RUDI VAN DRUNEN

peculiarities of radio devices

Rudi van Drunen is a senior UNIX systems consultant with Competa IT B.V. in The Netherlands. He also has his own consulting company, Xlexit Technology, doing low-level hardware-oriented jobs.

rudi-usenix@xlexit.com

IN THIS ARTICLE I WILL DESCRIBE THE peculiarities of using Radio Frequency (RF) devices in your systems. Most prominently, the WiFi subsystem in your laptop or access point will be covered, but I will also be touching upon Bluetooth and 3G (Cellular) devices.

Radio Systems

All wireless systems depend on radio waves to transmit and receive data. Radio wave frequency and the way data is encoded in radio waves differ from technology to technology.

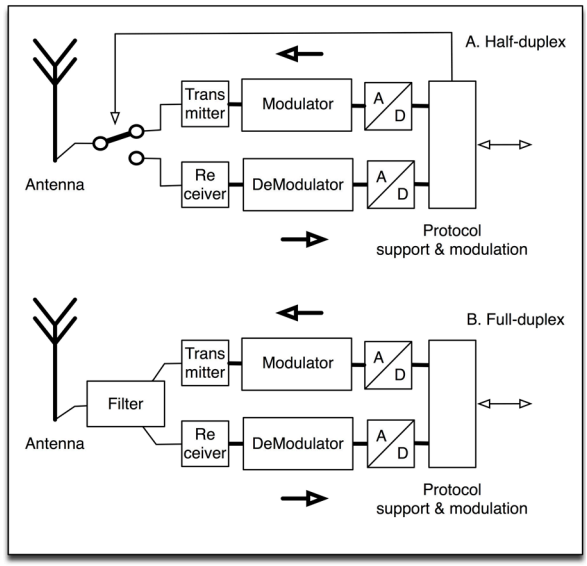

FIGURE 1: RADIO SYSTEMS: A: HALF DUPLEX; B: FULL DUPLEX

As illustrated in Figure 1, some radio systems switch between receiving and transmitting (half-duplex systems), and some radio systems are able to receive and transmit at the same time (full-duplex systems) using the same antenna. Full-duplex systems are often more complex, since designers need to prevent the high-power transmitter signal from getting into the sensitive receiver front end. The high-power signals, even if they are on another frequency, can make the receiver "deaf" or can even destroy the sensitive input amplifier electronics (we're talking about nano Volts sensitivity). Full-duplex systems usually get more throughput than half-duplex systems.

In the data encoding/decoding process we move from digital to analog technology; radio waves are all in the analog domain. Part of the encoding and

decoding is traditionally done in hardware using discrete electronics and integrated circuits by functional blocks known as modulators and demodulators, but as digital signal processors get more common and powerful, one can do this completely in software as well, resulting in software defined radio as shown in Figure 2. (More info on software defined radio can be found in [1].)

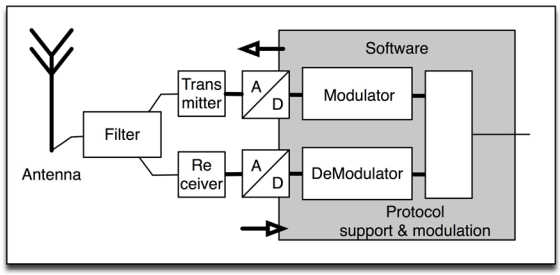

FIGURE 2: SOFTWARE DEFINED RADIO (SDR)

Almost all modern chipsets that provide WiFi, Bluetooth, or 3G capabilities nowadays are more or less software defined radios. By having the software doing the modulation/demodulation, it is much easier to adhere to newer protocols with the same hardware through simple firmware updates. Also calibration, which can be a tedious process in analog systems, can be skipped or replaced by simple software routines that calculate the different parameters to be used in the modulation and demodulation algorithms.

Wavelength

The radio waves, traveling at the speed of light (c = 300,000 Km/s through vacuum or air) have a specific wavelength (λ), which relates to the frequency the system operates at as $f = c / \lambda$. A typical WiFi system operating at 2.4 GHz has a wavelength of 12.5 cm (4.9 in). (See the first article in this series for more about wavelengths [2].)

Antennas

The last (or the first) part of a radio system is always the antenna. This device translates the electrical energy into electromagnetic (radio) waves, or the other way around. We can look at an antenna as a piece of wire ("the element") that can be brought into resonance, and often as a reflector to reflect and direct the radio waves. When the wire comes (partly) into resonance, at a particular frequency, it starts to efficiently transmit or receive radio waves. The frequency at which the wire element gets efficient (the resonance frequency) is dependent on its size (there is a relation to the wavelength), shape, and material. All different frequencies need different antennas.

The frequency a system operates on differs from standard to standard. Table 1 shows the commonly used frequencies (some technologies can and are designed for operation in multiple frequency bands) in everyday data communication use.

Technology	1st frequency band	2nd frequency band
WiFi	2.4 GHz	5 GHz
Bluetooth	2.4 GHz	
WiMax	2.5 GHz	3.5 GHz
3G (UMTS)	1.8 GHz	1.9 GHz
ZigBee	900 MHz	2.4 GHz

TABLE 1: FREQUENCIES USED BY DIFFERENT TECHNOLOGIES

Next to the typical frequency an antenna is designed for there is the gain an antenna provides. The gain of an antenna works both in the transmitting mode and in the receiving mode and is highly dependent on the size and shape of the antenna.

It you have an omni-directional antenna, then the gain versus direction in the x-y (horizontal) plane is the same in all directions, whereas a directional antenna has a distinct direction in which the gain is at its largest. Figure 3 shows a simulation of two different antennas and their gain characteristics.

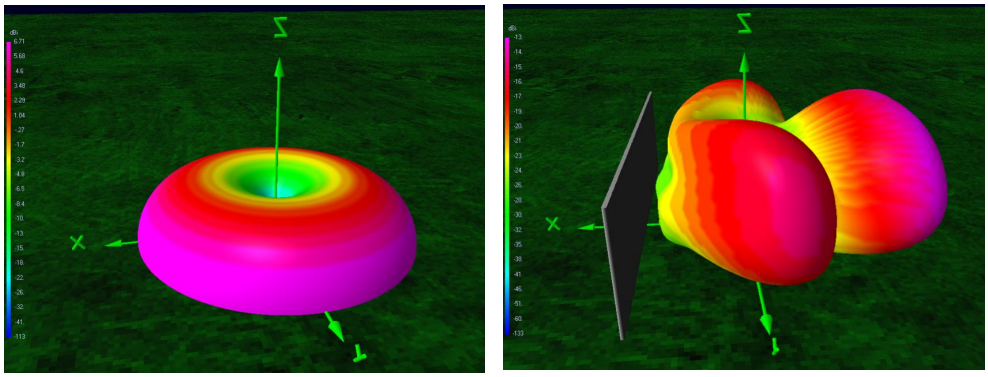

A: OMNI-DIRECTIONAL **B: DIRECTIONAL**

FIGURE 3: SIMULATED RADIATION PATTERNS

Omni-directional antennas are often used by access points where you want to cover as large an area as possible. You want the clients to be able to connect from all directions.

If for some reason you need a segment of the donut shape as coverage area, there are special sector antennas for this purpose. Mobile operators providing 3G service often use sector antennas with different opening angles (60, 90, or 120 degrees). By using different sectors they can shape the coverage area to their needs or use more transmitters in their base station, each serving a particular direction. The sector antenna is just a special case of the omni-directional antenna used to cover an area for clients.

Directional antennas, on the other hand, have a very small coverage angle: they often have opening angles of less than 10 degrees. Typical use of a directional antenna is in making point-to-point connections where the location of the other antenna is static and exactly known.

A special, mixed form of omni and directional antennas can be found in an active antenna that is able to direct a particular signal to a particular client within the 360 degrees of the omni antenna. This antenna requires pretty complex RF electronics to do phase shifts of the signals and uses multiple elements to transmit. This so-called phased-array antenna is often used by WiMax providers in their base-stations.

Figure 4 shows the actual real-life pictures of two antennas with their characteristics.

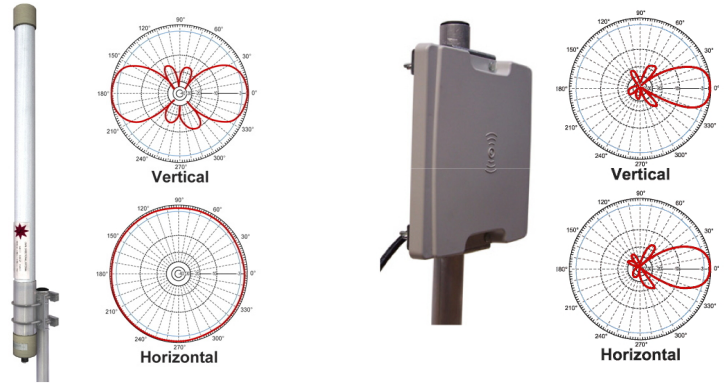

A: 8DBI OMNI-DIRECTIONAL B: 12 DBI DIRECTIONAL PANEL

FIGURE 4: PICTURES AND DIAGRAMS OF AN OMNI (A) ANTENNA AND A DIRECTIONAL (B) ANTENNA (SOURCE: WWW.GANDALF.NL)

The dB

In Figure 5 we introduce a new unit of measurement, the dB (decibel). The dB is not an absolute unit but denotes the relationship between two signals. When talking about power levels in W(atts), the ratio of two signals can be expressed in dB by calculating the 10-base Logarithm ($\log_{10}$) of the ratio between the actual level to a reference level, multiplied by 10 (since 10 dB (decibel) equals 1 B (Bell).

If you want to express a twofold increment in power, moving from 100 mW to 200 mW, for example, the dB increase equals

$$10\log_{10}\left(\frac{200mW}{100mW}\right) = 3dB$$

So if an antenna amplifies a signal 100 times in a particular direction, it has a 20 dB attenuation in that direction.

Often when talking about antennas we use the word "gain," which is synonymous with "attenuation." For antennas, we compare the gain of an antenna with a hypothetical isotropic radiator which radiates uniformly in all directions, and express the gain in dBi, whereas if we talk about (absolute) power levels we denote it by dBm. With dBm, we compare to 1 mW of power, so 100mW corresponds to

$$10\log_{10}\left(\frac{100mW}{1mW}\right) = 20dBm$$

Now let's put all this theory into action. We all know about WiFi systems and their peculiarities. To be able to make a point-to-point connection (building to building) using WiFi, we need to select equipment such as antennas, cabling, and access points/client devices in such a way that the connection is reliable but also within the limits defined by the FCC or its European Counterpart, ETSI, in maximum RF power output. Here we will show how to lay out such a system that will perform while operating within limits.

An Example

Having set up many point-to-point links in a wireless community network [3], I will show the theory in action on designing a 3 Km (a little less than 2 miles) link connecting two hops in the network. This example is shown in Figure 5. Here all components that add or reduce gain in the RF path are shown. The system we are designing uses 802.11b technology on an 11 Mb link rate. The connection is using the 2.4 GHz band.

802.11b technology has proven to be far better performing in outside environments than 802.11g due to the more robust RF modulation.

Making the calculations, we start off with calculating the loss that is imposed on the radio signal while traveling through free space:

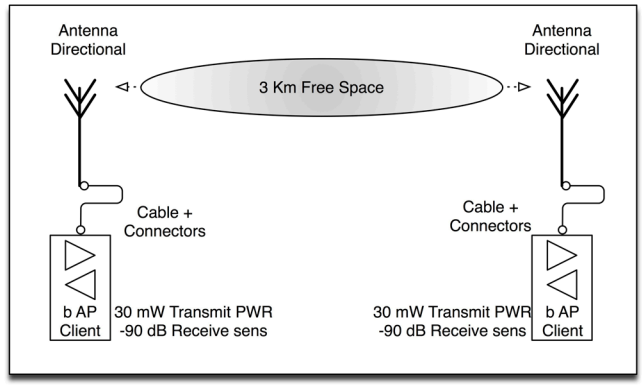

FIGURE 5. A BUILDING-TO-BUILDING OUTDOOR WIFI SETUP

The signal loss through air (free space loss) is defined as:

$$\text{Loss[dB]} = 92.5 + 20 \log f\,[\text{GHz}] + 20 \log d\,[\text{Km}]$$

with f the operating frequency in GHz and d the distance in kilometers.

In this case we have a loss of 92.5 + 7.6 + 9.5 = 109.6 dB.

We have access point and client devices that output 30 mW = 10 log10 (30 mW/1mW) = 14.8 dBm of RF energy and need a signal strength of –90 dBm for maintaining a good connection. All cabling and connectors in the system on each side have a loss of 1.2 dBi, recalling that we have cabling and connectors on both ends.

Now we can calculate the needed antenna gain. We need –90 dBm on the input of the device for a good connection. Allowing some margin, we define the minimum input signal as –85 dBm. This signal must come from the 14.8 dBm transmitter, going through a set of cabling twice, which gives a total of –2.4 dBi gain, and the free space, which gives –109.6 dB gain.

So the antennas must add another 14.8 – 109.6 – 2.4 – (-85) = 12 dBi, which can be accomplished using two antennas that have a gain in the desired direction of 6 dBi.

It is important to use the antennas in a symmetrical setup—that is, both sides having the same gain, as both devices send and receive; otherwise the power in front of the largest antenna will exceed maximum.

So a 6 dBi antenna for each side would do the job. To get some headroom, it is wise to select 7 dBi antennas, which gives you another 2 dBi extra margin. In a rainy climate, you might need some extra gain to cope with the extra loss you get from wet air. To check if we are within limits defined by the regulatory body, we have to calculate the output power of the antenna (the

field strength in front of the radiator), which in this case is 14.8 dBm (AP) + 7 dBi (antenna) – 2.5 dBi (cabling set) = 19.3 dBm = 85 mW of RF output power, which is allowed both in Europe and in the USA for the 802.11b band.

Another thing to take into account when designing a point-to-point WiFi link is that for the above free space loss calculation the RF beam needs to be unobstructed, so a line of sight is needed and the Fresnel zone must be clear. The Fresnel zone is an ellipsoid between the transmitter and receiver in which the radio waves live (see Figure 6).

The maximum radius of this Fresnel zone can be calculated by the following formula:

$$r[m] = \sqrt{\frac{d[Km]}{4 * f[GHz]}}$$

For the 3 Km link, the maximum radius of the ellipsoid would be about

$$17.32 * \sqrt{3/(4*2.4)} = 9.6 meters = 31\ ft.$$

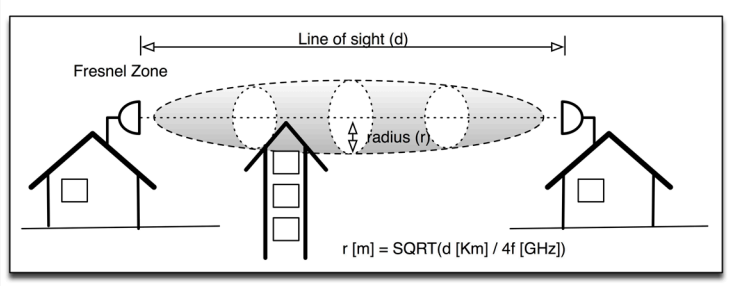

FIGURE 6: THE FRESNEL ZONE

Large objects covering (part of) the Fresnel zone will increase the free space loss, and therefore larger antenna gain might be needed. It is wise to have the Fresnel zone less than 10% occupied with trees or buildings. A larger percentage of occupation will seriously increase the free space loss.

If you make very long distance point-to-point connections, it is important not only to look at the Fresnel zone but also to take the curvature of the Earth into account. The Earth's curvature may result in partial coverage of the Fresnel zone when the antennas are not mounted high enough.

Range-Limiting Factors

There are a number of factors that limit the range of a wireless datacom system.

MULTIPATH

As with any waves, radio waves can be reflected by surfaces such as walls. These reflected waves, which are weaker than the direct wave, are received slightly after the primary signal (see Figure 7).

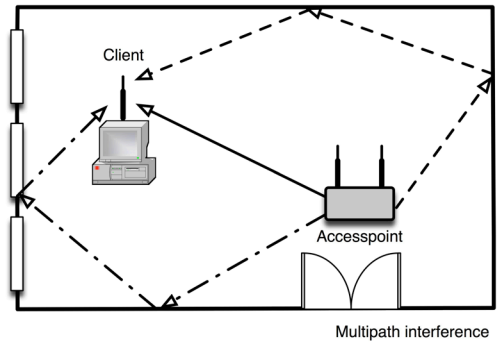

FIGURE 7: MULTIPATH INTERFERENCE

This causes interference and distortion of the resulting signal. This effect, commonly known as multipath interference, can distort the signal and make decoding/demodulating the RF signal particularly difficult for the hardware. Reflections can be of a different nature, as shown in Figure 8.

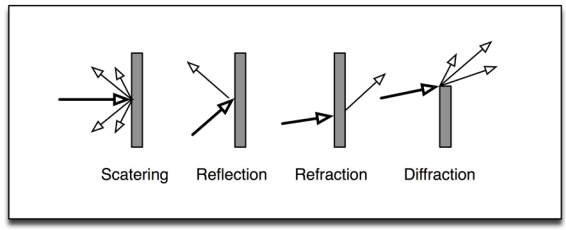

FIGURE 8: DIFFERENT WAYS REFLECTION WORKS

In some cases, a strong enough signal that is received out of phase with the direct signal can essentially create a blank, a spot where no signal is available, while only a few feet away you may have a strong signal. This is called a multipath null.

Using omni-directional antennas you will create large amounts of multipath interference, as the RF radiation is transmitted in all directions. Using directional antennas makes this less of a problem, as most RF energy is transmitted in one distinct direction.

ATTENUATION

An RF signal is reduced in strength once it passes through different materials. This effect is known as attenuation. The degree of reduction as related to material is shown in Table 2.

Material	Attentuation @2.4 GHz	Attentuation @5 GHz
Solid wood door	–6 dBi	–10 dBi
Steel fire exit/door	–13 dBi	–24 dBi
Concrete wall 18”	–19 dBi	–30 dBi
Interior wall 6”	–9 dBi	–4 dBi
Single pane window	–7 dBi	–6 dBi
Cubicle wall	–6 dBi	–2 dBi

TABLE 2: ATTENTUATION OF VARIOUS MATERIALS FOR WIFI SIGNALS

Table 2 shows that materials such as concrete walls or floors have much more attenuation than cubicle walls. In designing a radio network it is important to determine the number and material of the obstacles the signal must travel through. This has to be included in calculations of free space loss.

SIGNAL-TO-NOISE RATIO

The relation of the target signal of the radio system to other signals in the same frequency band is called the signal-to-noise ratio (SNR) and is expressed in dB as: SNR [dB] = Signal strength [dBm] – Noise floor [dBm] (see Figure 9). The more noise there is, the more difficult it becomes for the front-end electronics (or software) to filter the desired signal from the noise. Often the data rate the system is working on will drop if the SNR decreases. The more complex the modulation (encoding) of the RF signal is, the higher the bandwidth the system is capable of, and the higher the minimum SNR needs to be to get reliable communications.

Noise can be generated by different pieces of equipment. For 2.4 GHz, noise coming from stray microwave radiation, video transmitters, or ham radio operators is notoriously bad. Other sources of noise are X10 home automation equipment and cordless telephones.

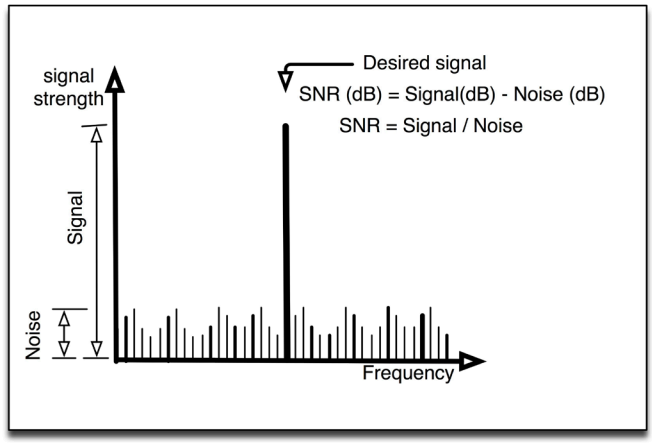

FIGURE 9: SIGNAL-TO-NOISE RATIO

An example here is the operation of Bluetooth and WiFi at the same time. The SNR of one system will decrease when the other system is turned on and the antennas are located close together. The Bluetooth RF signal is noise to the WiFi subsystem and vice versa. Other WiFi signals from adjacent channels are also noise to a client. Since the spectrum width of WiFi signals overlap, it is important for performance to use channels that are at least three apart.

GENERAL

These effects are applicable to all RF systems. WiFi (IEEE 802.11), Bluetooth (IEEE 802.15), ZigBee (IEEE 802.15.4), and 3G communication systems all use RF in roughly the same frequency space (2.5 GHz), and all have the same characteristics in the lowest layer. However, the higher layers of the different protocols differ quite extensively, and effects on the RF layer may be less profound for the user, since the other layers have the ability to correct errors introduced on the RF layer. An example here is that WiFi 802.11 a/b/g systems often use diverse techniques to overcome problems due to

multipath interference or multipath nulls. Here two antennas are used and the strongest signal is fed into the demodulator.

Conclusion

Building a reliable RF link is not straightforward. Radio frequency signals have their own special ways of traveling from the transmitter to the receiver and have to be handled carefully. Taking everything into account, RF signals are an easy and convenient way to transmit data. Because the medium is shared, however, you need to be aware of eavesdropping and use protocols or settings that match your security policy.

ACKNOWLEDGMENTS

I want to thank Rik Farrow for his comments on this article. It is always nice to have comments from people outside the field, since engineers get dragged into the heavy stuff too easily. My employer, Competa IT, also deserves to be mentioned here for the freedom I get to write these articles.

[1] http://www.sdrforum.org/pages/aboutSdrTech/whatIsSdr.asp.

[2] http://usenix.org/publications/login/2009-04/pdfs/vandrunen.pdf.

[3] http://www.usenix.org/event/usenix03/tech/freenix03/full_papers/vandrunen/vandrunen_html/index.html.