"You Believe You Understand What You Think I Said…"

The Truth About 802.11 Signal And Noise Metrics

A Discussion Clarifying Often-Misused 802.11 WLAN Terminologies

by Joshua Bardwell

It's not uncommon for 802.11 WLAN network administrators, designers, and engineers to talk about RF signal strength, signal-to-noise ratio, and other measurement metrics related to radio signal propagation. Unfortunately, it's also not uncommon for someone to use terminology that they don't fully understand. When their colleagues try to make sense out of the conversation they, too, often interpret some technical terms based on their own incomplete understanding. The result is a double obfuscation of the true meaning and significance of some very important aspects of RF design.

Due to a number of factors, there is a great deal of inconsistency and confusion regarding the definition of common 802.11 terms, such as "signal strength," "signal to noise ratio," "signal quality," and "receive sensitivity." This present paper, "You Believe You Understand What You Think I Said – The Truth About 802.11 Signal And Noise Metrics", describes standard definitions for these terms using definitive sources, such as the IEEE 802.11 standards for Wi-Fi wireless networking.

Table of Contents

"You Believe You Understand What You Think I Said…"
The Truth About 802.11 Signal and Noise Metrics

About The Author

Joshua Bardwell holds a BS in Computer Science (Networking Specialization) with a Minor in Technical Communications from the Georgia Institute of Technology. From 1995 to 2000, he was a Technical Instructor and Network Engineer with Optimized Engineering Corporation and a certified (CNX) instructor with Network General's Sniffer University. In 2000, Optimized Engineering partnered with WildPackets, where he specialized in wireless and wired protocol analysis training, courseware development, and consulting. In 2004, after the expiration of the OEC/WildPackets partnership, he continues to work with Optimized Engineering (now known as Connect802) as an RF engineer.

Since 1995, Joshua has continuously taught courses on diverse areas of protocol analysis, including: Ethernet, Token Ring, ATM, TCP/IP, and 802.11. Since 2001, he has focused primarily on 802.11 wireless LAN analysis, including acting as the lead developer for WildPackets' 802.11 analysis course and acting as primary author for the CWNA Study Guide, 3[d] edition.

Connect802 Corporation
www.Connect802.com
(925) 552-0802

Turnkey Wi-Fi Wireless Networking
Virtual Site Survey Services

Executive Summary

Due to a number of factors, there is a great deal of inconsistency and confusion regarding the definition of common 802.11 terms, such as "signal strength," "signal to noise ratio," "signal quality," and "receive sensitivity." This paper describes standard definitions for these terms using definitive sources such as the 802.11 standard. Its conclusions are as follows:

- "Signal strength" is defined in 802.11 as Received Signal Strength Indicator (RSSI). RSSI is intended to be used as a relative value within the chipset. It is not associated with any particular mW scale and is not required to be of any particular accuracy or precision. Therefore, the signal strength numbers reported by an 802.11 card will probably not be consistent between two vendors, and should not be assumed to be particularly accurate or precise.

- "Signal quality" is defined very briefly in 802.11. Common definitions have arisen, but they are usually incorrect. The correct definition hinges on the term, "PN code correlation strength," which is a measure of the match (correlation) between the incoming DSSS signal and an ideal DSSS signal. The common equation of "signal quality" and "signal to noise ratio" is incorrect.

- "Signal to noise ratio" is a general term that is used in a novel way by 802.11 administrators. Most usages of the term refer to the strength of the signal relative to thermal noise within a circuit, but 802.11 administrators use the term to refer to the strength of the signal at the receive antenna relative to the ambient, non-802.11 RF energy at the same frequency as the signal. While this definition isn't wrong, per se, it may lead to confusion when 802.11 administrators communicate with engineers who are using the more traditional definition.

- "Receive sensitivity" refers to the weakest power level that the card's internal thermal noise will allow it to receive. It is unrelated to the ambient, non-802.11 RF energy in the environment.

Section 1: Introduction

"You believe you understand what you think I said..."

The title of this paper is borrowed from the popular quotation, "I know you believe you understand what you think I said, but I'm not sure you realize that what you heard was not what I meant." The essence of this quote is that you might think you understand a certain topic, but without a common definition of terms, your understanding of the topic is flawed, and you won't be able to communicate about the topic.

This is particularly relevant to 802.11, which combines RF communication at the physical layer with a data link layer that is much more complex and sophisticated than traditional, wired data links like Ethernet. This places 802.11 LAN administrators in a difficult position, since they are doubly unfamiliar: both with the quirks of RF propagation and with the new data link layer. This unfamiliarity has led 802.11 administrators to use terms such as, "signal strength," "signal quality," and "signal to noise ratio," in an inconsistent and usually colloquial manner. While some of these terms are defined in the 802.11 standard, others come from the realm of RF engineering or electrical engineering. This document will explore these and other related terms and will provide standard definitions, backed up by references.

Are You Using The Right Terminology?

It's natural for an 802.11 administrator to want to communicate about 802.11. In lieu of well-defined, standard terms, it's equally natural to use common terms in whatever way is convenient. Confusion arises when common terms are used with inconsistent definitions, or when a misunderstanding of common terms, such as "signal strength," leads to false conclusions. You might hear an 802.11 administrator say:

- "My Cisco ACU says I have 100% signal strength and -10 dBm, but my Symbol client manager says I have 80% signal strength and -50 dBm. Which is right?!"
- "Let's design the network with 10 dB signal to noise ratio so that if the signal loses strength, there will be some buffer before the clients lose signal."
- "This 802.11b chipset requires 20 dB signal to noise ratio to achieve 11 Mbps. Even if signal strength is high, you won't get 11 Mbps without 20 dB signal to noise ratio."
- "Signal quality is the same thing as signal to noise ratio."

Although these are common statements, all of them demonstrate misunderstandings about the definitions of 802.11-related terms, as we will see in this paper.

The Consequences Of Inconsistent Terminology

In the 1970s, the I-95 interstate highway was extended between Wilmington, Delaware and Philadelphia, Pennsylvania. As part of the project, a bridge had to be built to allow freight ships to access the port of Philadelphia. As one crew worked north from Wilmington, and the other crew worked south from Philadelphia, the teams consistently measured to ensure that the two halves of the bridge would be properly aligned when they met. But as the ends of the halves got closer together, it became apparent that they were completely misaligned!

It seems that the instruction had been to align the "left side" of each half of the bridge with each other. The engineer on each side of the bridge used his own "left" instead of a common "left," with the result that the eastern side of the northern half lined up perfectly with the western side of the southern half! To this day, that bridge has a dog-leg turn in the middle, where this misalignment was compensated for.

This story demonstrates the potentially dire consequences of inconsistent definition of terms. The same consequences could arise if an 802.11 administrator, not understanding the implications of a term like "signal strength," attempted to design a WLAN using an incorrect definition.

Section 2: Signal Strength

Most 802.11 analysis tools and vendors' client management utilities provide a representation of signal strength (see figure 2.1 for examples). Four units of measurement are used to represent RF signal strength in 802.11. These are: mW (milliwatts), dBm ("dB"-milliwatts), RSSI (Received Signal Strength Indicator), and a percentage measurement. All of these units are related to each other, some more closely than others, and it's possible to convert from one unit to another. Let's explore how these units are related to each other before we discuss the most correct definition of "signal strength."
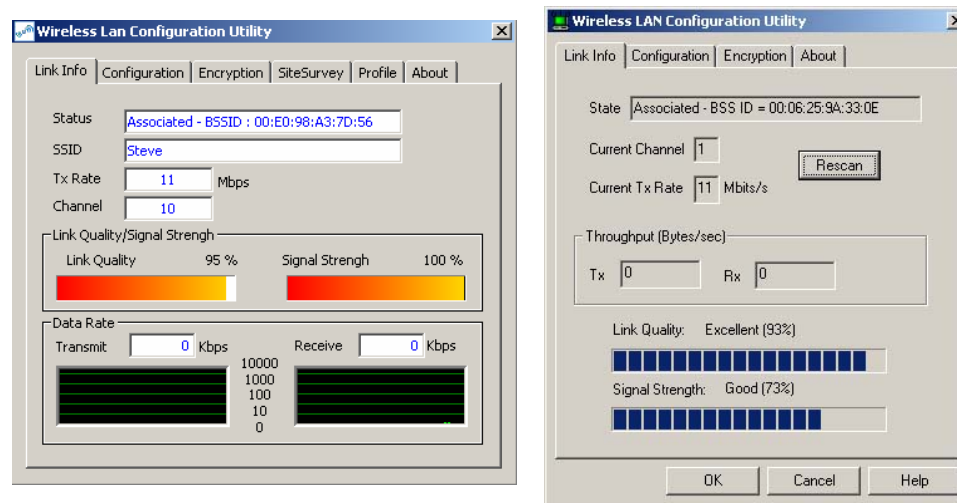
Figure 2.1 – Sample client utilities showing signal strength.

The mW And dBm Units Of Measure

The first two units to consider are the mW and the dBm (pronounced "dee-bee-em" or spoken as "dee-bee milliwatts"). Although these are not the most common units in 802.11, we discuss them first because they are the most basic. Just like a pound is a basic unit for measuring weight, a watt is a basic unit for measuring energy (and, in keeping with metric conventions, a mW is one one-thousandth of a watt). An electrical engineer or physicist could explain "energy" in more detail, but it's not necessary to go to that level of detail in this paper. Just for your perspective, a typical wireless access point has a power output of between 1 and 100 mW, while a typical wireless client has a power output of between 1 and 30 mW.

It turns out that measuring RF energy in mW units is not always convenient. This is due, in part, to the fact that signal strength does not fade in a linear manner, but inversely as the square of the distance. This means that if you are a particular distance from an access point and you measure the signal power, then you move twice as far away, the signal power will have decreased by a factor of four.

You double the distance and the signal decreases by $1/4^{th}$ the power, hence the "inverse square law." This relationship can be represented mathematically by the logarithm function. Since milliwatts directly measure the power of the RF energy, they will decrease logarithmically along with the power. This is inconvenient because logarithmic functions get very small, very fast. For example, an 802.11 RF signal transmitted at a power of 100 mW might be received by a client at a power of 0.000 000 01 mW. Not only does that make the math of calculating RF link budget slightly more complicated, but it's difficult to say all of those zeroes when talking about RF power! The fact that exponential measurements are involved in signal strength measurement is one reason why the use of a logarithmic scale of measurement was developed as an equivalent, but alternative way of representing RF power.

The "dBm" is a logarithmic measurement of signal strength. Since it is logarithmic, just like the power of the RF signal, as the RF signal's strength changes (logarithmically), the dBm value changes linearly. To put it more generally, if you measure a quantity that changes logarithmically (RF power) with a linear unit (mW), the unit will change logarithmically, which is inconvenient. If you measure a quantity that changes logarithmically with a logarithmic unit (dBm), the unit will change linearly, which is more convenient. If this doesn't one-hundred percent make sense now, hopefully it will after we do some examples later.

dBm values can be converted to and from mW values. Just like miles and kilometers can be converted directly, so can mW and dBm. To perform this conversion, the mW measurement is converted to a base-10 logarithm. It turns out that the logarithm values are quite small and so convention multiplies this value by 10 with the resulting value being called "dBm." Here are some examples to help clarify this relationship (since the logarithm is the inverse of the exponent, exponential relationships are shown in parentheses after each logarithm as another way of expressing the logarithm):

100 mW	log 100 = 2 (10^2=100)	20 dBm = 100 mW
50 mW	log 50 = 1.698 (10^1.698 = 50)	15.9 dBm = 50 mW
25 mW	log 25 = 1.397 (10^1.397 = 25)	13.9 dBm = 25 mW
13 mW	log 13 = 1.113 (10^1.113 = 13)	11.1 dBm = 13 mW

You can confirm these relationships with your scientific calculator. Notice that each time the actual mW power level becomes half as great, the dBm measurement goes down by (roughly) 3 dBm. It's convenient shorthand to remember that a decrease of 3 dBm yields roughly half the original signal strength and, conversely, an increase of 3 dBm yields roughly twice the original value.

Of course, power is always a positive quantity. You can't have "negative energy" (unless you're studying quantum mechanics and virtual particles!) so the mW measurement will always be something greater than zero. You can, however, have very small values—much less than 1. When representing a fraction less than 1

(but greater than zero) it can be shown that the corresponding logarithmic value is negative. You can confirm the following relationships on your calculator if you desire:

1 mW	log 1 = 0 (10^0=1)	0 dBm = 1mW
.5 mW	log .5 = -0.3010 (10^-0.3010 = .5)	-3.01 dBm = .5 mW
.25 mW	log .25 = -0.602 (10^-0.602=.25)	-6.02 dBm = .25 mW
.13 mW	log .13 = -0.886 (10^-0.886=.13)	-8.86 dBm = .13 mW

Notice, again, that a decrease of roughly 3 dBm yields a change of roughly half in the mW value. As you can see, it's much easier to say, and write, "-96 dBm" than have to figure out where all the zeros are and whether you're talking about "pico-watts" or "fempto-watts", and that's why we talk about dBm: it's a much more useful way of measuring signal strength at very low levels.

We can now say, "An 802.11 NIC transmits power at roughly 20 dBm (100 mW)." You should realize that while it's reasonable to talk about 20 dBm as being 100 mW it would be cumbersome to talk about –96 dBm as being .000 000 000 251 1 mW (which is, incidentally, about the weakest signal that the most sensitive 802.11 cards can receive). You should realize that convenience and ease-of-understanding are two fundamental reasons why the dBm metric is used for RF signal strength, rather than mW.

The actual formula used for the conversion is:

$$\textbf{dBm = log (mW) * 10}$$
$$\textbf{mW = 10\textasciicircum(dBm/10)}$$

Received Signal Strength Indicator (RSSI)

The IEEE 802.11 standard defines a mechanism by which RF energy is to be measured by the circuitry on a wireless NIC. In 802.11b, g, and a, this numeric value is an integer with an allowable range of 0-255 (a 1-byte value) called the Received Signal Strength Indicator (RSSI). 802.11 doesn't require that a chipset vendor use all 255 values, so each vendor's 802.11 NIC will have a specific maximum RSSI value ("RSSI_Max"). For example, Cisco chooses to measure 101 separate values for RF energy and their RSSI_Max is 100. Symbol uses an RSSI_Max value of 31. The Atheros chipset uses an RSSI_Max value of 60.[1] The RF energy level reported by a particular vendor's NIC will range between 0 and RSSI_Max.

Notice that nothing has been said here about measurement of RF energy in dBm or mW. RSSI is an arbitrary integer value, defined in the 802.11 standard and intended for use internally by the physical and data link layers (the hardware in the card and its drivers). For example, when an adapter wants to transmit a packet

---

[1] These values are based on the vendors' 802.11b chipsets circa 2002. They may not be current and are intended only as examples.

it must be able to detect whether or not the channel is clear (i.e. nobody else is transmitting). If the RSSI is below some very low threshold then the chipset decides that the channel is clear. This is the "Clear Channel Threshold" and some particular RSSI value is associated with it. When an 802.11 client is moving, there may come a point when the signal level received from the station's current access point drops to a somewhat low value (because the client is moving away from the access point.) The client may decide to roam to an access point with a higher signal strength. The signal strength below which the client will attempt to roam is called the "Roaming Threshold," and some intermediate (but low) RSSI value is associated with it. Different vendors use different signal levels for the Clear Channel Threshold and the Roaming Threshold and, moreover, the RSSI value that represents these thresholds differs from vendor to vendor because different RSSI_Max values are implemented.

RSSI In The 802.11 Standard

Although the term, "signal strength" is commonly used by those who work with 802.11, the 802.11 standard does not define "signal strength" specifically. Rather, it defines only RSSI. RSSI is defined as a parameter of a function called *PMD_RSSI.indicate(RSSI)*, which allows the PHY to report RSSI values. Below are the relevant sections of 802.11. Underlining is added to emphasize particular sentences. We will discuss the implications of these sentences after the quotes.

Section 14 (802.11 2.4 GHz FHSS)

14.2.3.2 RXVECTOR RSSI

The receive signal strength indicator (RSSI) is an optional parameter that has a value of 0 through RSSI Max. This parameter is a measure by the PHY sublayer of the energy observed at the antenna used to receive the current PPDU. RSSI shall be measured between the beginning of the start frame delimiter (SFD) and the end of the PLCP header error check (HEC). RSSI is intended to be used in a relative manner. Absolute accuracy of the RSSI reading is not specified.

14.5.5.8 PMD_RSSI.indicate

14.5.5.8.1 Function

This primitive transfers a receiver signal strength indication of the physical medium from the PMD sublayer to the PLCP sublayer. This value will be used by the PLCP to perform any diversity or clear channel assessment functions required by the PLCP or other sublayers.

14.5.5.8.2 Semantics of the service primitive

The primitive shall provide the following parameter: PMD_RSSI.indicate (STRENGTH). <u>The STRENGTH parameter can be a value from 0 to 15.</u> This parameter is an indication by the PMD sublayer of the magnitude of the energy observed at the selected antenna. This reported value is used to generate the RSSI term in the PHY-RXSTART.ind(RXVECTOR) primitive and might also be used by any diversity function. <u>Since RSSI is only used in a relative manner by the MAC sublayer, this parameter is defined to have no more than 16 values, ranging from 0 through RSSI_Max. The value zero is the weakest signal strength, while RSSI_Max is the strongest signal strength.</u>

Section 15 (802.11/802.11b 2.4 GHz DSSS)

15.4.5.10 PMD_RSSI.indicate

15.4.5.10.1 Function
This optional primitive, which is generated by the PMD sublayer, provides to the PLCP and MAC entity the received signal strength.

15.4.5.10.2 Semantics of the service primitive
The primitive shall provide the following parameter: PMD_RSSI.indicate(RSSI). <u>The RSSI shall be a measure of the RF energy received by the DSSS PHY. RSSI indications of up to 8 bits (256 levels) are supported.</u>

802.11a (Section 17) and 802.11g (Section 19) PHYs mirror section 15 (802.11/802.11b 2.4 GHz DSSS). Several important implications should be drawn from these definitions.

Precision In RSSI Measurements

The concept of "granularity" or "precision" is important to consider here. Since the RSSI value is an integer, it must increase or decrease in integer steps. The RSSI will vary between 0 and RSSI_Max, but remember that each vendor can define its own RSSI_Max. For example, Symbol provides 32 separate "steps" (RSSI_Max = 31), Cisco provides 101 (RSSI_Max = 100). Whatever range of actual energy is being measured, it must be divided into the number of integer steps provided by the RSSI range. Therefore, if RSSI changes by 1, it means that the power level changed by some fraction of the measured power range. A vendor that uses an RSSI scale with more values can measure received signal strength with higher precision than a vendor that uses an RSSI scale with fewer values.

Although it might intuitively seem like higher precision is better, bear in mind that 802.11 cards don't really need to be high-precision devices. Ask yourself, "how does an 802.11 card use signal strength?" Answers include: to determine whether another station is transmitting or not; to determine which antenna is the best antenna to use for diversity; to determine whether one access point is stronger than another. In every case, adequate results can be achieved with low precision. For example, if two access points are relatively close in signal strength, then a low-precision card might end up choosing the one that was slightly weaker—technically, a "wrong" choice. But is it really wrong? After all, if the APs are relatively close in signal strength, then does it really matter if the station

picks one that is just a little bit weaker? A vendor who decides that more precision will improve performance can choose to use a larger RSSI scale (up to 256 values), but most, if not all, vendors have chosen an RSSI scale of 100 values or less.

The unspecified and usually-low precision of 802.11 cards' RSSI scale means that the WLAN administrator should not assume that the "signal strength" as reported by an 802.11 card is a reliable method of determining whether one signal is stronger than another. If the signals are relatively close in power, the 802.11 card will probably say that they are the same strength. This is only acceptable if you are trying to model the card's behavior in an environment, since you're seeing the same power measurements that the card will use when it makes decisions about roaming, data rate, and so on.

Under no circumstances should you compare the signal strength measurements of cards using two different chipsets. All signal strength measurements in 802.11 are based on the RSSI, but 802.11 doesn't mandate how RSSI should be calculated, so different vendors will almost certainly measure it differently.

The Choice Of A Suitable Energy Range For Measurement

Notice that the RSSI parameter is "intended to be used in a relative manner." At no time does 802.11 require any specific relationship between RSSI and dBm or mW. This means that it is entirely up to the vendor to choose the power level (mW or dBm) that corresponds to the lower and upper bounds of the RSSI scale (RSSI = 0 and RSSI = RSSI_Max). Of course, each vendor does this differently, but some general statements can be made.

As you can confirm with the dBm-to-mW formulas given above, as the power varies between 1 mW and 100 mW, the dBm value changes by 20 units (0 dBm through 20 dBm). On the other hand, the power varies by less than 1 mW as the dBm value plunges 100 dB from 0 dBm (= 1 mW) to -100 dBm. Bear in mind that, realistically, stations will nearly always receive signals below 1 mW, unless they are very close (within a few feet) of the access point. Therefore, the power levels below 1 mW are critical to the operation of an 802.11 chipset. Those power levels are where data rates have to change or stations have to reassociate with new access points.

For example, consider that the difference between -85 dBm and -95 dBm—a difference of approximately 0.000 000 003 mW—might mean the difference between receiving a packet at 11 Mbps and totally losing connection to the access point! By contrast, the relatively enormous difference between 0.001 mW and 0.01 mW received signal power is inconsequential, since at 0.001 mW, the signal is more than strong enough to receive the packet at the highest data rate possible. Increasing power so that the station receives 0.01 mW doesn't improve things, because things were already as good as they could be! Realistically, 1 mW is about the maximum power that a station is likely to receive, and the station would

have to be within inches of a powerful AP to receive even that much! (That might surprise you, but remember that 802.11 NICs can receive signals down to about 10^{-7} to 10^{-9} mW, so a "weak" 1 mW signal is actually millions of times stronger than the weakest signal the card could discern.) Consequently, the RSSI ranges used by the 802.11 chipset manufacturers don't measure at all above 1 mW. The maximum RSSI value for all 802.11 chipsets and NICs represents a power level that is less than 1 mW. Everything between 1 mW and 100 mW is simply considered "100 %" RSSI .

This provides another potential reason to avoid comparing signal strength numbers between two vendors' chipsets. Both chipsets might receive exactly the same signal power, but one might report "100%" signal strength, and the other might report "90%" signal strength, depending on where the two vendors have chosen to set the "100%" point on their RSSI scale.

Using A Percentage Signal Strength Metric

To circumvent the complexities (and potential inaccuracies) of using RSSI as a basis for reporting dBm signal strength, it is common to see signal strength represented as a percentage. The percentage represents the RSSI for a particular packet divided by the RSSI_Max value and multiplied by 100 to derive a percentage. Hence, a 50% signal strength with a Symbol card would convert to an RSSI of 16 (because their RSSI_Max = 31). Atheros, with RSSI_Max = 60 would have RSSI = 30 at 50% signal strength. Cisco ends up making life easy with an RSSI_Max = 100 so 50% is RSSI = 50. [2]

If all vendors used the formula above for converting RSSI to signal strength percentage, then percentage for signal strength provides a reasonable cross-vendor metric for use in network analysis and site survey work. If signal strength is 100%, that's great! When signal strength falls to a particular vendor's threshold, the station will try to roam. Ultimately, when signal strength is down some very low value, the channel is going to be assumed to be clear. This conceptualization obviates the need to consider dBm, the RSSI_Max, or the "knee" in the logarithmic curve of mW to dBm conversion. It allows a reasonable comparison between environments even if different vendor's NICs are used to make the measurements. Ultimately, the generalized nature of a percentage measurement allows the integer nature of the RSSI to be overlooked.

However, if vendors do not consistently use the formula above, then we once again end up in a scenario where it's impossible to compare numbers from different vendors. For example, a vendor might hypothetically use a logarithmic function to map RSSI to signal strength, which would cause the signal strength to stay at high values longer as RSSI decreased, and then to drop off very rapidly as

---

[2] As we mentioned earlier in the paper, these RSSI numbers are based on the vendors' 802.11b chipsets, circa 2002. These numbers are provided only as an example of an issue that is still relevant today, and may not be correct for the vendors' current chipsets.

RSSI approached zero. Frankly, we don't know how vendors map RSSI to signal strength percentage, so it's difficult to draw concrete conclusions on this matter.

The Effect Of Signal Fade On Real-World Signal Strength Measurements

If an engineer stands in a particular location and uses a wireless network analysis tool (like AirMagnet, Sniffer Wireless, or AiroPeek) to measure signal strength from an access point the tool will not report a single, unchanging value even though it is standing still. Rather, the signal strength metric (whether it's represented as mW, dBm, RSSI, or percentage) will fluctuate across a range. The fluctuation is due to the fact that the signal being measured is constantly being affected by environmental influences in the dynamic electromagnetic spectrum in which the RF energy exists. The power contributed to the environment by an 802.11 device is very small. Consequently, other sources of RF energy (seen as noise by the 802.11 receiver) can disrupt the 802.11 signal.

The term *fade margin* refers to the magnitude of the maximum reduction in signal strength resulting from the various environmental influences. When a wireless network is designed, a fade margin value (often 10 dB) is included to account for possible environmental degradation. For example, if a particular 802.11b radio can operate at 11 Mbits/sec with a received signal level of -85 dBm then a network design would be based on signal propagation to the -75 dBm level.

Because of signal fading you will find that the signal strength reported by an analysis tool will not remain fixed at any particular value. It will "hover," going up and down throughout the fade range.

Final Conclusions About RSSI

The primary conclusion about RSSI, as detailed above, is that, due to the fact that 802.11 doesn't specify a required method of measuring RSSI, signal strength numbers from different vendors should not be compared to each other, since they are probably measuring it in different ways. In addition, given that 802.11 does not require any particular accuracy, and that the integer nature of RSSI limits the precision of signal strength measurements, in general, we recommend that signal strength measurements given by an 802.11 card be taken with a grain of salt. The RSSI value is intended for the card to use in its internal logic, an application that does not require either high precision or accuracy. Its use as a means of reporting signal strength to the administrator should be considered only a convenient extra. Signal-strength-related guidelines from a vendor, about that vendor's chipset, may still be useful, however. For example, if Vendor A says that its chipset will roam when signal strength drops below 40%, that information should be considered to be reliable, since the same mechanisms that are used to report signal strength to the client utility are used by the card to determine whether it should attempt to roam.

Section 3: Signal Quality

Alongside signal strength, "signal quality" is commonly reported by vendors' client utilities. For example, Cisco's client utility, seen in figure 3.1, presents a graph with signal strength on one axis and signal quality on the other axis. Anyone who is exposed to 802.11 for any length of time will pick up on the idea that these are the two general metrics of the "goodness" of an 802.11 signal. If either one of them is too low, then data rates may drop or the connection may be lost entirely. But what is "signal quality" really measuring? How can we use this term in planning and administering wireless LANs?
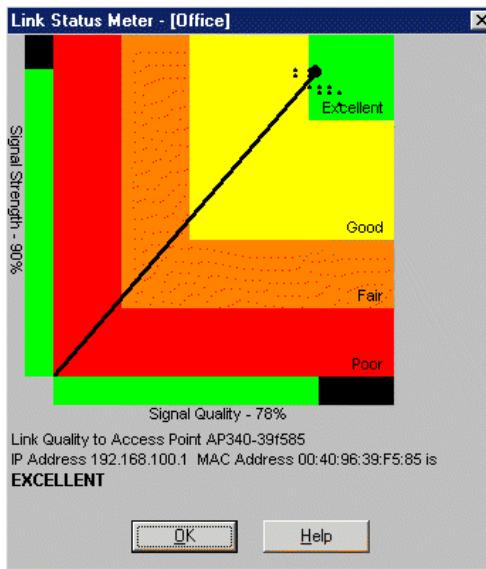
Figure 3.1 – Signal Quality in the Cisco Aironet Client Utility.

Signal Quality In The 802.11 Standard

Colloquial definitions of signal quality vary widely. A web search on the term reveals that it is usually used as a general indicator of the "goodness" of a wireless link, without any more specific definition. In some cases, "signal quality" appears to be being used as a substitute for "signal strength." This confusion is in no small part because "signal quality" is a more obscure term than "signal strength" (it only appears three times in 802.11-1999). 802.11 defines "signal quality" as "PN code correlation strength." Additionally, 802.11 defines the *PMD_SQ.indicate* primitive, which can be used by the PHY to report "the [signal quality] of the DSSS PHY PN code correlation." On the following page you'll find the text from the 802.11 standard.

The 802.11 standard describes the PMD_SQ_indicate primitive as follows:

15.4.5.11 PMD_SQ.indicate

15.4.5.11.1 Function

This optional primitive, which is generated by the PMD sublayer, provides to the PLCP and MAC entity the SQ of the DSSS PHY PN code correlation. The SQ shall be sampled when the DSSS PHY achieves code lock and shall be held until the next code lock acquisition.

15.4.5.11.2 Semantics of the service primitive
The primitive shall provide the following parameter: PMD_SQ.indicate(SQ). The SQ shall be a measure of the PN code correlation quality received by the DSSS PHY. SQ indications of up to 8 bits (256 levels) are supported.

15.4.5.11.3 When generated
This primitive shall be generated by the PMD when the DSSS PHY is in the receive state and code lock is achieved. It shall be continuously available to the PLCP, which, in turn, provides the parameter to the MAC entity.

As the definition implies, 802.11 only defines this primitive for the 2.4 GHz DSSS PHY (section 15) and not the FHSS PHY. 802.11a does not define this primitive. 802.11g mentions "signal quality," but points out (19.9.5.11) that it does not apply to ERP-OFDM modulation. (802.11g defines several different types of modulation types, of which ERP-OFDM is the only one that uses OFDM for the entire packet, including the PLCP header.) Therefore, we conclude that "signal quality" does not apply to non-DSSS modulation types—specifically, FHSS and OFDM. This makes sense, since DSSS is the only type of modulation that uses a PN code.

Unfortunately, other than the above sections, 802.11 does not provide any further definition of "signal quality", therefore, further investigation must hinge on the phrase, "PN code correlation". DSSS modulation is a complex topic, and although this phrase may be well-understood in DSSS circles, few WLAN engineers run in those circles! Further obfuscating the issue is that fact that mathematics is the language of DSSS, and most attempts to explain DSSS attempt to leave out the math, oversimplify, and end up making mistakes or being inaccurate in their definitions. It's no wonder that the term, "signal quality," is so misused.

In order to avoid introducing any mistakes, I will describe DSSS at as high a level as possible.

PN Codes In Direct Sequence Spread Spectrum

DSSS mathematically combines one or more bits of data with a sequence of bits known as a "PN code," which is equal in length to or longer than the data bits to be encoded. The combination of the data bit(s) and the PN code, known as a "symbol," is what is ultimately transmitted into the air. The process of combining is known as "spreading." For example, in 802.11/802.11b DSSS at 1 Mbps and 2

Mbps, a single bit of data (a 1 or a 0) is XOR'ed with the 11-bit-long PN code, '10110111000'. This specific PN code is known as the "Barker Sequence" and is also sometimes represented as '+1, -1, +1, +1, -1, +1, +1, +1, -1, -1, -1'. XOR'ing a single bit with the 11-bit Barker Sequence results in an 11-bit symbol. In this algorithm, each 11-bit symbol corresponds to one data bit. At higher data rates, more sophisticated algorithms are used, which may change the length of a symbol or the number of bits represented by a single symbol, but the nature of the PN code is the same.

PN Code Correlation Strength

Now that we know what a PN code is, what does "PN code correlation strength" mean? "Wireless Communications and Networks" (William Stallings, Prentice Hall, 2002) defines correlation thusly:

> *In essence, correlation is the concept of determining how much similarity one set of data has with another. Correlation is defined with a range between -1 and 1, with the following meanings: 1 = the second sequence matches the first sequence exactly; 0 = there is no relation at all between the two sequences; -1 = the two sequences are mirror images of each other.*

In addition, DSSS texts refer to the "correlator" as that part of the DSSS receiver that reverses the process of combining the original data stream with the PN code (known as "de-spreading"). In other words, a DSSS correlator receives symbols from the air and combines those symbols with the correct PN code to recover the original data. We were not able to find a specific definition of "PN code correlation strength" in our research, but given the above two definitions, we believe that we can speculate one. Although we aren't able to eliminate speculation about the term, "signal quality," hopefully, our speculation is more accurate and based on more facts than previous attempts to define the term. Regardless, this definition should be taken with a grain of salt until it is corroborated or debunked.

Consider an example of spreading the bit sequence '10' using the method described above: Barker coding, as applied in 802.11 DSSS at 1 Mbps and 2 Mbps. As seen in figure 3.2 (on the following page), the "result" would be two symbols, '10110111000' and '01001000111'. In fact, since Barker coding always uses the same PN code and encodes one bit per symbol, these are the only two symbols that will ever be modulated onto the air! A '1' will always spread to '10110111000', a '0' will always spread to '01001000111', and those are the only two possible values that can be spread! This logic can be generalized as follows: for a given spreading algorithm, define the number of bits encoded by a single symbol as n. A bit-string of length n has 2^n possible unique values. Therefore, the minimum number of unique symbols is also 2^n.
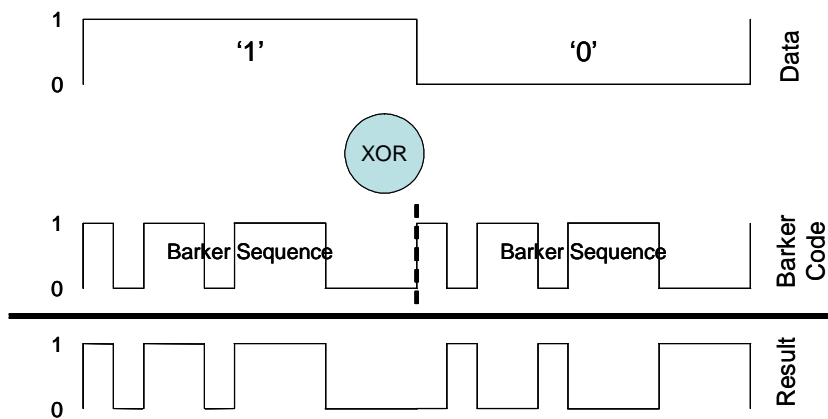
Figure 3.2 –DSSS "spreading" using the Barker sequence PN-code.

So, for any DSSS algorithm, a set of "valid" symbols, containing at least 2^n elements, can be defined. A transmitter should never transmit an "invalid" symbol, but interference and attenuation might cause a "valid" symbol to appear to be "invalid" to a receiver. For example, suppose the symbol '1011011100<u>1</u>' was received. This symbol is identical to the correct symbol for a '1', except that the last bit has been changed from '0' to '1'. The receiver could ask itself, "is this symbol closer to a zero or a one?"

Symbol for a '1':	1011011100<u>0</u>	
Received Symbol:	1011011100<u>1</u>	Different bit positions: 1
Symbol for a '0':	<u>01001000111</u>	
Received Symbol:	<u>1011011100</u>1	Different bit positions: 10

The receiver could determine that the symbol is closer to a '1' than a '0' even though one of the bits of the symbol was corrupted. In fact, given an 11-bit symbol, up to five bits could be corrupted without negating the receiver's ability to recover the original data. If six bits or more are corrupted, then the receiver could choose the incorrect symbol, which would result in an incorrect data-link layer checksum and a corrupt packet.

Final Conclusions About Signal Quality

The most likely definition of "signal quality," or "PN code correlation strength" is that it is some metric of the correlation between the correct symbol-stream and the actual symbol-stream received. For example, the PHY might count the average number of "wrong" bit positions over a window of some number of symbols, where zero "wrong" bit positions equals 100% signal quality and more "wrong" bit positions results in lower signal quality. This definition should be treated as speculation until it is corroborated or debunked.

If that definition is correct, then signal quality should be used as a metric of the amount of corruption in the environment between the access point and the client.

This use is consistent with manufacturers' statements about signal quality. For example, a manufacturer might state that its 802.11b chipset requires a signal quality of 20 dB or more to achieve an 11 Mbps data rate, above and beyond the minimum power level required to achieve the same data rate. In other words, an 11 Mbps signal received at -50 dBm should be more than strong enough for all 802.11 cards on the market, but if the quality of the signal is below 20 dB (in this example), the packet may be corrupted, regardless of its high signal strength.

It's common to equate signal quality with "signal to noise ratio," which is sometimes defined as, "the power of the 802.11 packet relative to the background noise in the environment." This definition is incorrect because it's not accurate to associate "PN code correlation strength" solely with the power of the data packet relative to background noise. Certainly, increased levels of ambient RF radiation could cause corrupted bits and decreased signal quality, but a correct definition of signal quality must take into account other possible causes of corrupted bits.

The final conclusions regarding signal quality are similar to those regarding signal strength: 802.11 does not define a specific method of calculating and reporting signal quality, meaning that vendors will most likely implement it in inconsistent ways. Signal quality measurements will most likely be inconsistent between different vendors. As with signal strength, signal-quality-related guidelines from a vendor, about that vendor's chipset, may still be useful.

Section 4: Signal To Noise Ratio

As the last section of this essay mentioned, "signal to noise ratio," or SNR, is commonly defined by 802.11 administrators as "the ratio of the power of the data signal to the power of the ambient RF energy," where "ambient RF energy" refers essentially to any RF emitter that is not an 802.11 transmitter. Although this definition could be correct under some interpretations of the term, "signal to noise ratio," it's important for an 802.11 administrator to realize that slightly different interpretations of that term are used in other areas, and confusion may arise if it's not clear exactly what kind of "signal to noise" is being discussed.

Figure 4.1 – SNR as reported by Cisco's Aironet Client Utility.

Defining "Noise"

In general, "signal to noise ratio" refers to the power level of an incoming signal relative to some type of background noise. The definition of "noise" can vary, depending on the field in which "noise" is being measured. For example, audio engineering refers to the SNR of an analog recording medium—that is, the strength of the loudest undistorted signal the medium can carry relative to the background "hiss" that would be heard if a blank medium were played back. Audio amplifiers also have an SNR rating, which essentially measures the level of "hiss" that would be heard if the amplifier were turned up all the way without a signal being fed through it (you can try this at home with your stereo if you want). Although we have used audio examples, since everyone has some experience with them, the same "hiss" exists in any analog circuit, including the RF receiver in your 802.11 card. The point here is that, because the definition of "noise" can vary, so does the definition of SNR. In each case, the "signal" is the power level of the incoming signal, but the "noise" changes depending on your perspective.

From the perspective of an RF chipset engineer or an electrical engineer, "signal to noise" ratio probably refers to the strength of the incoming signal relative to the "noise" within the RF chipset itself. Electronic circuits are subject to a type of noise known as Boltzmann noise, which is caused by thermal effects. Essentially, the heat within the chipset causes a certain amount of electrical distortion that manifests itself as noise. Electrical circuits are also subject to induced noise from outside sources—for example, it's common to find noise in the 60 Hz range (the frequency of AC current in the U.S.). Notice that this definition of "noise" and

"signal to noise ratio" is completely different from the 802.11 administrator's definition mentioned previously. This definition focuses on the strength of the signal and noise within the electrical circuit itself, as opposed to in the air.

Measurement of SNR, then, depends on the definition of the "noise floor," which is the level of "noise" in the environment being measured.
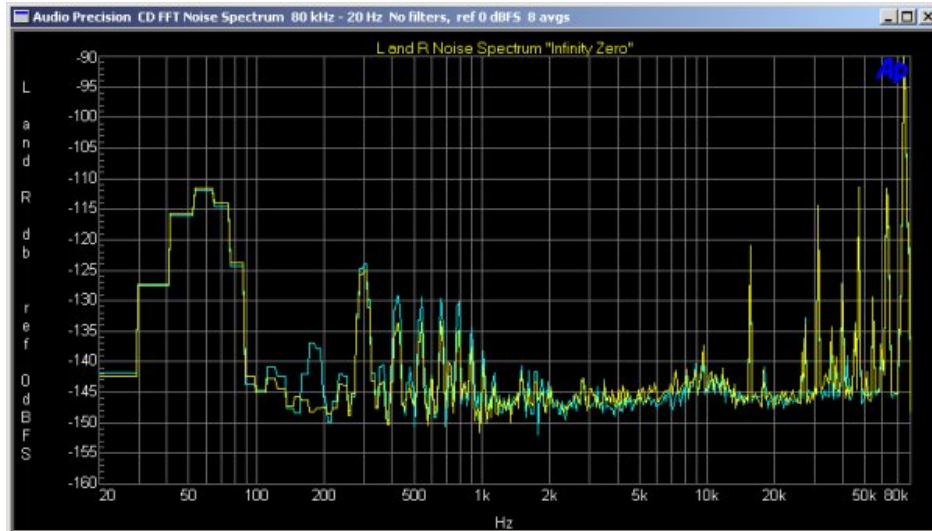

Figure 4.2 – "Internal" SNR measurement for a DVD player's audio outputs.

For example, the screen-shot above shows the "noise floor" of the audio outputs of a certain DVD player. This was created by measuring the output while nothing was being played through the DVD player, so "noise" in this case must refer to noise within the chipset. While some of this noise must be thermally-induced, notice the peak around 60 Hz—the frequency of AC electrical power. This peak is probably induced by factors outside the chipset. "Noise" within the electrical circuit is closer to the electrical engineer's definition—what we might call, "internal noise"—than the 802.11 engineer's definition—what we might call, "external noise."

In summary, to an RF chipset engineer, the "noise" might be the background noise within the circuit of the chipset, while, to an 802.11 administrator, "noise" might be the ambient, non-802.11, RF energy in the environment. To avoid a double-use of the term "noise" and "signal to noise ratio," we propose the term, "interference" to refer to the ambient, RF energy in the environment, and the term, "signal to interference ratio," or "SIR," to refer to the strength of the signal relative to the ambient, RF energy in the environment. This leaves "noise" and "SNR" to refer exclusively to the thermal noise within the chipset and avoids confusion when 802.11 administrators talk to RF engineers or electrical engineers. Finally, bear in mind that SIR is probably not what the card is measuring when it reports "signal quality," even though it's common to (incorrectly) equate signal quality with SIR.

Final Conclusions About Signal To Noise Ratio

Since 802.11 cards do not typically report SNR, it is hard to put the concept of SNR to use practically. As a WLAN administrator, you should understand that the received signal must be above a certain minimum threshold relative to the ambient, non-802.11 RF energy in the environment or it will be swallowed up in the noise. The effect is that environments with lots of background noise will require more powerful access points. Unfortunately, only a spectrum analyzer— typically an expensive piece of hardware—can accurately report the amount of background noise in an environment.

Section 5: Receive Sensitivity

A fundamental specification of an 802.11 card is its receive sensitivity. The receive sensitivity is the minimum power level at which a signal can be reliably received. For example, a NIC manufacturer may indicate that their particular card has a receive sensitivity of –96 dBm at 1Mb/sec. If the actual RF energy present at that card were less than –96 dBm, then the card would no longer be able to differentiate between signal and noise (where, in this case, "noise" refers to thermal noise internal to the chipset, not noise in the environment). The NIC would not detect the incoming packet at all, and the packet would be lost.

How Receive Sensitivity Is Measured

We asked a major vendor of 802.11 hardware how they measured receive sensitivity in their cards. They told us that to measure receive sensitivity, the WLAN card is placed into an RF-shielded room. This guarantees that the test signal will be the only RF transmission in the room, and no background noise in the environment will interfere with the test. The vendor then transmits packets at weaker and weaker power levels. As the power level decreases, the bit error rate as measured by the card increases. The receive sensitivity of the card will be the minimum power level at which the bit error rate remained below a certain threshold. Therefore, a lower receive sensitivity value (-93 dBm) is <u>better</u> than a higher one (-85 dBm), since it means that the card was able to "reliably receive" data at lower power levels.

Some WLAN designers attempt to ensure that the signal strength seen by a worst-case client is at least so many dB above the client's receive sensitivity. The rationale is that the card requires a certain minimum SNR to achieve a certain data rate, therefore, if the signal is even <u>close</u> to the card's receive sensitivity for that data rate, the card won't be able to achieve the data rate. This is incorrect, since the testing methodology used by the vendors already takes into account the SNR required by the card. As the signal in the shielded room gets weaker and weaker, it approaches the card's internal noise floor. As it gets close to the noise floor, the bit error rate goes up. When the signal is too close to the noise floor, the bit error rate exceeds the vendor's threshold, and that signal strength is the receive sensitivity of the card. As an example, imagine a card with an internal noise floor of -100 dBm. If the card requires an SNR of 10 dB to maintain an 11 Mbps data rate, then the receive sensitivity at 11 Mbps will be -90 dBm. All you need to know is that the signal strength must be above -90 dBm at the receiver. You don't need to care at all about the card's noise floor or its minimum SNR.

However, since the testing of receive sensitivity is done in a shielded room, there isn't any interference at all in the shielded room where receive sensitivity is measured. Therefore, signal to <u>interference</u> ratio is not taken into account by receive sensitivity. Clearly, if the interference (ambient RF energy in the area between the receiver and the transmitter) is too strong relative to the data signal,

the data signal will be "swallowed up" and will be lost. But we don't know of any reliable metrics for how strong is "too strong"! We have certainly seen some guidelines—e.g. the signal must be 20 dB stronger than the background noise to maintain 11 Mbps—but those guidelines are dubious. We don't know where they came from, whether they are based on signal quality instead of SIR, and so on. This is a fascinating area for future exploration.

The dBm value for a NIC's receive sensitivity is very close to the dBm value associated with an RSSI value of 0. Hence, the receive sensitivity of the adapter determines the lower end of the necessary measurement range for signal strength. It should be noted that, typically, if RSSI = 0 the dBm signal measurement is below the Receive Sensitivity level. It's impossible for any tool using a standard wireless NIC to measure signal strength below the NIC's Receive Sensitivity threshold.

Data-Rate Dependent Receive Sensitivity

Of course, different data rates, having more and less complex encoding and modulation methods, and being more and less resistant to corruption, will result in different receive sensitivities. As data rate increases, receive sensitivity decreases. To put it another way, the higher the data rate, the stronger the signal strength must be for the packet to be reliably received. This is why 802.11 cards drop to lower data rates when interference is present or when they are at the edges of their coverage range. For example, an 802.11b card might have specifications like this:

Data Rate	Receive Sensitivity
1 Mbps	-95 dBm
2 Mbps	-91 dBm
5.5 Mbps	-89 dBm
11 Mbps	-85 dBm

A Receive Sensitivity Caveat

While receive sensitivity might seem like a reliable way of comparing two vendors' cards, we know of no organization that certifies the veracity of the vendor's results. Therefore, there is the potential for vendors to manipulate the thresholds of their tests to influence their chipset's receive sensitivity numbers. For example, a vendor that uses a BER threshold of 10^{-9} (one error in every 1,000,000,000 bits) will end up with lower receive sensitivities than a vendor that uses a BER threshold of 10^{-8} (one error in every 100,000,000) even though the second vendor's card may actually be better at receiving bits.

Section 6: Fade Margin

We have heard of cases where designers of 802.11 networks referred to the card's receive sensitivity as the "noise" in the term "signal to noise ratio." In other words, the designer might say, "We have to ensure a minimum signal strength of -65 dBm, so that we keep at least 20 dB SNR." What the designer means is, "…so that we keep at least 20 dB over the card's receive threshold." Although one could argue that this was a valid use of the term SNR, it creates confusion, since SNR is already used in at least two other ways in 802.11: to measure the signal strength versus the amount of thermal noise in the chipsets and to measure the signal strength versus the amount of background RF energy in the environment (interference). "Fade margin" is a more appropriate term for this concept.

Fade margin is a standard telecommunications term that refers to the amount of "extra" energy in a communications channel. To put it another way, fade margin is the amount by which the signal can fade below its normal power level without disrupting the communications link. For example, if it's determined that a certain point-to-point RF bridge link requires a minimum output power of 18 dBm, and the link is actually putting out 24 dBm, the link has a fade margin of 24 dBm – 18 dBm = 6 dBm.

Causes Of Fading

In RF, fading is caused by changes in the environment between the transmitter and the receiver. For example, a person might walk between you and the AP, you might accidentally put your hand over the antenna of your PCMCIA card, you might turn in your chair so that your laptop is between the PCMCIA card and the access point, and so on. The environment between an 802.11 transmitter (an access point) and an 802.11 client is infinitely complex and constantly changing. RF interactions like reflection, refraction, and scattering mean that in indoor environments, there is never a single, simple line of sight between a transmitter and a receiver. Therefore, the signal strength that a client perceives can fade unpredictably, even when the client is sitting still.

The Importance Of Fade Margin

Fade margin is important because the signal strength will vary due to changes in the environment. But it's also important because the real implementation of a WLAN will always differ somewhat from even the best predictions. Figure 5.1 shows predictions of an access point's coverage, similar to predictions that Connect802 might use to plan the installation of a WLAN. In this example, the desired coverage is -75 dBm. Based on certain factors, Connect802's engineer determines that a 10 dB fade margin is appropriate. Therefore, the entire WLAN is planned based on a coverage of -65 dBm throughout. This ensures that, even if the real implementation differs from the predictions, it's likely that coverage will still be as good or better than the desired minimum signal strength.

In figure 6.1 below, the inner contour represents a -65 dBm signal level, while the outer contour represents the minimum desired signal strength, -75 dBm. The difference between these two contours is the fade margin..
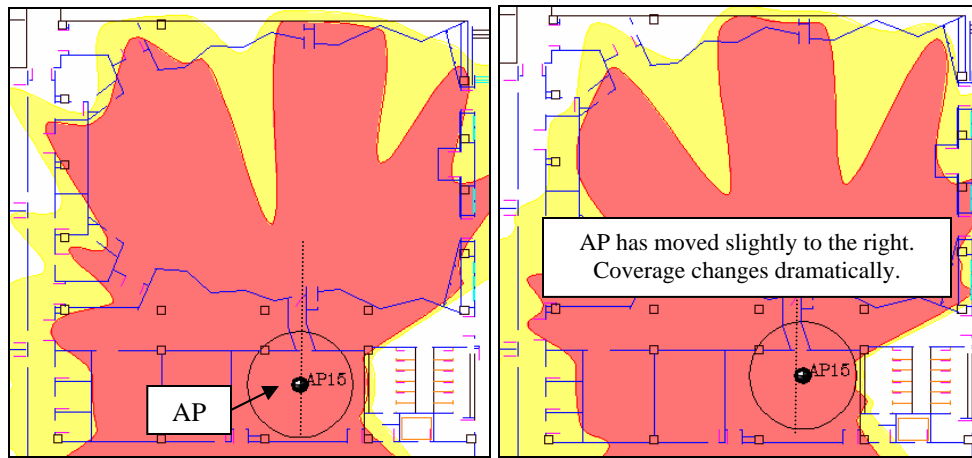

Figure 6.1 – How fade margin protects against discrepancy between predicted and actual WLAN coverage.

Notice, first that a very small change in the AP's position causes drastically different coverage within the large "auditorium" at the top of the screen shot. Without looking closely, you might not even notice that the AP is in a different location in the two screen shots (note the change in position relative to the vertical dotted line). If the installer doesn't place the AP underline{exactly} where the plan says he should, there could be holes in the WLAN's coverage. We essentially assume that the -65 dB contour is the "best-case" scenario, but that reality probably won't match that scenario. Next, notice that, regardless of the AP's placement, the 10 dB fade margin ensures that the entire room is covered.

Final Conclusions About Fade Margin

Like the designer mentioned above, most designers of 802.11 WLANs (including Connect802) engineer a certain amount of extra power into their systems. That is simply good practice. Instead of referring to this extra power as SNR or some other term, it is most correct to use a standard term that already exists for this concept. "Fade margin" is that term.

Appendix: Practical Conversion From Percentage To dBm

For reasons described in this paper, it is questionable as to the effectiveness or reasonability of using dBm measurements obtained from a standard wireless NIC as part of a real-world network troubleshooting exercise. This is because most NICs only provide RSSI in a range that is below –10dBm and everything above that is mapped to RSSI_Max (or, 100% signal strength). In addition, RSSI has unspecified accuracy and, generally, low precision. Nonetheless, it may be informative to study the methods that real manufacturers have used to convert between RSSI and dBm.

Following are conversion tables, based on information obtained from various NIC manufacturers, which will provide a mapping between RSSI and dBm. There is a two-step process to go from a percentage signal strength report in an analyzer to the dBm value in a vendor's table. First it's necessary to know the RSSI_Max for the vendor and, from that, the RSSI that corresponds to the current percentage value can be obtained. (i.e.: x% of RSSI_Max = RSSI). Once the RSSI value has been obtained from the percentage it's only necessary to plug it in to the vendor's table (or formula) and get a dBm value. You should notice, in each description that follows, how the values in the tables don't always increase in a linear manner. Sometimes a table value will go up by 5, other times by 6, and so forth. This is to account for the logarithmic nature of dBm measurements. Embodied in these "gaps" in the table, and exacerbated by the integer nature of the RSSI, are inherent potential inaccuracies that must be recognized.

NOTE: The formulae and tables described below may vary from one revision of a vendor's equipment to the next, but the derivation of dBm or signal strength percentage remains algorithmic in every case because it's always based on the RSSI value reported by the chipset. These numbers are based on the vendors' 802.11b chipsets circa 2002. We have not obtained information on more modern chipsets.

Conversion For Atheros

Unlike the other vendors described, Atheros uses a formula to derive dBm.
RSSI_Max = 60
Convert % to RSSI
Subtract 95 from RSSI to derive dBm
Notice that this gives a dBm range of –35dBm at 100% and –95dBm at 0%

Conversion For Symbol

RSSI_Max = 31
Convert % to RSSI and lookup the result in the following table:
RSSI <= 4 is considered to be –100dBm
RSSI <=8 is considered to be –90 dBm
RSSI <=14 is considered to be –80 dBm
RSSI <=20 is considered to be –70 dBm
RSSI <=26 is considered to be –60 dBm
RSSI greater than 26 is considered to be –50dBm
Notice that this gives a dBm range of –50dBm to –100dBm but only in 10dBm steps

Conversion For Cisco

Cisco has the most granular dBm lookup table.
RSSI_Max = 100
Convert % to RSSI and lookup the result in the following table. The RSSI is on the left, and the corresponding dBm value (a negative number) is on the right.

0	= -113	68	= -41
1	= -112	69	= -40
2	= -111	70	= -39
3	= -110	71	= -38
4	= -109	72	= -37
5	= -108	73	= -35
6	= -107	74	= -34
7	= -106	75	= -33
8	= -105	76	= -32
9	= -104	77	= -30
10	= -103	78	= -29
11	= -102	79	= -28
12	= -101	80	= -27
13	= -99	81	= -25
14	= -98	82	= -24
15	= -97	83	= -23
16	= -96	84	= -22
17	= -95	85	= -20
18	= -94	86	= -19
19	= -93	87	= -18
20	= -92	88	= -17
21	= -91	89	= -16
22	= -90	90	= -15
23	= -89	91	= -14
24	= -88	92	= -13
25	= -87	93	= -12
26	= -86	94	= -10
27	= -85	95	= -10
28	= -84	96	= -10
29	= -83	97	= -10
30	= -82	98	= -10
31	= -81	99	= -10
32	= -80	100	= -10
33	= -79		
34	= -78		
35	= -77		
36	= -75		
37	= -74		
38	= -73		
39	= -72		
40	= -70		
41	= -69		
42	= -68		
43	= -67		
44	= -65		
45	= -64		
46	= -63		
47	= -62		
48	= -60		
49	= -59		
50	= -58		
51	= -56		
52	= -55		
53	= -53		
54	= -52		
55	= -50		
56	= -50		
57	= -49		
58	= -48		
59	= -48		
60	= -47		
61	= -46		
62	= -45		
63	= -44		
64	= -44		
65	= -43		
66	= -42		
67	= -42		

Notice that this gives a range of –10dBm to –113dBm. Bearing in mind that a Cisco card will have a Receive Sensitivity of –96dBm at its lowest and it's impossible to obtain an RSSI value of less than 16. Note, also, that all RSSI values greater than 93 are assigned –10dBm and that there are multiple places in the table where two adjacent RSSI values are assigned the same dBm value.

There is another aspect to interpreting the Cisco RSSI. The Cisco device driver code indicates that if the RSSI value converts to less than –90dBm then it should be converted to a fixed value of –75dBm. This leaves a question as to the exact interpretation of an RSSI value that converts to –76dBm to –89dBm. A measurement of –75dBm would be reported as 36% signal strength. It seems sufficient to leave this quandary unanswered since whether a signal is at –75dBm or at –92dBm, the entire low-end range is less than what would be desirable for normal 802.11 WLAN operation.