation.

Governmental Regulations

Before comparing DSSS and FHSS systems in more detail, you must also understand the effect of governmental regulations on the two digital communication schemes. The FCC, CEPT, and MPT not only define available bands for unlicensed usage, but also impose other requirements on digital wireless transmissions. Moreover, the three governmental bodies have defined substantially different guidelines for such transmissions including different specific constraints on DSSS and FHSS systems.

The differences start with the available frequency spectrum. Although all three agencies have allocated spectrum in the 2.4-GHz band, the FCC and CEPT make available a band of 83.5 MHz while the MPT agency in Japan limit usage to a band of 26 MHz. For DSSS systems, a government-specified spreading factor or processing gain of at least 10 limits the total bandwidth available. The FCC and MPT explicitly define the spreading factor. In Europe, an allowable ratio of total transmit power and power density imposes essentially the same limit.

Dividing the available spectrum by the required spreading factor provides a simple estimate of bandwidth for DSSS systems -- 8.35 MHz in the US and Europe, and 2.6 MHz in Japan (before spreading takes place). More accurately determining the actual available bandwidth requires an examination of the ability to overlap or collocate wireless cells operating at different center frequencies.

Constraints on FHSS systems include a maximum bandwidth of 1 MHz than can be used during any hop period in the US. The US requirements also require that FHSS systems occupy at least 75 different hop frequencies over time. In Europe, only 20 hop frequencies are required which could make as much as 4 MHz of bandwidth (83.5 MHz/20 hops = 4 MHz) available.

DSSS and FHSS characteristic comparisons

Having set the stage, DSSS and FHSS can now be compared relative to the following characteristics: cost of implementation, performance, capacity, coverage, and immunity to impairments. Comparisons will include a theoretical discussion and an implementation discussion.

In the wireless LAN arena, performance and capacity are without question both the most misunderstood characteristics and the characteristics most subject to blatant and unsubstantiated specmanship by some vendors. To put the confusion in perspective, consider the two. The performance of a communication scheme is typically measured in terms of maximum data rate or bps, and latency or how quickly a system can access the available bandwidth. Capacity, meanwhile, should refer to the number of users that can be served, but in some cases -- especially in wireless LANs -- is also defined in terms of the aggregate throughput measured in bps that can be shared by users.

Wireless LAN performance and capacity

In LANs, performance and capacity are certainly interrelated. Relatively speaking, a shared-medium LAN can serve a few users that need significant bandwidth, or many users with sparser bandwidth needs. Hard capacity limits are rarely reached in any actual LAN installation because of the potential for degraded performance. Lucent Technologies' WaveLAN, in fact, has no hard limit on the number of nodes that can share a single channel.

The relative affect of maximum data rate and latency on realized performance varies by application. A LAN regularly used to transfer large files such as graphics images or to transfer real-time data streams such as live video, benefit most from a high maximum data rate. A more typical transaction-processing application such as a point-of-sale terminal in a retail store depends more on short data bursts and therefore benefits more from a scheme that offers low-latency access to the media. Office oriented applications such as running Windows applications fall somewhere in the middle.

As discussed in the first section of the backgrounder, a single 2-Mbps WaveLAN access point can offer 30 to 40 PC users running typical Windows applications performance that rivals a 10-Mbps Ethernet LAN. The lower-speed wireless LAN achieves such results due to an efficient MAC (medium access control) scheme called CSMA/CA (carrier sense multiple access with collision avoidance).

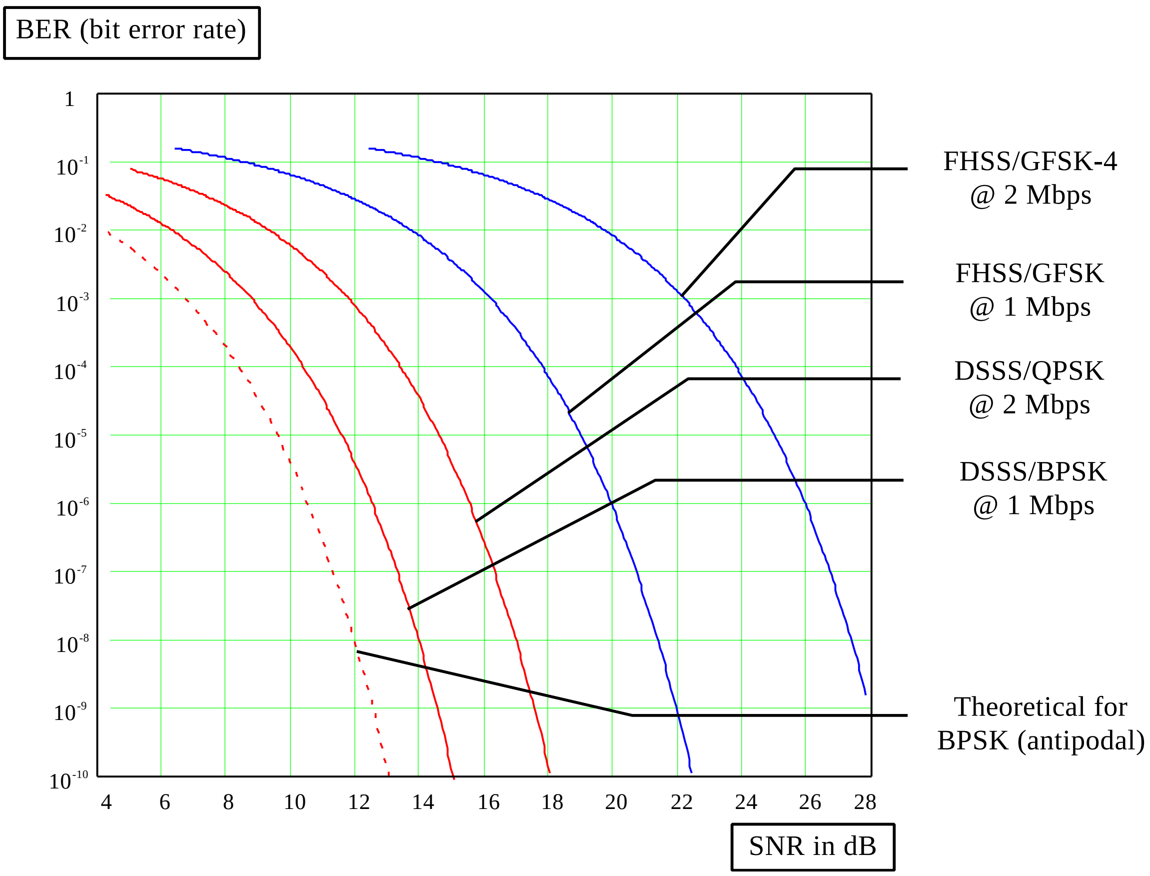

Fig. 1: The ratio of bit error rate (BER) to signal-to-noise ratio (SNR) for BPSK, QPSK and GFSK
(Lower error rate is better, lower SNR for the same error rate indicates more robustness to noise)
Theoretical for BPSK (antipodal) or DSSS/BPSK @ 1Mbps (without implementation loss)
BPSK-based DSSS @ 1Mbps (with implementation loss)
QPSK-based DSSS @ 2Mbps (with implementation loss)
GFSK-based FHSS @ 1Mbps (with implementation loss
GFSK-based FHSS @ 2Mbps (with implementation loss).

To compare DSSS and FHSS LAN performance, start with peak data rate and latencies. DSSS systems have a definite advantage in terms of peak data rate. Available DSSS systems such as WaveLAN can transfer data at 2 Mbps while most FHSS implementations operate at 1 Mbps. Moreover, the IEEE 802.11 draft standard currently describes FHSS operating at 1 Mbps, with optional 2 Mbps speeds possible in optimal-quality conditions. Conversely, the spec defines 2 Mbps as the standard DSSS rate with a 1 Mbps option reserved for particularly noisy environments. Based on theoretical calculations, Figure 1 graphs bit error rate (BER) relative to signal-to-noise ratio (SNR). The graphs show that 2-Mbps DSSS offers a more robust wireless link than 1-Mbps FHSS. More specifically, for a channel with a given SNR level, DSSS offers a link with much lower error rate.

DSSS systems not only offer better performance transferring large files, but also provide lower-latency medium access. Ultimately, the IEEE 802.11 standard will prescribe that both DSSS and FHSS systems use similar medium access schemes. FHSS systems, however, suffer additional latencies on each frequency hop. Moreover the recommended packet size is only 400 bytes for FHSS, while it is 1500 or 2400 bytes for DSSS. This means that FHSS systems will have to break up almost all-long data packets into fragments. A transmission preamble and a MAC header are needed for each fragment as well as a separate acknowledgment frame per transmission. These packet issues significantly add to the overhead of FHSS systems transmitting long packets and allow DSSS channels to support net throughput advantages that are significantly greater than the 2x peak rate advantage.

Capacity and channels

The faster data rate and lower latency offered by DSSS translates directly into the ability to serve more users from a single wireless access point or on a single wireless channel. WaveLAN, in fact, has been successfully deployed in point-of-sale applications with more than 120 nodes connected to a single access point.

While the capacity comparison between DSSS and FHSS systems is straightforward when a single access point is concerned, multiple channels complicate the matter significantly. In fact, multiple channels can be collocated in the same area to boost capacity in either type of system. An additional channel can't boost the maximum rate at which a single data file is transferred. A second channel can extend the overall bandwidth that is shared among users.

Multiple, collocated channels or access points are the basis for claims by some vendors to offer 10 Mbps or more capacity. A close look will reveal just how much capacity can be reasonably expanded through multiple channels.

Consider systems that meet the US FCC requirements. A DSSS system such as WaveLAN can break the available 83.5 MHz of bandwidth into eight separate 11-MHz channels. Adjacent channels overlap and access points using adjacent channels would not typically be collocated. The WaveLAN frequency allocation does result in three channels that are separated by more than 22 MHz. Access points using these three channels can be collocated thus tripling the available aggregate bandwidth. The remainder of the channels is used to complete the layout of a cell structure.

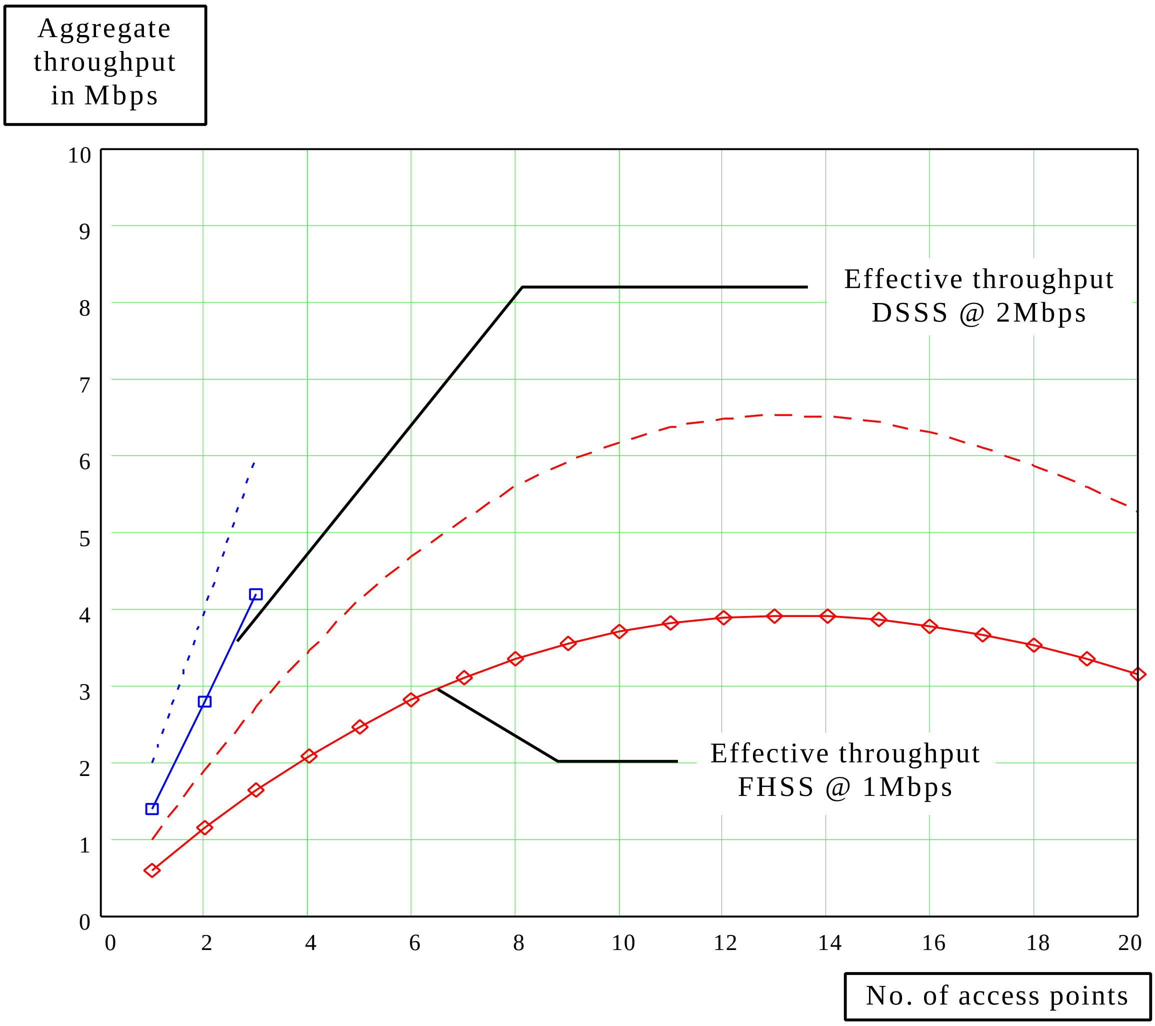

Fig. 2: Aggregate throughput vs. number of collocated access points
Dotted lines give the gross capacity.
Solid lines with boxes and diamonds gives effective net throughput.
The difference between gross and net throughput comes form overhead in frame headers, the requirement for transmission of acknowledgement frames, and losses due to collisions.

Each FHSS access point, meanwhile, uses almost the full 83.5-MHz band for hopping. To collocate FHSS access points, you must acknowledge that the access points will occasionally collide on a hop. Moreover, collocated access points will also suffer collisions due to interference when using directly-adjacent hop frequencies. Different and unsynchronized hopping patterns can minimize the collision risk to a degree. You can achieve a significant increase in total shareable throughput with a low number of collocated access points.

Some FHSS proponents have even claimed the ability to collocate 10 or more access points and claimed a 10x or greater linear ramp in aggregate throughput. To date, such capabilities haven't been proven in real-world demonstrations.

Practical FHSS limitations

Theoretical calculations show substantially different results (See Figure 2). The tests reveal degradation that occurs, as more access points are collocated thereby increasing the instances when two access points simultaneously attempt to use the same or adjacent frequency during a hop. Such collisions require regular retransmission by both nodes.

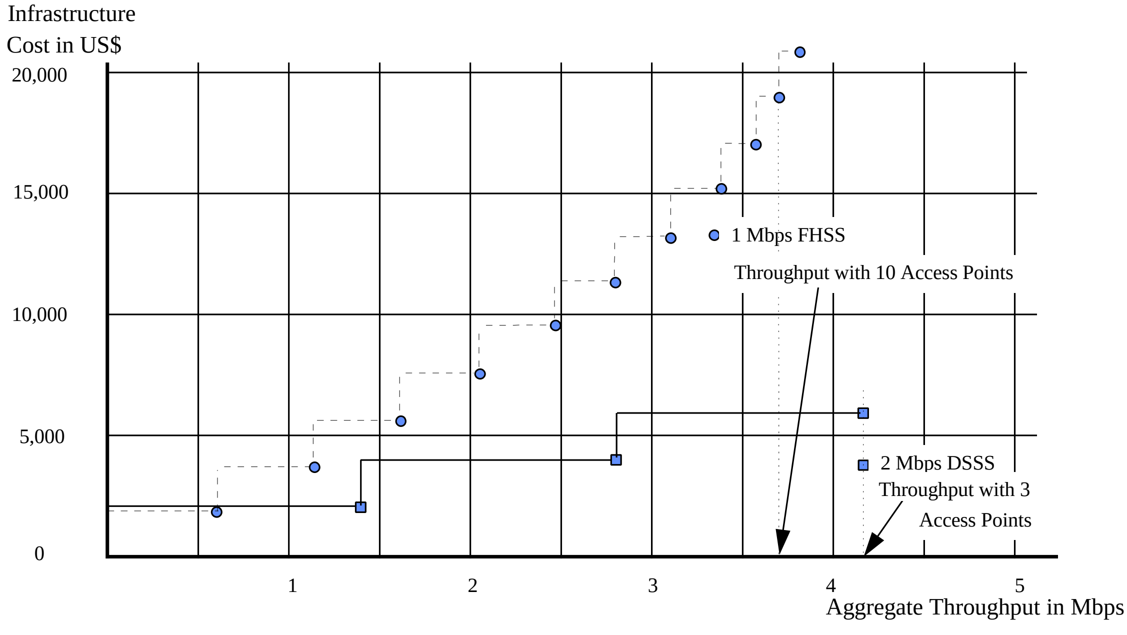

Fig. 3: Cost of Access points as a function of required throughput
Based on data in Fig. 2 with access points deployed within the same cell to boost aggregate throughput. Note that FHSS can never go beyond 4 Mbps (adding access points would just decrease throughput and add to the cost) while DSSS can get up to 4.2 Mbps (additional access points would give the same throughput and add to the cost) with 3 access points.

The calculations show that when using 1-Mbps FHSS access points, aggregate bandwidth (taking latencies and packet overhead into account) never exceeds 4 Mbps regardless of how many APs are used. Moreover, more than 10 channels are required to reach 3.7 Mbps in aggregate throughput. Due to practical limitations for the isolation between adjacent hop frequencies and the actual ratio between desired signal and adjacent hop frequency signal, the aggregate throughput growth at an increase of collocated access points stagnates. Few customers will be likely to collocate and pay for more than five or six access points with such significant degradation. Figure 3 graphs the cost associated with each technology relative to aggregate throughput. DSSS access points prove to be the bargain in reaching 6 Mbps aggregate throughput even if the DSSS equipment cost slightly more.

Not only do calculations show that FHSS systems don't scale significantly better than DSSS systems, but the latter can also be packed more closely together. The QPSK and BPSK modulation used in DSSS systems is more robust relative to cochannel interference than is GFSK modulation. This advantage means that the same channel in DSSS systems can be used in nearby cells without interfering. The IEEE 802.11 draft spec recognizes this fact in how it defines a difference in capture with respect to power level. The spec has set the DSSS defer level for valid transmissions 15 dB above the FHSS defer level. This defer threshold advantage allows DSSS cells using the same channel to be packed four to eight times as dense as FHSS when it reuses the same hop sequence. In total, the advantages in more closely packed cells, MAC efficiencies and packet efficiencies result in a 4 to 10x throughput or capacity advantage for DSSS systems.

DSSS and FHSS implementation costs

A straightforward cost comparison of DSSS and FHSS is almost impossible. Some will argue that a stand-alone DSSS radio cost more than an FHSS radio, but as discussed above, available DSSS implementations also deliver more throughput. A cost per bits transferred comparison would clearly favor DSSS. To get a complete view of cost issues, you must consider costs relative to all other characteristics of the two technologies. Therefore this backgrounder will address cost issues from a number of perspectives throughout the technology comparisons.

Generally, DSSS has been considered a more expensive technology to implement than FHSS. The implementations cost differences come down to two key areas. First, the DSSS scheme requires a more complex signal spreading and modulation scheme that results in the need for a more powerful DSP in the receiver to handle the decorellation function. The abundance of low-cost but powerful DSPs, however, eliminates this obstacle for companies with expertise in signal processing. Second, the DSSS modulation techniques require a more expensive power amplifier in the transmitter than do the available FHSS implementations.

Before you can consider the power amplifier, and how it affects other characteristics further, you must understand the differences in modulation technology.

Modulation characteristics

DSSS systems must use some form of phase modulation to implement the chipping scheme. The IEEE 802.11 implementation will likely be QPSK for 2-Mbps data transfers as implemented by Lucent in WaveLAN. The current draft of the standard also prescribes BPSK (binary phase shift keying) modulation for 1-Mbps transfers when noise or other impairments require a fall back to a lower rate. In the future, it's certainly possible that more complex schemes such as QAM (quadrature amplitude modulation) could be adopted to code more bits into a symbol and boost available throughput.

FHSS systems could actually use any modulation technique during the short bursts. The IEEE draft standard prescribes GFSK for use in 1-Mbps FHSS implementations. A FSK-based scheme was chosen because it is generally considered to be simple to implement. On the other hand, FSK lacks the robustness of QPSK, which results in advantages in terms of range, capacity, and immunity to noise for DSSS implementations.

The implementation differences in FSK and QPSK center on the RF transmitter. FSK radios uses different frequencies to represent different symbols or binary states. The radios are referred to as constant envelope radios which means that the amplitude of the transmitted signal never changes. The content in the signal is stored strictly in the frequency component. The constant envelope radio, therefore, can use a low-cost, non-linear power amplifier that saturates and clips the peaks of the transmitted signal.

At first glance, you might think that QPSK transmission would be equally invariant in amplitude while techniques such as QAM obviously depend on amplitude variation to convey data. Actually, QPSK radios generate amplitude ripples when making instantaneous phase changes. A linear amplifier is required to accurately transmit these amplitude fluctuations to maintain spectral purity and simplify the filtering required in the receiver.

Both a cost and a power-conversion efficiency advantage can be attributed to non-linear power amplifiers. The efficiency advantage would seem significant because many applications for wireless LANs require portability and batter-powered operation. Better power efficiency should prolong battery life.

In the case of DSSS and FHSS LANs, however, the affects of power efficiency are practically negligible. DSSS requires slightly more power during transmissions, but the typical wireless LAN node spends very little time transmitting. Moreover, when DSSS nodes transmit, they do so at a significantly higher data rate that do most FHSS nodes. Therefore, the DSSS node spends far less time with transmit power turned on than does an FHSS node.

In the case of cost, you get what you pay for. Designing a linear amplifier is still today a non-trivial task. In fact, FSK was chosen for FHSS systems, in part, so that companies lacking the mixed-signal expertise necessary could still participate in the wireless LAN market.

Clearly, state-of-the-art communication products are moving toward modulation techniques that require a linear amplifier. Examples can be found in products ranging from dial-up data modems to cellular phones. For instance, analog cellular phones (AMPS phones) use a non-linear amplifier while the more capable digital-cellular phones just emerging require a linear amplifier yet yield greater capacity per channel and more immunity to noise. As you will see shortly, the cost of a linear amplifier pays off relative to coverage area and immunity to noise.

Coverage area and reliability

In general DSSS systems offer a receiver sensitivity advantage discussed previously, and that advantage should result in greater range and more reliable links. In fact, coverage area in a wireless LAN is determined by a combination of three factors -- total output power, SNR, and total receiver noise level. The third factor is normally independent of DSSS and FHSS implementations, since it is mainly based on the noise introduced in the first stages of the receiver (RF-input stage: filter and gain stages) and/or the environmental noise, whichever is stronger. The design of the DSSS or FHSS receiver front-end is subjected to the same quality-price constraints with regard to the introduction of receiver noise factor, and with regard to implementing a large dynamic range for the receive level. The same environmental noise level applies for both systems. Therefore total receiver noise is not important for the purpose of a comparison.

The governmental regulations can play a role in the transmit power component of range. The FCC allows either DSSS or FHSS systems to use 1 Watt. In Europe, both systems can use 100 mW maximum transmit power, and power density is limited to 10 mW/MHz for DSSS and 100 mW/MHz for FHSS.

Actually other practical consideration prove more important in determining transmit power. Both DSSS and FHSS implementations must support the PCMCIA form-factor used in portable, battery-powered systems. These implementations provide a transmit power in the 50 - 100 mW range. These types of products have to deal both with current consumption limits not only in the transmit mode but also in the receive mode (battery lifetime). So in practice, the transmit power in real products is limited by the amount of current that can be drawn from a PCMCIA slot, rather than the regulatory power constraint.

The practical limitations on transmit power mean that SNR differences in the DSSS and FHSS largely determine differences in range. The difference between DSSS and FHSS in the allowable path loss between transmitter and receiver, is dominated by the difference in the minimum required SNR. As Figure 1 depicts the QPSK modulation used in DSSS systems yields a 4 to 10 dB SNR advantage relative to GFSK as defined by the current draft of the IEEE 802.11 standard. The SNR advantage translates into a 50% advantage in coverage area for a DSSS system operating at the same bit rate and transmit power as an FHSS systems. Figures 4 and 5 depict range and reliability advantages of DSSS systems.

Immunity to impairments

Wireless LAN systems must overcome several types of signal impairments to offer users the reliable data transfers that they've come to expect from wired LANs. The impairments come from multipath reflections of the transmitted signal and from external narrow- and wide-band noise sources.

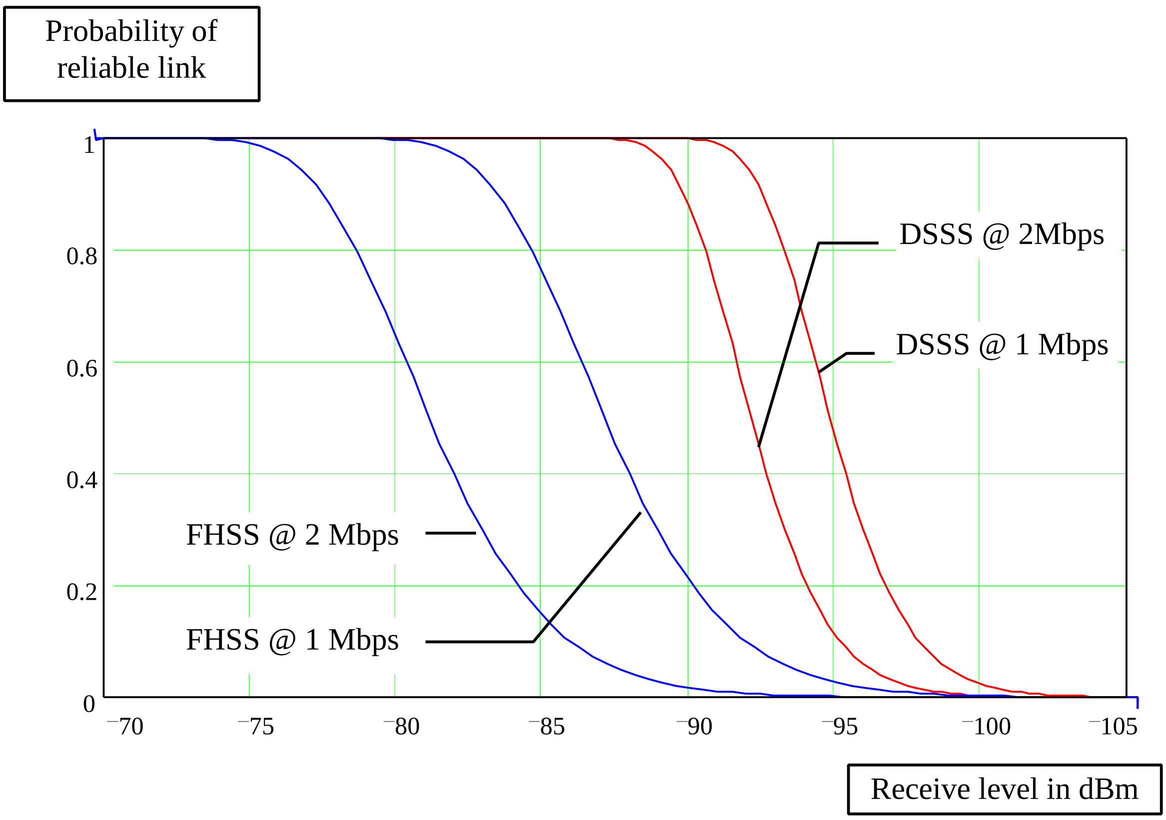

The multipath phenomena due to signal reflections -- especially troublesome in environments such as a warehouse -- results in multiple copies of a signal arriving at the receiver. The effect is called intersymbol interference and can result in frequency-selective fading of the transmitted signal.

DSSS and FHSS systems combat frequency-selective fading in different ways and each are effective. FHSS systems simply retransmit the same information a second time at the next hop frequency when fading disrupts a transmission. A DSSS receiver, meanwhile, integrates the signal across a much wider spectrum. The receiver can therefore operate with low responses in part of the spectrum with no need for retransmission.

Spread spectrum systems work similarly in the presence of narrow-band noise. FHSS systems retransmit at a different hop, while DSSS systems are immune to narrow-band noise. Both types of system can recover, but FHSS suffer a slight degradation in performance due to retransmission.

Wide-band noise can cause problems for both FHSS and DSSS systems. Luckily, most wide-band is sporadic and comes from sources such as a microwave oven. Therefore, wide-band noise typically just causes sporadic retransmissions on either FHSS or DSSS systems.

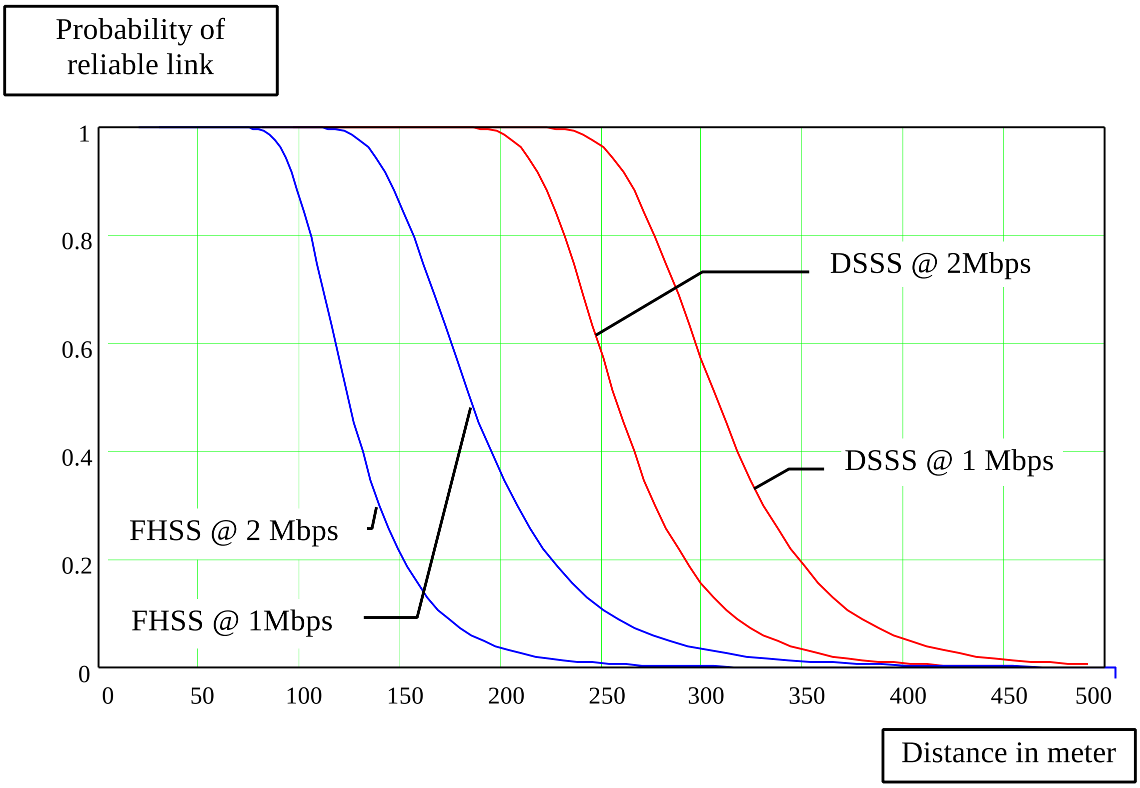

Fig. [illegible]: Probability of a reliable link vs. [illegible]
Based on FHSS and DSSS systems both having a transmitter output power of 50 mW, antenna-diversity (2 antennas), total receiver noise of 7 dB above thermal noise, BER performance as in Fig. 1, and an indoor channel having a delay spread of 50 nsec. [illegible] Reliability of link; the higher in the graph the better, 1 is reliable link in all cases. Horizontal axis: Distance)

Cost conclusions

After an examination of performance, capacity, range, and robustness to noise, it's clear that DSSS is a significantly better choice unless FHSS implementations offer substantial cost advantages. And the implementation cost differences appear significant at first glance, but a detailed look reveals little if any real cost difference. First and foremost, a single chip implementation of both FHSS and DSSS modems are required to minimize cost and maximize deployment. Neither is available as a single chip yet but both certainly will reach that level of integration in the future.

The keys to a single chip implementation are state-of-the-art fabrication capabilities, mixed signal IC expertise, and high-performance DSP cores. Lucent lists all three keys as core competencies and will certainly be able to offer a single-chip DSSS implementation.

Today, the best way to compare the two may be the price and configuration of end-user products. Lucent, for example, offers a WaveLAN node on a half-size ISA-bus card that sells for $ 695, a PCMCIA card for $ 695 and a WavePOINT access point for $ 1995.

Ancillary considerations

In addition to technical concerns, some other factors are worth considering when comparing wireless LAN products. For example, compatibility with the ongoing IEEE 802.11 effort can be important. Next year, all vendors will begin transitioning products to comply with the standard. Those products that closely track the standard now will simplify the transition for customers with installed products. Much of the DSSS implementation specified in the IEEE 802.11 draft is based on proposals from Lucent Technologies.

Customers should also consider market share leaders when choosing a wireless LAN technology. The installed base is directly indicative of the ongoing sales volume. Booming volumes will be key to lower prices and more pervasive use of wireless LANs.

Today, Lucent holds well over 50% of the market for wireless LANs with its DSSS-based WaveLAN family. Moreover, Aironet is second in the overall market and also uses DSSS. In total, DSSS represents 75% of the installed base, and that market share is not declining.

Looking to the future

A look to the future will also confirm DSSS as the best technical choice for wireless LAN investment today. Lucent Technologies projects that DSSS has the potential to make a jump to 10-Mbps data rates using a single channel. Lucent Technologies has extensive experience in modulation algorithms and radio designs that will enable a move to 10 Mbps. Even if FHSS vendors can double or triple peak data rates down the road, the gap between DSSS and FHSS in terms of performance will widen.

Conclusion

All of the technology, cost, and forward-looking comparisons point to DSSS as the most appropriate choice for virtually every application. In rare simple applications, FHSS may prove more cost effective than DSSS but the former can limit future expansion.

Any cost advantage enjoyed by FHSS systems is also sure to erode as DSP Imps continue to drop in price. DSPs are now used in everything from cellular phones to toys. The ICs are becoming sufficiently inexpensive that DSSS systems will match FHSS in price.

The more robust DSSS DSP implementation can provide other advantages as well. Spare MIPs can be used to add other features such as support for multiple bit rates or multiple modulation schemes.

Today, DSSS is also the only choice backed by world-class IC design and fab capacity. The companies promoting FHSS are fabless and lack the DSP and mixed-signal design skills that Lucent Technologies uses in a host of products. Synergies in these different products and unmatched fab capacity ensure that DSSS ICs, boards, and systems will be available in quantity at a low price.