Interference Immunity of 2.4 GHz Wireless LANs

Of the three major technologies available for this band, only HomeRF is designed with a frequency agile physical layer and robust upper layer protocols to combat 2.4 GHz interference. This is what makes HomeRF the ideal wireless LAN technology for the home environment.

HomeRF is the most interference immune of the three major 2.4 GHz wireless LAN standards

The unlicensed 2.4 GHz band has become one of the most active communications bands in the RF spectrum. As market demand for local area networks (LANs) extends from the corporate to the home environment, 2.4 GHz wireless LAN equipment is being produced and purchased in greater volumes than ever before. In the home environment the HomeRF standard[1] is emerging as the leading technology. The IEEE 802.11b standard[2] is making strong inroads into the corporate environment. And finally, though initially designed as a cable replacement technology and not a wireless LAN, Bluetooth[3] devices are expected to be showing up in the 2.4 GHz band by the hundreds of millions[4]. As more devices are deployed into this unlicensed band, the interference immunity of those devices will become increasingly important to technology adopters[5]. As is demonstrated below, by thoughtful design in multiple layers of its construction, the HomeRF standard produces products which are more robust against interference than any other wireless LAN standard for the 2.4 GHz band.

How wireless LANs Work

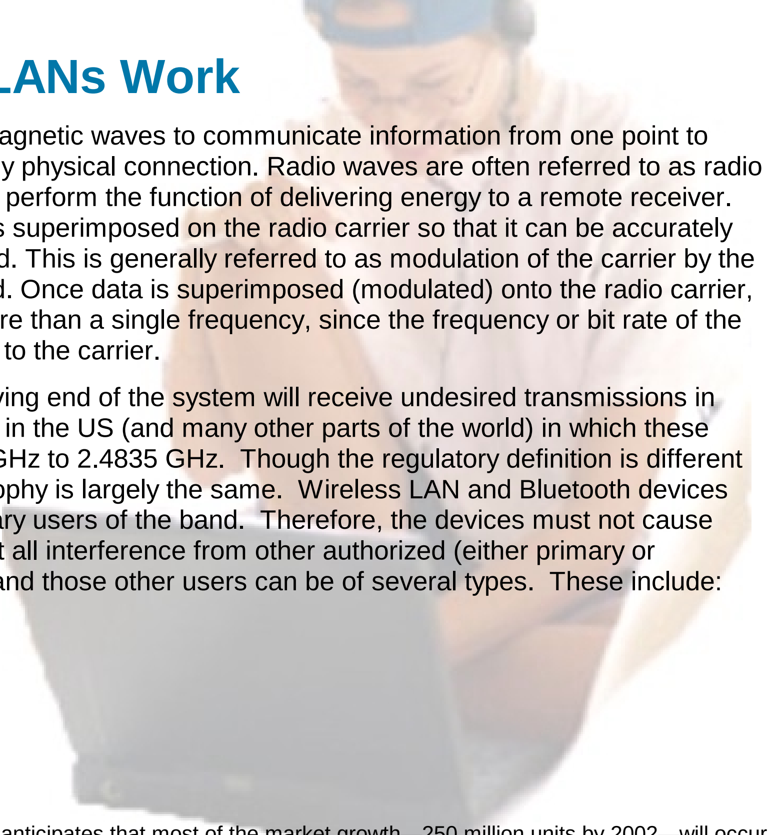

Wireless LANs use electromagnetic waves to communicate information from one point to another without relying on any physical connection. Radio waves are often referred to as radio carriers because they simply perform the function of delivering energy to a remote receiver. The data being transmitted is superimposed on the radio carrier so that it can be accurately extracted at the receiving end. This is generally referred to as modulation of the carrier by the information being transmitted. Once data is superimposed (modulated) onto the radio carrier, the radio signal occupies more than a single frequency, since the frequency or bit rate of the modulating information adds to the carrier.

A wireless system must anticipate that the receiving end of the system will receive undesired transmissions in addition to the desired transmissions. The band in the US (and many other parts of the world) in which these unlicensed devices can operate is from 2.4000 GHz to 2.4835 GHz. Though the regulatory definition is different in different parts of the world, the general philosophy is largely the same. Wireless LAN and Bluetooth devices operating in the band are secondary to the primary users of the band. Therefore, the devices must not cause unacceptable interference, and they must accept all interference from other authorized (either primary or secondary) users of the band. In the 2.4 GHz band those other users can be of several types. These include:

[1] http://www.homerf.org

[2] http://grouper.ieee.org/groups/802/11/main.html

[3] http://www.Bluetooth.org

[4] Market-analyst company Dataquest (www.dataquest.com) anticipates that most of the market growth—250 million units by 2002—will occur in the cell-phone market. This was the original focus of the Bluetooth technology developers, the replacement of cables used to connect cellular phone with headsets or PDAs, for example.

[5] The issue of interference in the 2.4 GHz band was the subject of a Wall Street Journal article, “Raft of New Wireless Technologies Could Lead to Airwave Gridlock”, by Jared Sandberg, January 8, 2001.

- ISM devices. The ISM (Industrial, Scientific, and Medical) bands are established by treaty at the International Telecommunications Union (ITU) and are, therefore, allocated for primary use in most countries in the world, including the US. The ISM devices operating in this band have uses from heating body tissues, magnetic resonance imaging, radiometry for cancer detection and even thawing of frozen organs for transplant to industrial RF heaters for treating food, chemicals, packaging, textiles and construction materials. Microwave lamps are used to manufacture fiber optic cable, automobile glass and headlamps, no-wax floor tiles and to dry printing on containers. The National Air and Space Museum in Washington, D.C. is illuminated by microwave lighting. By far, however, the most important source of ISM interference in the band are commercial microwave ovens, of which there are over 100 million in use in the US alone.
- Military radar and radiolocation.
- Amateur radio.
- Other unlicensed communications technologies.

The ability to operate in the presence of these unwanted emissions should be a key characteristic of a 2.4 GHz unlicensed communications device. Communications systems developers often use spread spectrum technology to allow operation in the presence of interference.

Spread Spectrum Technology

Spread spectrum refers to a wideband radio frequency technique originally developed by the military for use in reliable, secure, mission-critical communications systems. Spread-spectrum is designed to trade 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 properly, a spread-spectrum signal looks like background noise. There are two types of spread spectrum radio: frequency hopping and direct sequence. The illustration in Figure 1 of how frequency hopping and direct sequence systems use the spectrum is more fully explained below.

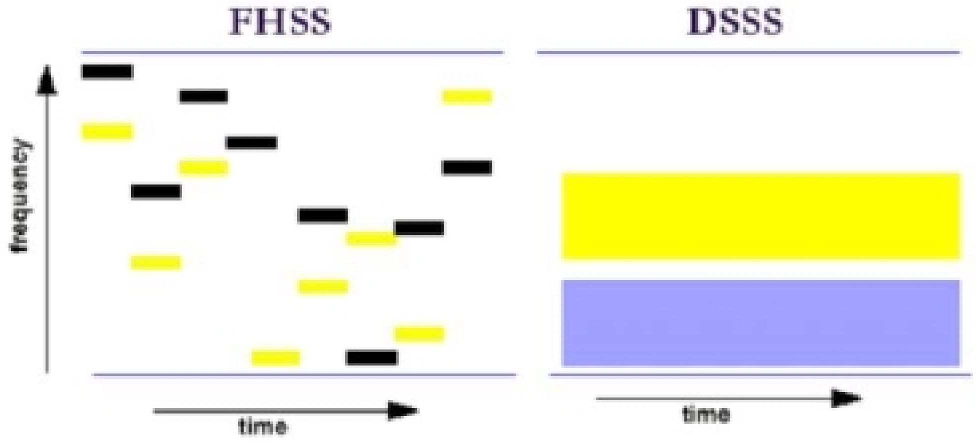

Figure 1: Spectrum Use by FHSS (Left) and DSSS (Right) Technologies

Frequency-Hopping Spread Spectrum Technology

The HomeRF and Bluetooth standards are both frequency hopping spread spectrum technologies.

Frequency-hopping spread-spectrum (FHSS) uses a narrowband carrier that changes frequency in a pattern known to both transmitter and receiver. Both the access point and client "hop" between frequencies based on the same pseudorandom pattern, transferring a piece of data during each hop. Properly synchronized, the net effect is to maintain a single logical channel. To an unintended receiver, FHSS appears to be short-duration impulse noise. In Figure 1 the FHSS side of the figure shows two different hopping sequences and how they use different, small slices of the spectrum for short periods of time.

Whenever interference corrupts the signal, the devices can resume their data transfer after the next hop to a new frequency that is clear. Bandwidth drops each time the device encounters a blocked frequency. However, interference does not break a connection. In the presence of interference, the connections do not fail and throughput will degrade gracefully. Adaptive hopping (avoiding frequencies that are known to be blocked) can be used to increase throughput.

The hopping pattern (frequencies and order in which they are used) and dwell time (time at each frequency) are restricted by most regulatory agencies. For example, for operation in the 2.4 GHz band in the US the FCC requires that 75 or more frequencies be used with a maximum dwell time of 400 milliseconds.

All FHSS products on the market allow users to deploy more than one logical channel in the same area. They accomplish this by implementing separate channels on different, pseudo-random, hopping sequences. Because there are a large number of possible sequences in the 2.4 GHz band, FHSS allows many non-overlapping channels to be deployed in the same space.

Direct-Sequence Spread Spectrum Technology

The IEEE 802.11b standard is a direct sequence spread spectrum technology.

DSSS transmitters spread the signal over a frequency band that is wider than required to accommodate the information signal by mapping each bit of data into a redundant bit pattern of “chips” known as a chipping code. The longer the chipping code used, the greater the probability that the original data can be recovered (and, of course, the more bandwidth required). Even if one or more bits in the chip are damaged during transmission, statistical techniques embedded in the radio can recover the original data without the need for retransmission. At the destination the chips are mapped back into a bit, recreating the original data. Transmitter and receiver must be synchronized to operate properly.

The ratio of chips per bit is called the "spreading ratio". A high spreading ratio increases the resistance of the signal to interference. To an unintended receiver, DSSS appears as low-power wideband noise and is rejected (ignored) by most narrowband receivers. A low spreading ratio increases the net bandwidth available to a user.

In practice, DSSS spreading ratios for wireless LANs are quite small. Virtually all manufacturers of 2.4 GHz products offer a spreading ratio of less than 20. The IEEE 802.11b standard specifies a spreading ratio of 8.

Several DSSS products in the market allow users to deploy more than one channel in the same area. They accomplish this by separating the 2.4 GHz band into several sub-bands, each of which contains an independent DSSS network. Because DSSS truly spreads across the spectrum, the number of independent (i.e. non-overlapping) channels in the 2.4 GHz band is small. The maximum number of independent channels for any DSSS implementation on the market is three. The DSSS portion of Figure 1 shows two separate DSSS channels accessing a wide bandwidth in a time static manner.

FHSS Devices are More Interference Immune than DSSS devices

One of the clear advantages that FHSS systems have over DSSS systems is their immunity to interference. With Figure 1, above, in mind, it is easy to understand the two attributes that make DSSS poor at rejecting outside interference.

1. DSSS products spread their transmission power thinly across the spectrum. Transmitted power in any specific segment of the band is low. As a result, low levels of interference can easily overpower the DSSS transmission. FHSS products, in contrast, use relatively high power in a narrow segment of the band for a short time. This allows the FHSS signal to overpower the interference in their segment of the band.
2. Multi-channel DSSS products use statically allocated pieces of the band. Interference in any significant piece of this allocated band will interfere with the transmission, possibly destroying it entirely.

All channels in an FHSS network hop around the entire 2.4 GHz band. Strong interference in a segment of the band may hamper some of the transmissions, but FHSS transmitters - being frequency agile - will use the remainder of the band effectively. Users may see a decrease in throughput but the network will continue to operate.

When interference occurs, it has a marked difference on the two different technologies. FHSS can overpower and/or hop around the interference, experiencing, at most, some limited degradation. The DSSS signal, on the other hand, can neither overpower nor avoid the interference. This is shown in Figure 2.

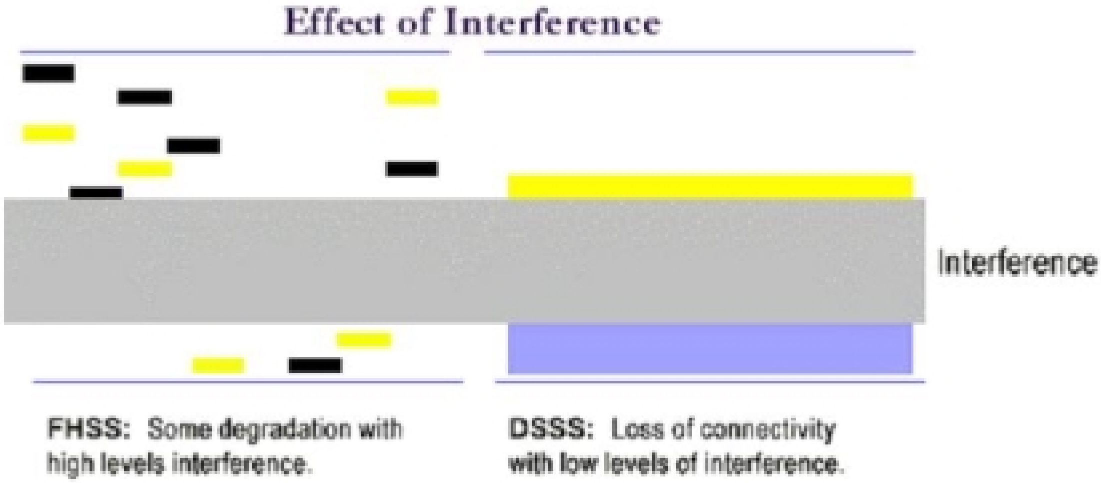
Effect of Interference
Interference
FHSS: Some degradation with high levels interference.
DSSS: Loss of connectivity with low levels of interference.

Figure 2: FHSS and DSSS interference coping strategies

All devices operating in the unlicensed 2.4 GHz spectrum band cause interference, to some extent, with all of the other devices in the band. By far the most problematic interference source for communications devices in the band, however, are consumer microwave ovens. As shown in Figure 1 DSSS products are assigned to a single specific frequency range. An interfering microwave oven can render the radio inoperable depending on the channel on which the DSSS unit is operating. Since DSSS units are frequency static, they have no mechanism for avoiding the interference from microwave ovens.

FHSS products have a built-in mechanism for avoiding microwave oven interference since they use the entire unlicensed band in a frequency agile manner. First, the interference occurs only when the carrier is on a channel that overlaps the microwave oven interference and is otherwise absent, so the fraction of link time that is actually affected is relatively low. Second, a smart system can adapt to the situation, realizing that the received signal often contains errors in a particular frequency range and decide to change the hopping sequence to avoid that range. As discussed below, the HomeRF frequency hopping technology employs a method of hopset adaptation that, in conformance with the FCC's frequency hopping rules, allows HomeRF systems to reduce their susceptibility to microwave oven, or other static, interference.

According to one analysis of this effect, "Current […] DSSS products on the market occupy 22 MHz of bandwidth at the first nulls of the main signal lobe. On the other hand, FHSS operates within 1 MHz of bandwidth as a consequence of an FCC mandate. As a result, the circuit design for the DSSS products requires very good passband characteristics for the 22 MHz width."[6] This same analysis holds even though the FCC has now allowed FHSS to operate with channel widths as wide as 5 MHz. Thus, DSSS networks may be crippled by outside interference that FHSS networks, because of their frequency-agile nature, can simply "hop around" without experiencing significant degradation. FHSS products spend only milliseconds at each frequency. Noise on any given frequency will typically be absent after hopping to another frequency.

Concern about the interference robustness of the DSSS 802.11b technology is the focus of a great deal of current industry activity.[7] Anticipating co-located deployments of 802.11b and Bluetooth devices, the wireless LAN industry has started to analyze the impact that a Bluetooth system will have on the operation of an 802.11b

[6] DSSS vs. FHSS, The Skinny. Solectek. http://www.solectek.com/tech-center/index.html

[7] The Wall Street Journal article mentioned in footnote **Error! Bookmark not defined.** begins with a story of a family that found that its 2.4 GHz cordless phone system was blocking operation of its 802.11b network.

device.[8] A co-existence group (IEEE 802.15, Task Group 2, the Coexistence Task Group) has been formed to address precisely this issue. Analyses have already started to appear which point to the extreme sensitivity of the 802.11b physical layer to external interference, especially interference from Bluetooth. According to one of these,

> *"In the scenarios measured, even Wi-Fi™ stations less than five to seven meters (free space) from their Access Point suffer more than 25 percent degradation in throughput. This degradation exceeds 50 percent by the 30-meter mark. Within an office environment with cubicles, the range associated with each throughput level will be significantly reduced. When cubicles must be penetrated, Wi-Fi™ loses almost one-third its expected throughput within the first couple of meters. Erosion of performance in excess of 50 percent takes place with stations less than eight meters from their Access Point."*[9]

Another study examined the impact of 100 mW Bluetooth devices on the throughput of an 802.11b system in various configurations.[10] As shown in Figure 3, a Bluetooth transmitter within 10 cm of an 802.11b receiver was found to result in at least a 60% reduction in throughput, and made communications between 802.11b devices impossible at distances beyond 150 feet. For a Bluetooth transmitter 10 meters from the 802.11b receiver, the throughput reduction was less pronounced, but still reached 45% with 250 feet between the 802.11b devices.

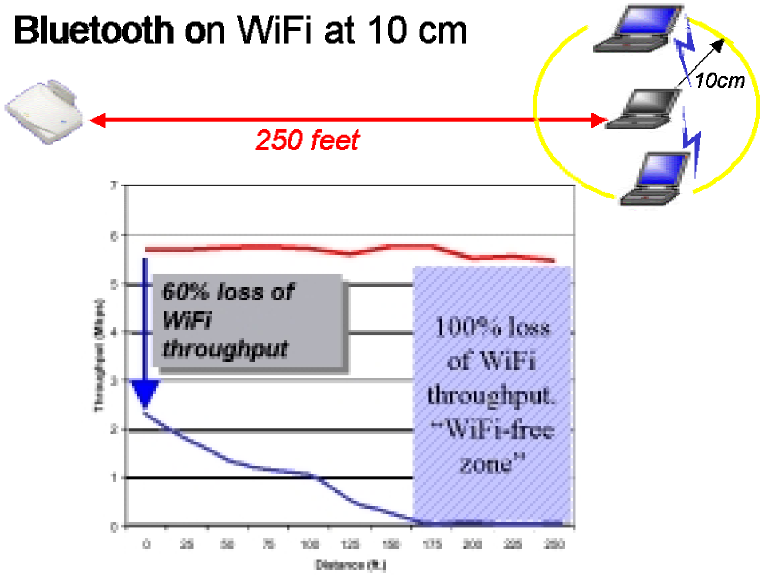

Figure 3: Throughput Degradation of 802.11b System Due to a Nearby Bluetooth Transmitter

As demonstrated by these analyses, the interference situation is severe enough to lead to the consideration of some very radical solutions, including the banning of Bluetooth devices in areas where 802.11b devices are likely to be in use.[11] Another attempt to solve this problem would require the FCC to change its rules for the 2.4 GHz band, and, in fact, a petition to do this was filed with the FCC on October 2, 2000.[12]

[8] The concern with Bluetooth, specifically, is related to the very high hopping rate of Bluetooth. While HomeRF changes frequency at most 100 times per second, Bluetooth will change frequency 1600 times per second. Therefore, in any given time interval, it is more likely that a Bluetooth device will transmit within an 802.11b spectrum range than that a HomeRF device will do so.

[9] "Wi-Fi™ (802.11b)and Bluetooth Simultaneous Operation: Characterizing the Problem", Mobilian Corporation. http://www.mobilian.com/documents/WPSG_2.pdf. Note that Wi-Fi™ is the marking name for the 802.11b technical standard.

[10] "Wi-Fi (IEEE 802.11b) and Bluetooth - Coexistence Issues and Solutions for the 2.4 GHz ISM Band", Texas Instruments, February 2001. http://www.ti.com/sc/wirelessnetworking.

[11] The interference between Bluetooth and 802.11b is probably not, on average, as severe as that reported in a recent article in which, in the presence of a Bluetooth transmitter, "the 802.11b connection was beaten down to about one-thirtieth of its normal speed." However, articles like this demonstrate the concern that users have over this kind of interference. Stephen Manes, Forbes Magazine, Dec. 11, 2000, http://www.forbes.com/forbes/2000/1211/6615224a.html.

[12] "Amendment of Part 15 of the Commission's Rules Regarding Spread Spectrum Devices, ET Docket Number 99-231, filed by multiple parties, October 2, 2000.

Proxim has also performed tests to study the sensitivity of the 802.11b technology to Bluetooth interference. In one of those tests, a Bluetooth interferer was placed 12 feet away from either an 802.11b node or a HomeRF node, and the resulting impact on data throughput of the victim node was measured. The test configuration is shown in Figure 4.

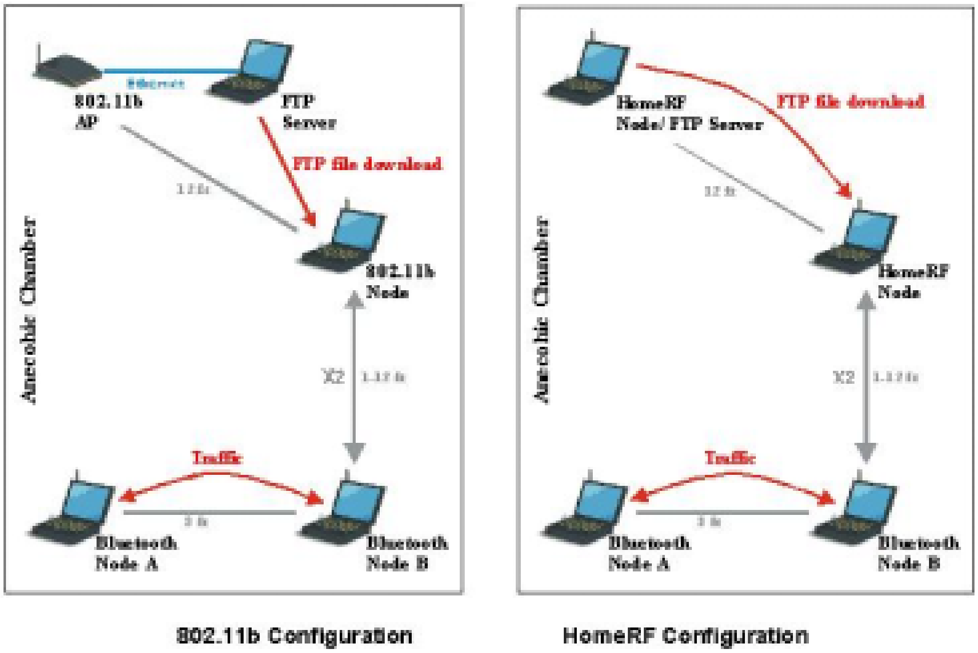

802.11b Configuration

HomeRF Configuration

Figure 4: Proxim Interference Test Configuration

These tests are consistent with the results of the Mobilian and Texas Instruments tests, referenced above. The 802.11b device, located 12 feet (about 4 meters) from its access point had its throughput degraded by 25% – 30% by the nearby Bluetooth device. On the other hand, Proxim’s tests showed that a HomeRF device in the same situation would experience a throughput reduction of only 10%.

It is reasonable to wonder whether it is only HomeRF’s frequency hopping physical that leads to this interference robustness. The answer is that it is not. HomeRF contains additional features at higher layers in the standard that make HomeRF products more immune to 2.4 GHz interference. Even though Bluetooth uses a frequency hopping physical layer, it is not as robust against interference as is HomeRF because it lacks the kinds of features discussed below.

HomeRF is designed to be extremely robust against 2.4 GHz interference

Though both Bluetooth and HomeRF are FHSS systems, they are not equally robust against 2.4 GHz interference. The reason for this is that higher layers of the protocol (beyond the frequency hopping physical layer) also add to HomeRF’s interference immunity.

The studies of Bluetooth interference on 802.11b performance cited above also show the effect of interference on Bluetooth connections. The Texas Instruments test, for example, found the results shown in Figure 5 for the case of a Bluetooth receiver within 10 cm of an 802.11b transmitter. The nearby 802.11b device caused a throughput degradation of at least 40% in the Bluetooth system. When the spacing between the 802.11 transmitter and the Bluetooth receiver was increased to 10 meters, the throughput degradation was reduced to about 10%, independent of the spacing between the Bluetooth nodes.

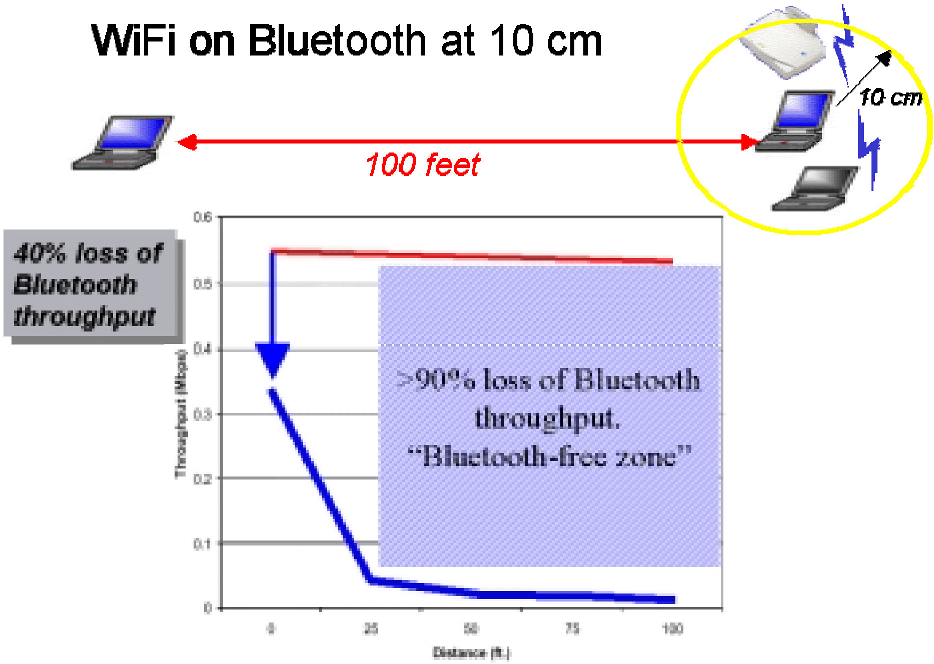

Figure 5: Throughput Degradation of a Bluetooth System Due to a Nearby 802.11b Transmitter

While Bluetooth is a frequency hopping system with the inherent interference avoidance mechanisms discussed above,[13] Bluetooth has also been designed to be extremely simplistic. At protocol layers above the frequency hopping physical layer it has only minimal mechanisms to cope with interference that it receives. The HomeRF frequency hopping system, on the other hand, contains MAC level retry mechanisms both for the data and isochronous voice connections. Therefore, while Bluetooth will operate well as the short-range "cable replacement" system for which it has been designed, it will not be very reliable for wireless LAN applications, especially those involving voice or other real-time media applications.

The HomeRF standard leverages the natural strength of FHSS to avoid interference and combines it with strong additional mechanisms to combat interference. Two of these mechanisms are "hopset adaptation" and "subframe hopping with retries" to ensure the integrity of voice communications.

HomeRF's Hopset Adaptation Gives High Reliability to the Retry Mechanism

Hopset adaptation is used to minimize the impact of long-term, static interferers. An example of such an interferer is a microwave oven. "The ovens emit up to 750 W of power at 50 pulses per second with a radiation cycle of about 10 ms… The emitted radiation sweeps from 2.4 to 2.45 GHz and remains stable for a short period at 2.45 GHz. Even though the units are shielded, a good amount of energy can still interfere with the transmission from wireless LANs."[14] Under the FCC's frequency hopping rules for the 2.4 GHz band hopping systems using hopping channels of at least 1 MHz must occupy at least 75 MHz of the 83.5 MHz available in the band. Therefore, there is no opportunity to avoid entirely the frequency range occupied by a wideband static interferer. HomeRF takes advantage of the fact that communications that are blocked by interference can be re-transmitted on the next hop. HomeRF's hopset adaptation mechanism ensures that *two adjacent hops* will not both be within a frequency range where interference has been identified.

This mechanism works as follows. When an "interference range" is identified, the hopset is examined to find out if two consecutive hops are both within the range. If such a pair of hops is located, an attempt is made to switch

[13] A point not lost on the authors of the Texas Instruments paper, who write, "Frequency hopping devices have an inherent level of robustness due to the fact that they do not continually transmit at the same frequency. The changing of the transmit center frequency or hopping means that the probability of colliding with the transmission of another device at any particular time is very small."

[14] Bing, Benny, "High-Speed Wireless ATM and LANs", Artech House, 2000, page 59.

one of those hops with a hop that is outside the interference range. Therefore, though the full set of 75 hops is still used, the hopping sequence now virtually guarantees that an “interfered with” hop will be followed by a hop without interference. This technique is very powerful, and leads to no consecutive “bad” hops in the presence of an interference of up to 31 MHz. This is shown in Figure 6.

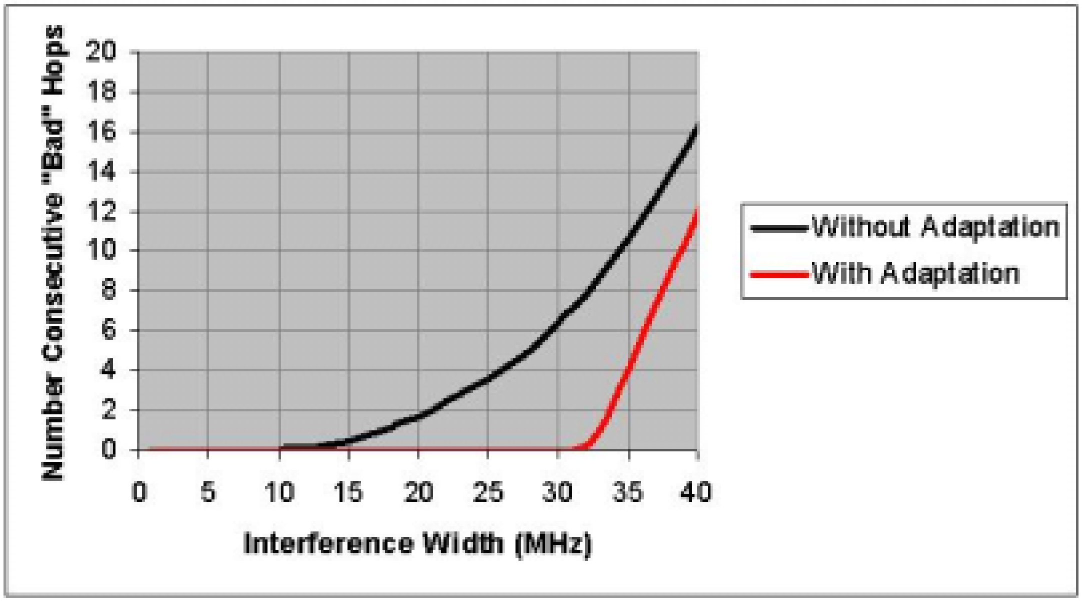

Figure 6: Performance of the HomeRF Hopset Adaptation Algorithm

The power of this algorithm is that in the presence of a wideband, static interferer (like a microwave oven or DSSS wireless LAN, for example) HomeRF can virtually guarantee that a hop that receives interference will be followed by a hop that is free from interference. Since HomeRF has the usual retry mechanism for data packets, as well as a unique retry mechanism for voice packets (see the discussion below), this leads to very robust HomeRF communications in the presence of this type of interference.

HomeRF takes special care to provide low bit error rate voice connections

For voice applications, HomeRF has improved upon the DECT specification, on which the voice component of HomeRF is based, to offer a retransmit mechanism for degraded voice packets and a subframe hopping mechanism to increase the interference immunity. That is, while HomeRF is based on a 20 ms frame structure (with one hop every frame), HomeRF moves to a 10 ms subframe structure (with one hop every subframe) whenever there are active voice connections. This structure is illustrated in Figure 7.

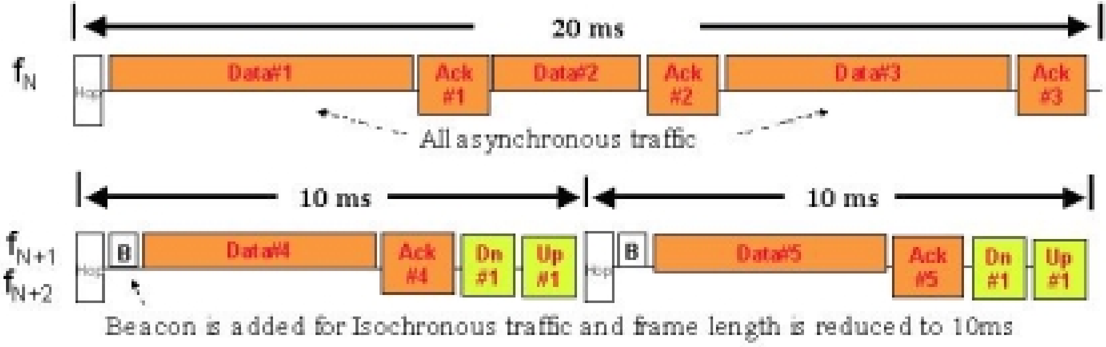

Figure 7: Basic overview of the HomeRF frame structure

The upper part of Figure 7 shows what the HomeRF frame looks like when only data is present on the network. The frame is 20 ms long, and all of the data accesses the medium using a contention-based access protocol. The lower part of the frame shows the frame structure when a voice communication is present. In this case, the part of the frame available for contention-based access is reduced, and part of the frame is set aside for TDMA access for the voice. The voice connection is full duplex using time division duplex (TDD), which explains the Dn#1 and Up#1 areas in the lower part of Figure 7. These are the downlink and uplink parts of the voice communication. Use of TDMA for the voice access methodology explains how HomeRF provides for low latency voice communications on the network. Low bit error rates are provided by having a robust, frequency hopping, physical layer, and by adding a unique retry mechanism for the voice frames. Voice packets that are blocked due to interference in one subframe are resent in the next subframe, less than 10 ms later. This is illustrated in Figure 8.

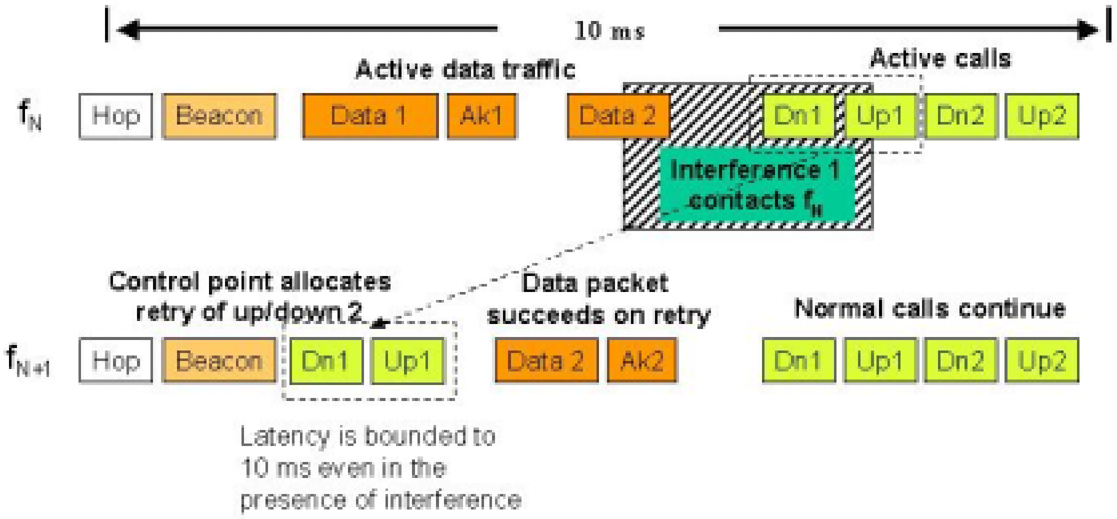
10 ms
Active data traffic
Active calls
f_N
Hop
Beacon
Data 1
Ak1
Data 2
Dn1
Up1
Dn2
Up2
Interference 1 contacts f_N
Control point allocates retry of up/down 2
Data packet succeeds on retry
Normal calls continue
f_N+1
Hop
Beacon
Dn1
Up1
Data 2
Ak2
Dn1
Up1
Dn2
Up2
Latency is bounded to 10 ms even in the presence of interference

Figure 8: HomeRF provides a unique mechanism to retry voice packets

As shown in Figure 8, if interference appears in the frequency range and time such that HomeRF voice packets experience interference, those packets can be sent again no more than 10 ms later. Since the next hop is virtually guaranteed to be outside the interference region due to the hopset adaptation algorithm described above, HomeRF voice has not only low latency, but low bit error rate characteristics.

Frequency hopping plus hopset adaptation lead to very low bit error rates in the presence of typical 2.4 GHz interference

HomeRF was designed with a clear understanding of the 2.4 GHz environment. The fundamental frequency hopping physical layer together with the addition of HomeRF specific features leads to very robust interference performance. An estimate of the interference immunity of HomeRF in the presence of a microwave oven is shown in Table 1. The result, as shown, is that the expected bit error rate is significantly less than 1%.

Probability of a hit by microwave oven interference (Assumed characteristics: interference covers 20% of the band for 50% of the time.)	Raw P_{hit} = 10% (This is the extreme failure rate that would be experienced by some WLAN systems.)
HomeRF time/frequency diversity reduces this probability (Results in acceptable voice quality)	Independent trials result in $P_{2\ hits} = P_{hit} \times P_{hit}$ = 1%

HomeRF also uses hopset adaptation to further reduce this probability in the presence of persistent interferers (Results in excellent voice quality)	The probability of a second hop into the interference zone is reduced further. N_i = channels with interference N_T = total channels $P_{2\ hits} \sim P_{hit}^2 \times (N_i/N_T) << 1\%$

Table 1: Voice Packet Failures in the Presence of a Microwave Oven

The difference in interference robustness between HomeRF and the other main frequency hopping technology, Bluetooth, is that Bluetooth was not designed to operate as a wireless LAN in the home environment. A typical Bluetooth scenario might be as a connection between a cellular phone and a wireless headset used in a car. In that case, the Bluetooth devices can be expected to be very close together, and it is unlikely that an interferer like a microwave oven or a wireless LAN network would be present. Therefore, Bluetooth was not designed at its upper layers to be robust against this type of interference.

Conclusion

Unlicensed technologies are a perfect fit for consumer devices. The growth in the PC market, broadband deployment, and the popularity of the Internet are causing an explosion in the number of wireless LAN products that use unlicensed spectrum. Though it is not always clear to the casual observer, these devices are not interchangeable. Many choices are made in the creation of these products that have an impact on how these devices will function in an inherently harsh and noisy environment like the unlicensed 2.4 GHz band. Of the three major technologies available for this band, only HomeRF is designed with a frequency agile physical layer and robust upper layer protocols to combat 2.4 GHz interference. This is what makes HomeRF the ideal wireless LAN technology for the home environment.