

"Wireless Shootout: Matching Form Factor, Application, Battery Requirement, Data Rates, Range to Wireless Standard

David John Urban
Comcast

Abstract

Wireless standards are developed to meet the needs of many different applications, data rates and form factors. The IEEE 802.11ac standard is optimized for data rates greater than 300 Mbps for ranges of 1 to 10 meters with notebook computers having all day battery life. IEEE 802.15 based Bluetooth low energy is optimized for wearable devices such as watches, heart rate monitors, and stride sensors with data rates of 1 Mbps in extremely small form factors with battery life of weeks, months, and even years. Sensors that are geographically dispersed requiring years of battery life and very minimal data transfer may be best served using a lower frequency band with modulation tailored for low data rates and long range.

This paper discusses these three different use cases and describes how modulation, spectrum, and bandwidth can be matched to the unique requirements of each application and form factor

INTRODUCTION

The first campus wide wireless local area network was built at Carnegie Mellon University, named Wireless Andrew. The project was the brain child of Alex Hills and Marvin Sirbu. The first equipment was installed in 1995, using 915 MHz unlicensed spectrum and direct sequence spread spectrum providing 1

to 2 Mbps data rates. At the time this was orders of magnitude faster than the 19.2 kbps speeds available with other technologies and services.

In 1997, the IEEE standardized WLAN in the 802.11 working group and Wireless Andrew was upgraded to 2.4 GHz conforming to the 802.11 standard so that equipment from multiple vendors could work together. This is described in an entertaining book by Alex Hills that is highly recommended reading[1].

An interesting aspect of the book is Alex's path from tapping out Morse code to fellow hams halfway around the world, to running an AM radio station in Alaska, to installing radio telephones in rural Alaska, to building Wireless Andrew at 915 MHz 1 Mbps and upgrading to 2.4 GHz 11 Mbps. Each progression involved much higher data rate, Morse code to voice broadcast to two way telephone to WLAN. And each progression involved operating at a higher carrier frequency in order to accommodate the higher data rate. And yet with each progression the range decreased from half way around the world, to 50 miles, to several miles, to hundreds of feet.

Today, as we look at a range of products that have different speed and data rate requirements a useful guide is how different parts of the spectrum and different bandwidths and modulations were used for specific applications in the past.

This paper begins by analyzing the latest WLAN standard readily available for the fastest data transfers for a notebook computer. Measurements are reported of the data transfer speeds and range of 802.11ac with an 867 Mbps notebook computer in a single family home. The next section takes a look at the WLAN used in many phones and tablets. This is followed by a discussion of the personal area wireless technology for wearable devices that can pair with a smart phone or tablet. Finally, the paper takes a look at spectrum and physical layer considerations for standalone sensors that require years of battery life and can be dispersed over tens of miles.

WLAN 802.11AC FOR NOTEBOOKS

IEEE 802.11ac is a wireless local area network standard for 5 GHz unlicensed band that is optimized for data transfer speeds greater than 300 Mbps over distances from 1 to 10 meters. Notebook computers using 802.11ac with two stream capability can operate for eight hours powered by a battery. A typical battery type can be Li-Ion Polymer 7.6V, 54.4 Whr, and 7150mAh. So if the battery is pretty much drained after eight hours of heavy use, which means that the average power consumption of the device is around 6.8 watts.

To understand the importance of RF transmit power in the overall budget for power consumption consider a rough estimate of DC power into the final stage transmit power amplifier. The Wi-Fi radio will support two streams of data so two transmit power amplifiers are required.

The output power of the notebook computer used for the tests could have an output power of up to +17 dBm in the 5.6 GHz band for an 80 MHz channel according to FCC test reports. The peak to average ratio of an OFDM signal is roughly 10 dB so that the final stage power amplifier must be linear for signal envelopes up to +27 dBm, a half watt of RF power. A standard design linear class A power amplifier may have a DC to RF efficiency of 10 percent, that is 5 watts of DC power are needed to support 0.5 watts of linear RF power. With two transmit chains, the DC power needed for the final power amplifier stages would be 10 watts using a standard class A power amplifier design and 100 percent duty cycle. This will not work since it exceeds the total average power consumption of the whole computer. Class AB power amplifier designs can have DC to RF efficiency as high as 50 percent which would take the final stage power amplifiers average DC power consumption down to 2 watts. This is a little bit closer to a workable amount of battery consumption. It is evident from these calculations that efficient power amplifier design is essential to be able to transmit at the power levels required by 802.11ac. These rough estimates suggest that lowering the duty cycle and output power of the Wi-Fi final stage power amplifier is an important consideration in improving overall battery life.

Devices supporting 450 Mbps have been commercially available for several years using the 802.11n standard. Getting to a 450 Mbps transmit rate with 802.11n requires three streams of data, 64-QAM modulation, a short guard interval, and a 40 MHz channel width. The details

are described in full in references [2], [3], [4]. 802.11ac client devices typically use only two antennas compared to the three antennas used by 802.11n. This simplifies the radio circuitry and reduces power consumption. Testing in a single family home has found that it is very rare to get three streams of data during steady transmission. Testing has also shown that two streams of data are quite common at normal usage ranges. The key to increasing the data rate with 802.11ac is the 80 MHz channel width. This doubles the data rate while only increasing the SNR requirement by 3 dB. 802.11ac adds a 256-QAM modulation rate with 5/6 code rate. So what parameters are required to get an 867 Mbps data rate with 802.11ac? The channel width must be 80 MHz. The modulation must be 256-QAM with 5/6 code rate. Two streams of data must be sent. The guard interval must be 400 ns. The transmit rate of 867 Mbps is calculated in equation (1).

$$867Mbps = \frac{(8bps/Hz \cdot \frac{5}{6}) \cdot 234subcarriers \cdot 2streams}{(3.2\mu s + 400ns)} \quad (1)$$

The subcarrier spacing for all 802.11 OFDM signals is 312.5 kHz. This ensures backwards compatibility for 20 MHz, 40 MHz, and 80 MHz operation. Since the subcarrier spacing is 312.5 kHz the useful symbol time or FFT duration can be calculated as the inverse of the subcarrier frequency spacing which is 3.2 μ s. A short guard interval of 400 ns is added to the FFT duration to determine the full symbol time. The FFT size can be calculated by dividing the channel width by subcarrier frequency spacing which for an 80 MHz channel width the FFT size is 256. This means that the 80

MHz channel is subdivided into 256 subcarriers spaced by 312.5 kHz. Of these subcarriers, some are nulled near the RF carrier and at the extreme low and high parts of the spectrum. Some subcarriers are used as pilots. This leaves 234 out of the 256 subcarriers that are used to transmit data as seen in equation (1).

The wavelength of a radio wave is the speed of light divided by the frequency. For our tests the frequency was 5,765 MHz. $\lambda = 0.052 \text{ m} = 2.047$ inches. The path loss between two isotropic antennas in free space is proportional to the distance divided by the wavelength, $L = 4\pi d/\lambda$ as shown in equation (2). The free space path loss at 1 meter is 47.7 dB for a frequency of 5.765 GHz. The free space path loss at 3 meters is 57.2 dB for a frequency of 5.765 GHz. The free space path loss at 10 meters is 67.7 dB for a frequency of 5.765 GHz.

The subcarrier spacing for 802.11 OFDM signals is 312.5 kHz. The same subcarrier spacing is used for 20, 40, and 80 MHz channels. This makes it easy for an 802.11a device that is limited to a 20 MHz channel width to work alongside an 802.11ac device that can operate with an 80 MHz channel width. Since the subcarrier frequency spacing is always the same the FFT size can be calculated as the channel width divided by the subcarrier spacing. For an 80 MHz channel width the number of subcarriers is 256. The 80 MHz channel width is divided into 256 subcarriers with each subcarrier separated in frequency by 312.5 kHz.

Of the 256 subcarriers in an 80 MHz channel a few subcarriers are nulled out at the frequency extremes

and in the center in order to make reception easier. By nulling out subcarriers at the center frequency, the receiver does not have to worry about DC offset when converting the signal to baseband. By nulling out subcarriers at the frequency extremes near the upper and lower ends of the channel, the receiver does not have to worry about interference from the spectral spillover of adjacent channels. In addition to the nulled subcarriers, some subcarriers are used as pilots. Pilot subcarriers do not carry data, they send a known pattern. Pilot subcarriers are important for frequency synchronization and channel estimation. By measuring two pilots at two different symbol times, the frequency offset can be determined. The time difference between two symbols is known, this gives us the change in time. So by determining the phase difference between two pilots at two symbol times the frequency offset can be determined as the change in phase over change in time. Likewise, by looking at two pilots within a symbol at different frequencies, the receiver can estimate the frequency response of the channel in between the two pilots and apply a correction. This is called channel estimation. After subtracting out the pilots and the nulled subcarriers from the 256 point FFT, the number of subcarriers remaining that can be used as data subcarriers is 234.

Each of the 234 data subcarriers are modulated. Modulating a subcarrier means that the magnitude and phase of the subcarrier is adjusted to represent a set of bits. The modulation can be BPSK, QPSK, 16-QAM, 64-QAM, or 256-QAM. For BPSK modulation a 0 bit corresponds to a modulation of

$\cos(2\pi ft)$ and a 1 bit corresponds to a modulated signal $-\cos(2\pi ft)$. A zero bit is distinguished from a one bit because the transmitted signals are sent 180 degrees out of phase from each other. In the same fashion QPSK maps two bits 00, 01, 10, 11 into sinusoidal signals with four distinct phases, 45° , 135° , 225° , 315° . For 16-QAM things are slightly more complicated, four bits are mapped into 16 amplitude and phase combinations. Since $\cos(0)=1$ and $\sin(0)=0$ and $\cos(90^\circ)=0$ and $\sin(90^\circ)=1$ with this pattern of the cosine function going to 0 when the sine function has a magnitude of 1 repeating every 90° , it is possible to send $\sin(2\pi ft)$ and $\cos(2\pi ft)$ at the same time. These two signals can each be amplitude modulated and as long as the combination is sampled every 90° they will not interfere with each other. This is called quadrature amplitude modulation, QAM. In the case of 16-QAM, 2 bits amplitude modulate the sine wave and 2 bits amplitude modulate the cosine wave. Likewise, 64-QAM maps 6 bits to 8 possible sine wave amplitudes and 8 possible cosine wave amplitudes.

The modulation that is new in 802.11ac is 256-QAM. 256-QAM maps each subcarrier to 8 bits. 256-QAM requires a signal to noise ratio of 30 dB for a bit error rate of 1 in ten thousand. The highest modulation in 802.11n is 64-QAM so that the maximum number of bits that can be mapped onto a subcarrier using 802.11n is 6 while the maximum for 802.11ac is 8 bits per subcarrier, thus the spectral efficiency of 802.11ac can be 33 percent more than 802.11n.

The bits that map to the data subcarriers are code words. LDPC is an option for 802.11n and 802.11ac.

The forward error correction can be convolutional coding or LDPC. LDPC coding is extremely powerful.

Tests were performed in a single family home in an attempt to see just how fast files could be transferred using 802.11ac Wi-Fi with real devices. The next sections describe these tests. The transmit rate of an 802.11ac signal can be as high as 1300 Mbps. A local server running an Apache2 web server using Ubuntu Linux operating system was connected directly to the wireless router. The wireless router was capable of 1300 Mbps 802.11ac speeds. The web server was connected to the wireless router with a CAT-5 unshielded twisted pair cable with 1 Gbps Ethernet ports at the computer and the router. An objective-C program was written using the CoreWlan class

Library to measure parameters such as signal, noise, transmit rate, channel width, PHY mode. The program was run on a notebook computer with an 802.11ac 867 Mbps client station. The program transferred a 600 MB file from the server to store on the notebook computer while measuring the time before and after the download and then calculating the average transfer rate. This was done twice with a waiting period of 20 seconds between downloads. The average download speed of the two measurements was then recorded in a test array. These tests were performed at many locations throughout a single family home. Fig. 1 to 9 of this paper records the main results from these battery of tests.

In order to check the validity of the program used to make the measurements of download speed from a local web server, first direct

TCP and UDP traffic between two 3by3 802.11ac Wi-Fi adapters was measured. The download speed measured by the program included such factors as web server settings and computer hard drive speeds so it was desired to get a maximum TCP and UDP throughput measurement that did not have the same dependencies.

An 802.11ac 1300 Mbps access point was connected to an 802.11ac 1300 Mbps wireless to 1 Gbps Ethernet adapter in bridge mode. Ubuntu computers running iperf were connected to the 1 Gbps Ethernet ports of the access point and adapter. The wireless link was 3 meters long through one wall. A TCP transfer rate of 512 Mbps was observed using the iperf default client and server setting for Linux. A UDP test was performed over

an eight hour period using a 400 Mbps transfer rate. The test was done overnight so that very little competing traffic would be present. The UDP packet error rate for a 400 Mbps data rate over an eight hour period was measured to be 4.5E-4. Over short time periods there was little packet loss for transfer rates at 500 Mbps and below. The packet loss for UDP transfer rates at 600 Mbps and above was very high, typically 25 percent.




Fig.1 Measurement of RSSI.

Fig.1 shows histograms of the RSSI values measured at distances of 1, 2, 6, and 10 meters between the AP and the STA. The RSSI reading was taken from the notebook computer in between downloads of 600 MB and after a pause of 20 seconds. Typically a hundred to two hundred measurements were made at the same location while the orientation of the notebook computer was adjusted. The RSSI decreases as the distance increases and even at the same location the RSSI can vary. Notice that while the distribution of RSSI at a given location appears to follow the expected Gaussian probability density curve, there seems to be two bell curves at a given location rather than one. A theory to help explain this observation is that receive antenna gain of the notebook computer varies with azimuthal orientation resulting in two distinct RSSI levels. The variation around these two levels indicate that the signal propagation can vary over even a short time period by several dB.

Free space path loss has a $1/r^2$ power density attenuation verses distance relationship while a simple direct ray and ground reflection path loss model has a $1/r^4$ attenuation verses distance relationship. The free space path loss equation is shown in equation (2) where L is the loss, d is the distance between isotropic antennas and λ is the wavelength. These free space and the ground reflection models are represented by the straight lines in Fig.2, the top one free space and the bottom one ground reflection.

$$L = \frac{4\pi d}{\lambda} \quad (2)$$

The plot indicated that measured levels tend to fall somewhere in between these two models. For a clear line of sight path between the AP and the STA with little in the way of scattering objects, the measurement can be very close to the free space path loss model. When there is an obstructed path and many scattering objects, the two ray model gives a more accurate prediction. While it may seem discouraging that we cannot pin point the receive level any closer than 20 dB with our two models and whichever model we choose an actual measurement is almost certain to differ from our prediction, things are not so bad. We can treat this as a statistical problem which will allow us to define a required fade margin for a given desired availability.




Fig. 2 RSSI vs. Distance.

$$L = A_0 + 10n \log_{10} \left(\frac{d}{d_0} \right) + N_{\sigma} \quad (3)$$

Equation (3) shows the formula for a log normal path loss model with free space path loss of A_0 at reference distance d_0 with a path loss exponent of n and the random variable N having a normal probability density function with zero mean and standard deviation

σ . L is the path loss and d is the distance between antennas. We split the difference between free space path loss and the two ray model and select a path loss exponent of 3. This means that we expect the path loss to increase by 30 dB per decade of distance increase and 10 dB per octave of distance increase. This leads us to Fig. 3 which shows a histogram of the difference between measured values of RSSI and predicted values using a logarithmic model with a path loss exponent of 3.




Fig. 3 Histogram of model error.

The model is not very good at predicting an actual measured value of RSSI since errors range from -10 dB to +10 dB. However, the important thing is that we can calculate a standard deviation of this error of 4.5 dB. This allows us to predict that allocating 9 dB fade margin in our link budget calculations will provide 99.8% availability over a range of times and locations. This is based upon the integration of the normal distribution from minus infinity out to two standard deviations above the mean.

This model of radio wave propagation is called a lognormal path

loss model. The path loss is modeled as free space path loss at a short reference distance, in this case 1 m. The path loss exponent is applied to calculate the median path loss at a distance away from the reference point, in this case the path loss exponent that best matched measurements was 3. Measurements are expected to have a normal distribution around the median value, in this case the standard deviation was found to be 4.5 dB.

Using the values found in FCC type acceptance reports of the devices used for testing, the transmit power is around 14 dBm with 5 dBi antennas for the transmitter and receiver. The free space path loss at 1 meter for 5,765 MHz is 48 dB. Thus, the receive level at a 1 meter distance is -24 dBm. At 10 meters, the receive level will drop by 30 dB if the path loss exponent is 3, taking the median receive level to -54 dBm at 10 meters. If we add a 9 dB fade margin to account for two standard deviations of variation, then the level is -63 dBm. With an RSSI of -63 dBm, the data rate and reliability of 802.11ac has been shown to be excellent. Thus, up to a 10 meter distance we can expect a 99.8% availability with solid reliability and high throughput. Stations will operate at ranges greater than 10 meters, but there will likely be times when the reliability and throughput is sacrificed.

The guard interval for 802.11 OFDM is either 400 ns or 800 ns. The guard interval of 400 ns is called the short guard interval and the guard interval of 800 ns is called the long guard interval. Surprisingly the guard interval was one of the parameters that

was often tweaked in the adaptive modulation process.

The speed of light is 299,792,458 m/s. By applying unit conversions to the speed of light, one can see that it takes light 1.0 ns to travel 1 foot. This is a convenient fact to remember when thinking about communication systems. A time delay of 400 ns corresponds to a radio wave traveling 400 feet, perhaps reflecting back and forth from a wall that is 200 feet away. This is why the short guard interval for an 802.11 WLAN signal is 400 ns; to account for reflections from objects up to 200 feet away. The long guard time is 800 ns which can eliminate intersymbol interference from reflections off of objects up to 400 feet away.




Fig. 4 Guard Interval versus Distance.

A simple analogy from everyday life, helps in understanding the role of a guard interval in OFDM in order to eliminate inter-symbol interference due to reflections. Suppose I have appointments for 1 hour private meetings with Joe at 1 o'clock and with Sue at 2 o'clock. If Joe shows up 15 minutes late and still wishes to meet privately for one hour, I have a problem. From 2 to 2:15 I will have both Joe and Sue in my office and will not be able to have a private conversation with either. If I make the meetings 45 minutes long and I set

aside the first fifteen minutes for small talk, then Joe can show up fifteen minutes late, we can have our forty five minute private conversation, we can have small talk with both Joe and Sue in my office from 2:00 to 2:15 and Sue can still have a forty five minute private meeting. While not the most time efficient method, it is a very simple way to prevent conflict without having to worry about the precise arrival times. It works as long as no one shows up later than the designated guard time of fifteen minutes.

Fig. 4 shows the measured guard interval versus the distance between the AP and the STA. This plot indicates that a short guard interval is used a large percentage of the time for distances less than 5 meters. For distances greater than 5 meters the short guard interval is never used. Only at a 1 meter distance is the short guard interval used all of the time. This makes sense since when the STA is very close to the AP then long reflections will experience much more relative attenuation than the direct ray. At longer distances there may not be a direct ray and long reflections will be closer in amplitude to the direct ray.

Fig. 5 shows the percentage that each of the 5 modulation types were reported during the testing. 256-QAM was the modulation used for about 30% of the tests. 64-QAM was used for 58% of the tests. Adaptive modulation is a key feature in 802.11 WLAN. Since the channel conditions change rapidly and widely, sometimes the best modulation is BPSK and sometimes it is 256-QAM. The WLAN radio needs to be able to adjust automatically for the best modulation.




Fig. 5 Modulation Usage of Tests.

While 802.11 has moved on from Alex Hill's direct sequence spread spectrum used by Wireless Andrew at Carnegie Mellon in the 90's to 80 MHz wide 256 point FFT OFDM with a 400 ns guard interval, 256-QAM modulation and LDPC coding, all of the set up and signaling for 802.11 is performed with 802.11b/g signals. This is the beauty of 802.11, it can update to ever higher speeds and still work with older devices since signaling is performed using signals that every device can understand while data to a specific device is tailored to the highest speed that device is capable of utilizing. Thus, a device capable of 2 stream 802.11ac may still be using older style modulation for request to send, clear to send, and acknowledgement messages.




Fig. 6 Download Speed and Transmit Rate.

In addition to signaling needed to set up a packet transfer, each packet

has a training sequence and overhead sent along with data symbols. This means that when the transmit rate is 867 Mbps, the actual download speed will be less than 867 Mbps due to training fields, signaling, set up, and other overhead. Fig. 6 shows the download speed measured versus the reported transmit rate. As expected, the general trend is that download rate increases with increasing transmit rate and that the average download speed is less than the transmit rate. We measured around 500 Mbps TCP throughput and 400 Mbps UDP throughput without packet loss using an 802.11ac 1300 Mbps adapter at very close range. The reported speed in Fig. 6 were made using a notebook computer with an 867 Mbps 802.11ac built in Wi-Fi adapter and downloading 600 MB files from a web server on the local area network and then averaging several measurements. During the testing the competing traffic on the local area network was small compared to the greater than 300 Mbps download speeds of the notebook. However, there were other users of the wireless network during measurements and this likely accounts for some of the cases where download speeds were much lower than transmit rate.




Fig. 7 Download Speed vs. Distance.

Fig. 7 show the average download speed from a local web server to a

notebook computer after downloading two 600 MB files as a function of the distance between the AP and the STA. The download speeds were not appreciably different for distances less than 10 meters but did fall off noticeably beyond 14 meters.




Fig. 8 SNR Vs. Transmit Rate.

Fig. 8 shows the transmit rate reported by the notebook computer as a function of the signal to noise ratio. The signal to noise ratio was calculated from the RSSI and the noise measurement provided by the CoreWlan library of the notebook computer. As expected the higher the signal to noise ratio the higher the transmit rate.

WLAN 802.11N MIMO FOR PHONES AND TABLETS

Many of the latest phones and tablets use 802.11n while moving to two transmit and receive chains for MIMO operation. With two transmit and receive chains these devices are capable of two spatial streams of data sent over the same frequency and at the same time. A SISO device has a single input and a single output; only one transmit chain and one receive chain. A MIMO device has multiple input and multiple output capability; two receive chains and two transmit chains. With a rich multipath environment signal processing is able to invert the 2by2 matrix describing

the four paths between the two receive antennas and the two transmit antennas. If the 2by2 channel matrix can be inverted then the two streams of data can be calculated at the receiver. Higher SNR is required in order to separate the multiple streams of data at the receiver.

When comparing the performance of two devices using 802.11n with one device having MIMO capability, there is no change in the maximum bandwidth of 40 MHz and the maximum modulation order of 64-QAM. The only difference is that the MIMO version is capable of transmitting and receiving two spatial streams while the SISO device can only transmit and receive one spatial stream.

The doubling of the data rate can easily be observed by performing a speed test using a broadband connection of 100 Mbps. The 802.11n MIMO device recorded a download speed of 114 Mbps while the 802.11n SISO device measured a download speed of 48 Mbps. Speed test results using standard smart phone and tablet applications vary widely, but in general the MIMO device downloads at twice the data rate compared to the SISO device. This is as expected since the MIMO device can send two streams of data while the SISO device can only send a single stream of data.

From experience while using SISO and MIMO devices, there are additional benefits beyond just two stream capability from the 802.11n MIMO update. These may be chalked up to the multiple transmit and receive chains and perhaps general RF performance such as better noise figure, antenna pattern, AGC.

Watching live TV from a public outdoor Wi-Fi hotspot at a line of sight distance of around 80 meters, the newer tablet with 802.11n MIMO Wi-Fi displays high quality video without losing connection or buffering. The older 802.11n SISO device showed the live video at very low quality and had connection and buffering issues making the application unusable.

WPAN 802.15 BLUETOOTH LOW ENERGY FOR WEARABLES

Bluetooth low energy is designed for personal area network applications, PAN. This means that it is intended for communications between devices around a person, such as wireless headphones to listen to music from a phone in your pocket.

The aim of the Bluetooth low energy specification is to provide connectivity over several meters with small devices having long battery life. The battery life can be prolonged by limiting the amount of energy consumed by the device. The energy is equal to the power integrated over time. Reducing the signal to noise ratio requirements allows the transmit power to be reduced which in turn reduces the energy consumption and elongates the battery life.

Another, even more powerful technique to improve battery life is to turn the device off. Of course, a device that is always turned off is not particularly useful, at least with regards to providing us with sensor data. With a certain amount of data to be sent, it is beneficial to be able to send the data quickly so that the device can be turned off quickly. This leaves us with competing factors, low transmit power reduces energy

consumption while the device is on, but on the other hand high transmit power improves the signal to noise ratio which allows higher spectral efficiency which in turn reduces the time that the device needs to be turned on. The right tradeoff between these two depends upon the amount of data to be transmitted, among other things. For transferring GB's of data between a notebook computer and a server, the 867 Mbps 802.11ac signal allows the data transfer to occur in less than a minute. For transferring smaller amounts of data from a wearable device to a phone, the 1 Mbps transmit rate of Bluetooth low energy allows the data to be transferred quickly.

802.11ac is designed for greater than 300 Mbps data rates over ranges of 1 to 10 meters using battery powered computers that are expected to last all day on a single charge. Bluetooth low energy is designed for connectivity over a distance of 1 meter with rates around 1 Mbps using wearable devices powered by coin cell batteries with capacity of 230 mAh.




Fig. 9 BLE vs 802.11ac SNR.

The modulations used are tailored to the application and the battery

requirements. 802.11ac has modulation rates that can go up to 256-QAM while Bluetooth low energy uses Gaussian frequency shift keying. A comparison of the bit error rate versus energy per bit over noise spectral density, E_b/N_0 , is shown in Fig.9. The bit error rate for 256-QAM is 2% for 15 dB E_b/N_0 . The number of bits per symbol for 256-QAM is 8 and $10 \cdot \log_{10}(8) = 9\text{dB}$ so that the SNR is 24 dB for a 2% error rate for 256-QAM. LDPC coding which is an option for 802.11n and 802.11ac is able to correct for a 2% error rate provided that the code rate and code word length are sufficient.

For FSK using noncoherent detection the bit error rate is 2% for an E_b/N_0 of 9 dB. FSK transmits 1 bit every symbol so that the SNR is also 9 dB for a 2% bit error rate. Bluetooth low energy does not use LDPC coding, however retransmissions of packets allow good throughput even with a 2% bit error rate.

This is illustrated in Fig. 10 which was recorded while listening to music using a Bluetooth headset. The music sounded fine during the test without any disruption. While walking away from the computer sending the Bluetooth transmission the RSSI changed from -50 dBm to below -80 dBm and this resulted in the retransmission rate going from 2% to 8%.

The modulation used for Bluetooth low energy can work at 15 dB lower SNR than the highest order modulation of 802.11ac. Working at a lower SNR requires less transmit power which reduces energy consumption and extends battery life.

A Bluetooth low energy device is required to have a receiver sensitivity of -70 dBm. Bluetooth low energy operates in the 2.4 GHz ISM band. The transmit power of Bluetooth low energy devices ranges from -20 dBm to +10 dBm. There are 40 channels for a BLE device with a center frequency beginning at 2402 MHz and ending at 2480 MHz separated by 2 MHz. The wavelength for 2.4 GHz devices is 4.9 inches (4).

$$\lambda = \frac{c}{f} = \frac{299792458 \text{ m/s}}{2402 \text{ MHz}} = 12.48 \text{ cm} = 4.914 \text{ in}$$

(4)

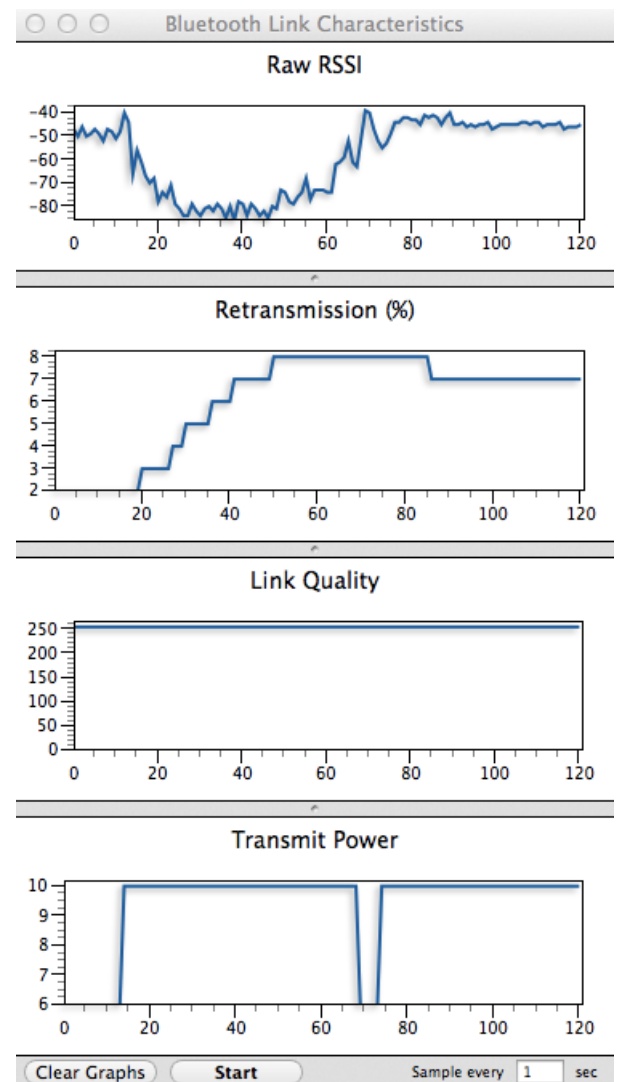


Fig.10 Retransmissions in Bluetooth.

A typical CR2032 button cell battery provides 3 volts with a capacity of 240 mAh. The CR denotes that the “Lithium coin” battery uses a chemical system of Lithium/Manganese Dioxide, Li/MnO₂. The 20 in the part number indicates that the diameter of the battery is 20 mm and the 32 in the part number indicates that the height or thickness of the coin cell battery is 3.2 mm.

The current draw for a Bluetooth low energy device is complicated. The current draw is different for various modes of operation such as sleep, wake up, transmit, and receive. The application note found in reference [8] carefully measured the current draw and time duration for various modes with an oscilloscope. As a rule of thumb, it was determined that the average current draw during continuous operation was 23 µA.

Bluetooth low energy uses Gaussian frequency shift keying modulation. This means that a zero bit is represented by one frequency and a one bit is represented by another. There are 40 channels in the 2.4 GHz band. BLE frequency hops between these channels.

$$\begin{aligned}
 h(t) &= \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{t^2}{2\sigma^2}\right) \\
 H(\omega) &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} \exp\left(-\frac{t^2}{2\sigma^2}\right) \exp(-j\omega t) dt \\
 \frac{t^2}{2\sigma^2} + j\omega t - \frac{\omega^2\sigma^2}{2} &= \left(\frac{t}{\sqrt{2}\sigma} + \frac{j\omega\sqrt{2}\sigma}{2}\right)^2 = \frac{(t+j\omega\sigma^2)^2}{2\sigma^2} \\
 H(\omega) &= \exp\left(-\frac{\omega^2\sigma^2}{2}\right) \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} \exp\left(-\frac{t^2}{2\sigma^2}\right) \exp(-j\omega t) \exp\left(+\frac{\omega^2\sigma^2}{2}\right) dt \\
 H(\omega) &= \exp\left(-\frac{\omega^2\sigma^2}{2}\right) \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} \exp\left(-\left[\frac{t}{\sqrt{2}\sigma} + \frac{j\omega\sqrt{2}\sigma}{2}\right]^2\right) dt \\
 H(\omega) &= \exp\left(-\frac{\omega^2\sigma^2}{2}\right) \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{+\infty} \exp\left(-\frac{(t+j\omega\sigma^2)^2}{2\sigma^2}\right) dt \\
 H(\omega) &= \exp\left(-\frac{\omega^2\sigma^2}{2}\right) \\
 \text{Fourier transform pair} \\
 \text{time domain } \longleftrightarrow \text{ frequency domain} \\
 \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{t^2}{2\sigma^2}} &\longleftrightarrow e^{-\frac{\omega^2\sigma^2}{2}} \\
 B \text{ is the 3 dB bandwidth} \\
 \sqrt{2} &= e^{\frac{(2\pi B)^2\sigma^2}{2}} \\
 \ln\sqrt{2} &= \frac{1}{2}\ln(2) = \frac{1}{2}(2\pi B\sigma)^2 \\
 \sigma &= \frac{\sqrt{\ln(2)}}{2\pi B}
 \end{aligned}$$

Fig. 10 Gaussian Filter Equation

The equation for the Gaussian filter that is applied to the bit stream of a Bluetooth low energy signal is shown in Fig. 10 along with the derivation of its frequency response. Fig. 11 shows the waveform at the output of the Gaussian filter with a bit stream input.

There are a couple of things to note. First, the Gaussian impulse response has the property that its Fourier transform is also a Gaussian response, both the impulse response and the frequency response follow a bell curve. The parameter σ , the standard deviation, can be calculated from the 3 dB bandwidth of the filter's frequency response and this term is in the numerator of the impulse response and the denominator of the frequency response. This gives the expected result that a narrow impulse response is required for a wide frequency response and a wide impulse response in time is required for a narrow frequency response.

For Bluetooth low energy the parameter B is set to one half the bit rate, BT=0.5 where T is the symbol period of 1 µs. The 3 dB down point of the Gaussian filter for Bluetooth

low energy is 500 kHz. Another thing to note is that a Gaussian filter is not a Nyquist filter. That means that it does not have an impulse response that goes to zero at every symbol time other than its own. A Nyquist filter eliminates intersymbol interference which is critical for high rate modulations such as 256-QAM. In order for a Nyquist filter to work in the presence of multipath, either an adaptive equalizer or a guard interval is required. A Gaussian filter reduces but does not eliminate intersymbol interference. For this reason a Gaussian filter does not support high spectral efficiency modulations. On the plus side, the Gaussian impulse response decays steadily with time so that for modulations such as FSK, there is no need for a guard interval or adaptive equalizer.




Fig. 11 Gaussian Filter Convolved With Bit Stream.

Standard Bluetooth also uses Gaussian frequency shift keying modulation. The frequency of the carrier is shifted by around 180 kHz to represent a bit. Low energy uses a longer 250 kHz frequency shift. With a 250 kHz frequency shift and a 1 MHz bit rate the two keyed signals are orthogonal. 250 kHz is the smallest frequency deviation for which the two signals are orthogonal. For this reason, this is referred to as minimum shift keying. Imagine a demodulator that splits the signal in two and runs the signal into two mixers. One mixer uses

a local oscillator locked to the frequency that represents a bit 0 while the other local oscillator is tuned to the frequency representing a bit 1. The output of the mixer is integrated over the symbol period of 1 μ s. If the frequencies match then the output of the mixer will be a DC voltage. If the frequencies do not match then the output of the mixer will have a low frequency component of twice the frequency offset, in this case 500 kHz. The period of a 500 kHz sine wave is 2 μ s. The sinusoidal signal will traverse a half cycle which will integrate to zero. The principal of minimum shift keying is illustrated in equation (5) and Fig. 12.

$$\begin{aligned} \cos(\omega t + \Delta\omega t) \cos(\omega t + \Delta\omega t) &= \frac{1}{2} + \cos(2\omega t + 2\Delta\omega t) \\ \cos(\omega t + \Delta\omega t) \cos(\omega t - \Delta\omega t) &= \frac{1}{2} \cos(2\Delta\omega t) + \cos(2\omega t) \end{aligned} \quad (5)$$




Fig. 12 Principal of minimum shift keying.

Frequency shift keying is used so that the signal can be demodulated at low input levels with poor signal to noise ratio conditions. This reduces the transmit power requirements and improves battery life while reducing the form factor. Bluetooth low energy uses a wider frequency deviation than standard Bluetooth so that the two waveforms are orthogonal. The wider frequency deviation increases the separation requirement between

channels from 1 MHz to 2 MHz and reduces the number of channels from 79 channels for standard Bluetooth to 40 channels for Bluetooth low energy. The standard Bluetooth channels are shown in Fig. 13.




Fig. 13 Bluetooth standard has 79 channels.

GFSK with minimum shift keying is a constant envelope signal so that a linear power amplifier is not required. This helps improve battery life.

WMAN INTERNET OF THINGS LONG RANGE LOW BIT RATE

The Internet of things has been earmarked by many as the next big thing in technology. Although opinions vary about future manifestations of the Internet of things, providing connectivity to things never before imagined to need connectivity is a common element. This includes things that are very far away from Wi-Fi, Bluetooth, and even cellular coverage. Examples are garbage bins, restroom soap dispensers, parking spaces, and all manner of environmental sensors. Other applications include providing wide area connectivity to Bluetooth low energy wearable devices for those times when it is inconvenient to carry your cellular phone.

Geographically dispersed environmental sensors occasionally sending data to a base station require coverage area measured in miles and battery life measured in years.

Applications for machine to machine communication along with a standard intended to ideally suit the applications is described in reference [10]. The group developing the radio interface for this application is called Weightless.

Operating in the UHF spectrum from 470-860 MHz gives the base stations greater range than the 2.4 GHz band used by Bluetooth low energy and the 5 GHz band used by 802.11ac. The wavelength at a center frequency of 470 MHz is 63.8 cm which is just over 2 ft. A half wave dipole antenna would be a foot long. At UHF band the antenna needs to be physically larger for an equivalent antenna gain compared to operating in the 2.4 GHz band or 5 GHz band.

The free space path loss at 1 GHz and 1 km is 92.4 dB. This leads to the familiar formula for path loss shown in equation (6). The path loss at 50 km which is 31 miles and 500 MHz is 122 dB.

$$L = 94.2 + 20 \cdot \log_{10}(d_{km}) + 20 \cdot \log_{10}(f_{GHz}) \quad (6)$$

The system link budget is 170 dB. For free space path loss conditions getting a 31 mile coverage range is no problem with such a large link budget.

The uplink channel width is 2 kHz. Using only a 2 kHz bandwidth improves the system link budget. The noise floor of an 80 MHz 802.11ac channel is 46 dB higher than the noise floor of a 2 kHz signal, if both have the same noise figure. The thermal noise spectral density at room temperature is -174 dBm/Hz. So with a 4 dB noise figure receiver and a 2

kHz channel width, the noise floor is -137 dBm. DBPSK, differential binary phase shift keying, is the lowest possible modulation. DBPSK requires a 10.5 dB SNR. A $\frac{1}{2}$ rate convolution code is used for error correction which has 7.5 dB coding gain. A spreading factor of 4 is used to reduce the required signal to noise ratio by 6 dB. With these three factors, the signal can be demodulated successfully at 3 dB below the noise floor, SNR=-3dB. The transmit power is 20 dBm for the uplink device. The antenna gain for the uplink device is -4 dBi and the antenna gain for the base station is 14 dBi. The noise figure of the base station receiver is 4 dB. Using these parameters the system gain is calculated to be 170 dB as shown in equation (7), that is the path attenuation allowed between the two antennas. A measurement of the spectrum showing the 2 kHz uplink channel and the 6 MHz wide downlink channel is shown in Fig. 14. The system uses time division duplex so that both the uplink and downlink use the same spectrum at different times.

$$\begin{aligned}
 G_s &= P_t + G_t + G_r + 174 - 10 \\
 &\quad \cdot \log_{10}(B) - F - \eta \\
 &= 20 - 4 + 14 + 174 \\
 &\quad - 10 \cdot \log_{10}(2e3) - 4 \\
 &\quad - (-3) = 170dB
 \end{aligned}
 \tag{7}$$

A benefit of the log-normal path loss model used in an earlier section for 802.11ac propagation is that the model is physics based with three intuitive parameters. The reference distance can be estimated based upon the expected distance for near line of sight transmission. For an indoor WLAN AP this will be around 1 meter. For a 100 foot high cellular base station antenna the reference

distance may be around 100 meters. For a 1,000 feet high broadcast tower, the free space reference distance may be around 1 km. When a strong reflection is expected the path loss exponent will be close to 4 and when a strong direct ray is expected the path loss exponent should be close to 2. The standard deviation may be close to 5 dB when obstructions are uniform throughout the coverage area while the standard deviation will be closer to 10 dB when obstructions vary throughout the coverage area.

Consider a deployment from a broadcast tower. Using a path loss exponent of 4 and a reference distance of 1 km and a 10 dB standard deviation, the maximum distance calculated with a log-normal path loss model at 470 MHz with a 20 dB fade margin is 40 km. This is shown in Fig. 15.




Fig. 14 Spectrum of WMAN IoT Signal.

In order to check the validity of the model, the calculation was compared to the prediction of TV coverage from a website using the Longely-Rice Irregular Terrain Model and the FCC data base. At a distance of 23.5 miles the predicted signal strength was -36 dBm from a broadcast station with an

output power of 739 kW and an antenna height above average terrain of 179.9 meters. This would be a path loss of 124 dB, just 5 dB below the free space path loss. In reality, the signal could be picked up with an indoor antenna only after careful positioning and even then it required re-adjusting every once in a while in order to maintain picture quality.




Fig. 15 Broadcast Tower Path Loss

Fig. 15 predicts a receive level at this distance of -62 dBm. In this particular case the log normal path loss model provided a prediction that was more accurate than the terrain based model. For prediction of UHF coverage over long distances the Longely-Rice Irregular Terrain Model and the FCC data base are critical in predicting coverage at a particular location. The log normal path loss model is of little help in predicting the actual coverage at a particular location because it does not account for the topography between the two antennas. In particular, if a hill shadows the RF signal then there is little chance of reception. On the other hand, if one can see the tower with a pair of binoculars then strong signal reception can be expected. The log normal path loss model is useful for statistically analyzing the likelihood of a system

working over a range of locations and times.




Fig. 16 Strand Mount Path Loss

Fig. 16 show that the log normal path loss model predicts a coverage of 1.2 km for expected parameters from a shorter tower on the order of 10-30 meters. The reference distance used is 1 meter and the path loss exponent is 4. The frequency is 470 MHz and the standard deviation used is 10 dB, thus requiring a 20 dB fade margin. The system gain used is 170 dB.

Note that the uplink uses a 2 kHz channel width with direct sequence spread spectrum having 4 chips. Thus, a 500 Hz rate is spread out to 2 kHz. The modulation is DBPSK with 1 bit per symbol and the code rate of error correction is $\frac{1}{2}$. Thus, the data rate is only a paltry 250 bits per second. For remote sensors, only having to occasionally report small amounts of data back to the base station, this low data rate may be sufficient. The bit rate may be low but the 170 dB system gain makes sure that the range is very long. With sensors placed throughout an area, many devices will not require the full 170 dB system gain. These devices can use higher order modulations. As with Bluetooth low energy and 802.11ac, sometimes a

high data rate can allow the device to turn off more quickly and thus preserve battery life.

CONCLUSION

Battery life is an increasingly critical parameter in consumer adoption and satisfaction of electronic communication devices. Great improvements to battery life can be expected in the coming years. Today, a notebook computer, tablet computer, and smart phone are expected to last all day before having to recharge the battery. Smart watches and fitness bands with Bluetooth low energy and Bluetooth headsets can last several days before having to recharge. Wearable devices using Bluetooth low energy can last weeks, months, sometimes almost a year before the coin cell battery needs to be replaced. Great care has been taken in order to achieve these impressive battery lifetimes. Even more attention to reducing energy consumption will be needed to meet the needs and expectations for future devices.

The two counteracting forces at work are lowering the SNR requirement to reduce transmit power and increasing data rate to reduce duty cycle. The data rate needs to be tailored to the size of data that needs to be transmitted, the range of transmission, and the form factor of the device. Computers with 120 GB hard drives need to send large data files. With an 80 MHz channel width, two streams of data, 256-QAM 5/6 LDPC coding, and a 400 ns guard interval, an 802.11ac signal can transmit a 1.2 GB file in 24 seconds. Tests have shown that files can be transferred in about this amount of time over a range of 1 to 10 meters.

Wearable devices using Bluetooth low energy need to send smaller size data files but it is still important that the time it takes to transfer data is short. A 1 Mbps data rate is realized with GFSK modulation having a 250 kHz frequency separation. The 3 dB bandwidth of the Gaussian filter is 500 kHz. No information is carried in the amplitude of the signal so that the constant envelope signal can use nonlinear high efficiency amplifiers. These and other features of Bluetooth low energy are designed to optimize battery life while still meeting the device data transfer requirements.

802.11ac uses the 5 GHz band while Bluetooth low energy uses the 2.4 GHz band. In order to collect data from millions of geographically diverse sensors having very small data transfer needs it may be advantageous to utilize UHF frequencies which have benefits for long range connectivity. By using very narrow band channels, for example 2 kHz uplink channel width, low order modulation with spread spectrum and forward error correction it has been shown that a 170 dB link margin is possible with reasonable transmit powers, noise figures and antenna gains in the 470-860 MHz UHF band. This is a highly coveted band but only a small sliver of it is needed for collecting sensor data. With the UHF band being repacked for broadcast TV and separated into FDD uplink and downlink blocks, this low power application can make use of small sections of spectrum that would otherwise need to be unused and reserved for guard bands. With 170 dB system gain, distance ranges from 40 km to over a km can be expected as antenna height goes from 1000 feet to 18 feet.

This paper has made measurements and analyzed the performance and applicability of many form factors and data rates. The channel width went from 80 MHz for a notebook computer to 2 MHz for a wearable sensor to 2 kHz for a remote sensor. The spectrum, likewise varied from 5 GHz to 2.4 GHz to 500 MHz. The data rate varied from 500 Mbps to 1 Mbps to 250 bps. It has been demonstrated that large size file transfers benefit from very high data transfer speeds using high order modulation, wide channel width and higher frequency spectrum bands. While at the same time applications requiring smaller data file transfers and smaller physical form factor can benefit from lower order modulation, narrower channel width and lower frequency spectrum bands.

REFERENCES

- [1] Alex Hills, "Wi-Fi and the Bad Boys of Radio: Dawn of a Wireless Technology", dog ear Publishing, ISBN: 978-145750-560-7, 2011.
- [2] Albano., Ong,I.,Gilson,R.,Urban. "Design of Wireless Home Network", SCTE CABLE-TEC Expo 11, November 15-17, Atlanta, GA, 2011.
- [3] David Urban, "Video Calling Over Wireless Networks", 2012 Spring Technical Forum, CableLabs, NCTA, SCTE.
- [4] Chris Albano, Dinesh Devotta, Dave Urban, "Delivering DOCSIS 3.0 Cable Modem Speeds over the Home Network", ", SCTE CABLE-TEC Expo 10, 2010.
- [5] Brittain, James E., "Harald T. Friis", Scanning the Past column, Proceedings of the IEEE, VOL. 83, NO. 12, December 1995, pp. 1674.
- [6] Matthew Gast, 802.11ac: A Survival Guide, O'Reilly, ISBN-13: 978-1449343149.

[7] Heydon, Robin, "Bluetooth Low Energy The Developer's Handbook", Prentice Hall, 2013, ISBN-13:978:0-13-288836-3.

[8] Kamath, Sandeep & Lindh,Joakim, "Measuring Bluetooth Low Energy Power Consumption, Texas Instruments Application Note AN092.

[9] Masuch,J., Delgado-Restituto,M., "A 190- μ W zero-IF GFSK Demodulator With a 4-b Phase-Domain ADC", IEEE Journal of Solid-State Circuits, Vol. 47, NO. 11, November 2012, pp.2796-2806.

[10] William Webb, "Understanding Weightless; Technology, Equipment, and Network Deployment for M2M Communications in White Space", Cambridge University Press, 2012, ISBN 978-1-107-02707-7.

ABREVIATIONS

WLAN wireless local area network

WPAN wireless personal are network

WMAN wireless metropolitan area network

OFDM orthogonal frequency division multiple access

AP access point the WLAN wireless router

STA station the WLAN client adapter

GFSK Gaussian frequency shift keying

FSK Frequency shift keying

QAM quadrature amplitude modulation

SNR signal to noise ratio

MIMO multiple input multiple output: antenna techniques used to send multiple streams of data over the same spectrum at the same time

SISO single input single output: with only a single antenna multiple streams of data are not possible

BLE Bluetooth Low Energy

QAM quadrature amplitude modulation

LDPC low density parity check coding