Phase Noise Effects on OFDM Wireless LAN Performance

This article quantifies the effects of phase noise on bit-error rate and offers guidelines for noise reduction

By John R. Pelliccio, Heinz Bachmann and Bruce W. Myers
Raytheon RF Networking

Consumers are looking to wireless data transceivers to convey all types of information. From 3G cell phones to wireless LANs, the convergence of voice, data and video is driving the demand for wireless gear that is capable of transmitting farther, faster and more efficiently than ever before.

In the wireless LAN industry, for example, the past few years have seen a migration from 1 and 2 megabit per second (Mbps) radios to the recent proliferation of 11 Mbps devices. Driven by the insatiable demand for bandwidth, manufacturers are rolling out plans for products capable of data rates as high as 54 Mbps at frequencies in the 5 to 6 GHz range. These products, based on industry standards such as the IEEE's 802.11(a) and the European Telecom Standards Institute's HiperLAN 2, use a unique and spectrally efficient modulation scheme known as orthogonal frequency division multiplexing (OFDM) to communicate.

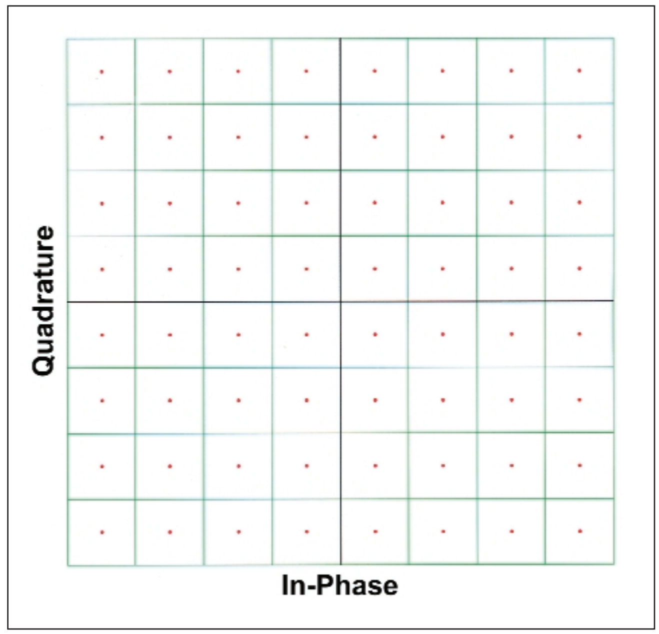

▲ **Figure 1. Constellation diagram of a 64-QAM subcarrier with no distortion.**

OFDM is essentially a series of orthogonally separated subcarriers each with its own modulated waveform. OFDM is very robust against multipath signals and amplitude and group delay variations in the channel. Although this waveform type is very well suited to an indoor environment, it presents some unique challenges to system designers. The waveform properties that most affect the analog design are accommodation of an inherently high peak to average power ratio (up to 21 dB), sensitivity to non-linearities of the analog components (i.e., gain compression and AM to PM conversion), and sensitivity to phase noise. A system designed without the appropriate amount of margin in any of these categories will generate an unacceptable number of bit errors.

An 802.11(a) transceiver uses a total of 64 separate subchannels, each spaced 312.5 kHz apart, for a total channel bandwidth of 20 MHz. Of these subchannels, 48 are used for data, while 12 are unpopulated to allow for guardbands at the channel edges, and four are reserved exclusively for pilot tones. The 48 data subchannels are each modulated independently. Modulations range from BPSK for 6 Mbps data transfer through 64-QAM for 54 Mbps service. A single QAM symbol of modulation level 2^n can

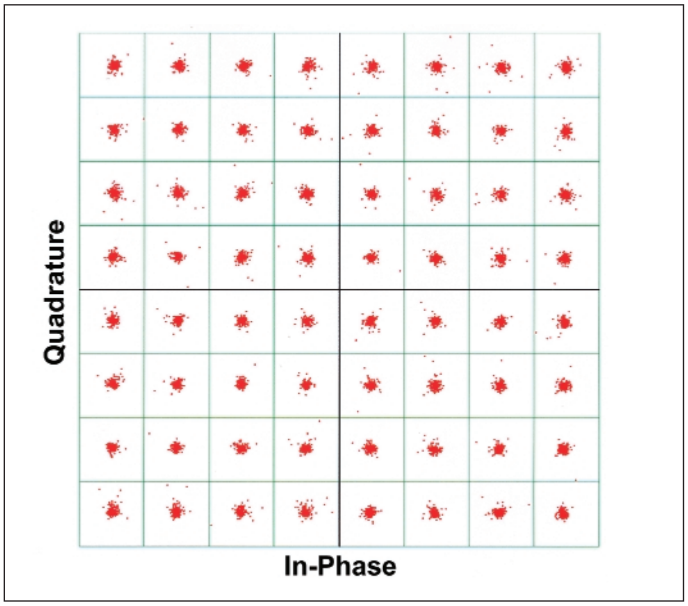

▲ **Figure 2. Constellation diagram of a 64-QAM subcarrier with distortion.**

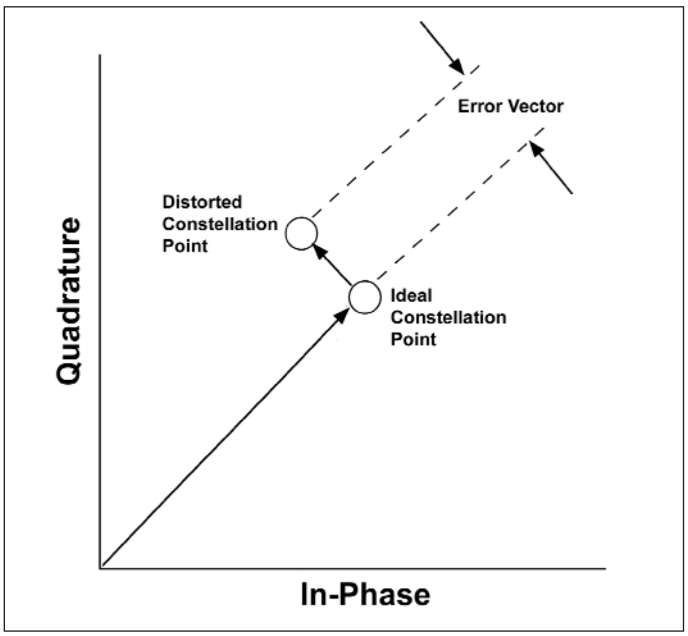

▲ **Figure 3. Constellation error vector definition.**

carry *n* bits, so one 64-QAM symbol can convey 6 bits of information. The constellation diagram of a 64-QAM subcarrier is shown in Figure 1.

The constellation diagram shows all possible states on the complex plane that a 64-QAM symbol can assume. However, additive white Gaussian noise (AWGN), transmitter and receiver nonlinearities and multipath effects affect QAM symbols. A received QAM symbol may look like the example in Figure 2.

Every QAM receiver has a processor that takes each bit and determines its position in the constellation relative to the origin and a set of predefined decision boundaries. When a symbol crosses a decision boundary, a bit error results. The system must therefore be designed so that this happens infrequently. Since the decision boundaries are placed closer together as the QAM order increases, the requirements placed on the modulated signal become more restrictive as higher order modulations are used.

In principle, the error contribution could be allocated between the transmitter and receiver in any proportion. However, 802.11(a) specifies that each transmitter must, over the course of a defined number of symbols and packets, provide an average error vector (defined as the distance from the ideal symbol point to the location in the constellation at which it is actually received) less than a certain magnitude. This magnitude decreases as the data rate (and QAM order) increase, from –5 dB at 6 Mbps to –25 dB at 54 Mbps. This "constellation error" specification allows interoperability between different vendors' products by ensuring that neither transmitters nor receivers are designed so that they produce so much error that an interoperable product will not be able to resolve the signals with a good degree of accuracy.

Several factors can impact this error vector. AM/AM conversion in the transmit amplifier due to gain compression can cause errors in the intended amplitude of the transmitted symbol. AM/PM conversion and phase noise can cause errors in the phase component of the error vector. There are two important points to note regarding the contributors to constellation error. One is that since the magnitude of the error vector is derived from a combination of all these factors, reducing the magnitude of one error contributor allows more latitude for the others. For example, if an extremely linear amplifier is used, amplitude distortion will be minor and so AM/PM conversion and phase noise can be allowed to be larger. The second important note is that while AM/AM and AM/PM conversion can be predicted and proactively corrected (for example, through digital predistortion), phase noise is by definition a random process. Thus, the amount of phase noise in a system will always have a direct and irreversible effect on the quality of the received signal.

Every active component in a transmitter and receiver can generate phase noise. However, for practical purposes the frequency synthesis components of the system tend to contribute far more noise than amplifiers and other types of circuits. Everything in the frequency synthesizer, from the reference frequency generator to the phase locked loops to the local oscillators, contributes to the overall phase noise power of the system. If the phase noise power is too high, the resultant error vector in the received constellation will be large, decision boundaries will be crossed, and bit errors will result.

Like most parameters in a system design, the amount

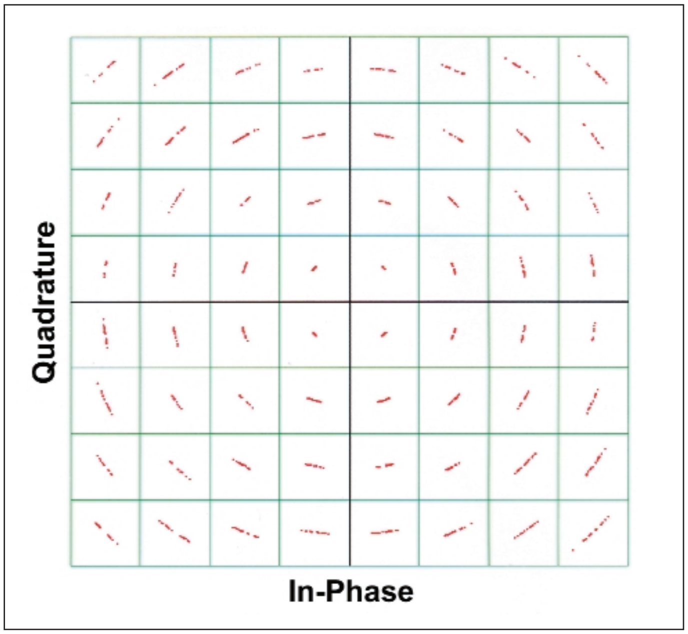

▲ **Figure 4. Constellation diagram of a 64-QAM subcarrier with phase noise.**

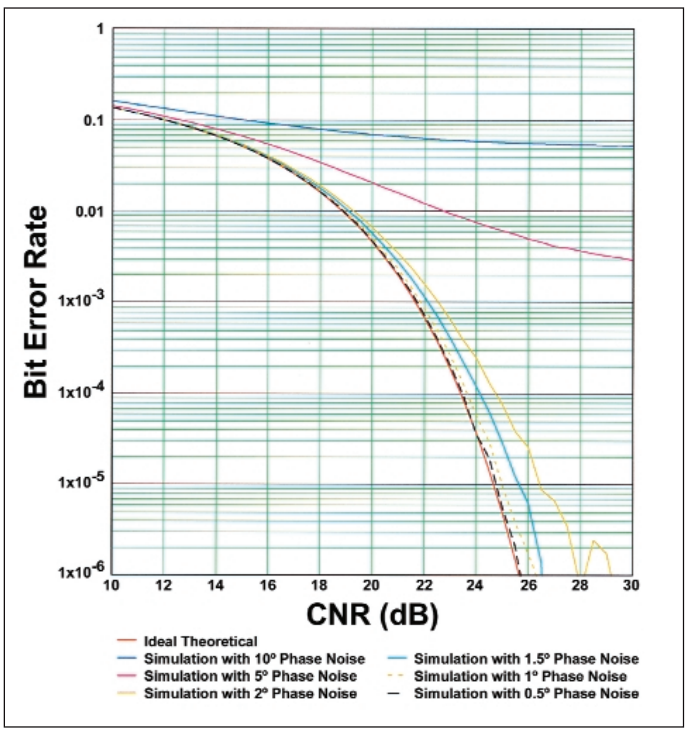

▲ **Figure 5. BER for OFDM signal with different amounts of phase noise.**

of allowable phase noise in an OFDM system becomes a question of compromise. Components with ultra-low phase noise specifications are readily available, but are often large and expensive. Conversely, trying to "over-integrate" a system to save on parts count or board space can cause problems if the processes or components used do not have phase noise specifications that will lead to an acceptable phase noise power in the final analysis.

How, then, can the effect of phase noise be determined? Phase noise has two effects on an OFDM system. The first is that it causes a phase shift in the received signal so that its constellation might appear as shown in Figure 4. The second effect of phase noise is to cause the receiver frequency reference to not align properly with the transmitted signal, causing loss of orthogonality and thereby introducing interchannel interference (ICI).

These effects are not difficult to alleviate. As discussed earlier, phase noise is mostly due to the synthesizer, with most of the noise power being near the nominal carrier frequency. To compensate for differences in the frequency sources of the transmitter and receiver, some tracking of the received signal must be employed. The tracking algorithm will also follow and thus compensate for any low-frequency phase noise. In fact, it has been determined that the presence of the frequency tracking algorithm allows us to negate the effects of phase noise located closer to the carrier than about 10 percent of the subcarrier spacing. IEEE 802.11a requires that the RF frequency and data clocks be derived from the same source. By tracking the carrier frequency, it is possible to also compensate for differences in data clocks at the same time, thus maintaining orthogonality.

It is evident, then, that the amount of allowable phase noise in an OFDM system needs to be quantified so that the frequency synthesis section of the OFDM radio is neither overdesigned nor underdesigned. Since the primary system design goal is a low bit error rate, the designer must decide upon an acceptable increase in the carrier to noise ratio (CNR) at the operating bit error rate due to phase noise. Figure 5, obtained through simulation, shows the effect of residual phase noise with a few different RMS values on an OFDM signal using 64-QAM subcarriers. Note that the CNR required for a specific bit error rate will decrease depending on the decoding algorithm employed. Figure 5 shows raw results, without the benefit of decoding.

The residual phase noise depends upon the implementation of the tracking loop. The following example may serve to illustrate the considerations needed to design the frequency sources used in the transceiver.

Consider the demodulation scheme illustrated in Figure 6 (note that this is an oversimplification presented for discussion purposes only). An OFDM signal $s(t)$ with quadrature amplitude-modulated subcarriers is passed to a tracking loop, a second order phase locked loop, with a corner frequency f_C. Because f_C is much smaller than the bandwidth of s, the VCO output is close to the ideal carrier frequency with no modulation. Multiplying this with the modulated signal and removing the high-side mixing product results in recovery of the baseband signal.

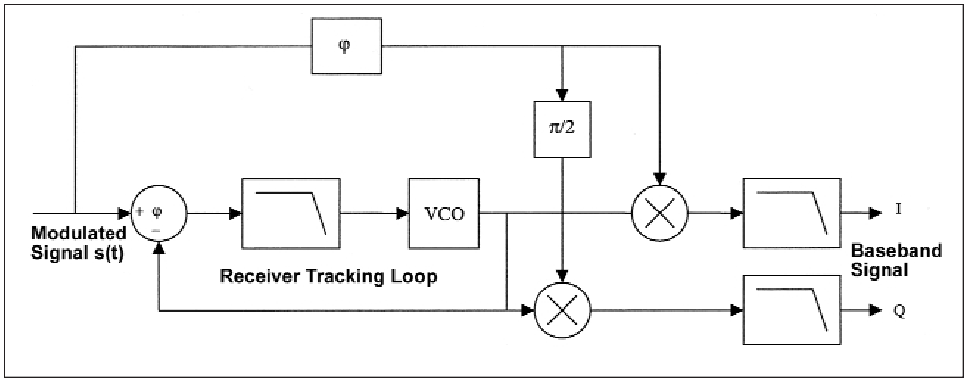

▲ **Figure 6. Demodulator functional block diagram.**

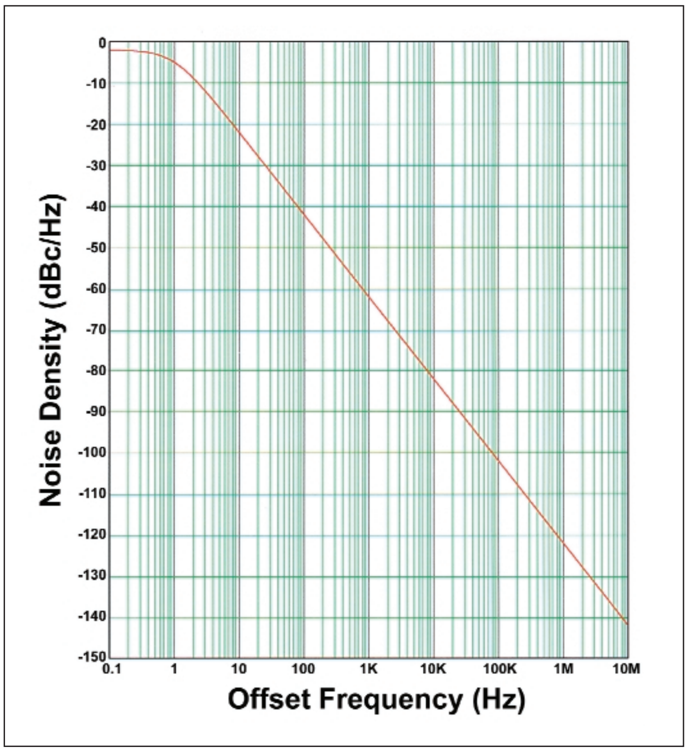

▲ **Figure 7. Lorentzian phase noise power spectrum with 1 hertz linewidth.**

For this discussion we will assume that the VCO itself does not generate any noise. Phase noise at frequencies that are smaller than the corner frequency will then be tracked by the PLL and therefore introduce no errors into the demodulation process. Phase noise outside of the loop bandwidth, however, will cause misalignment of the FFT spacing as well as phase rotations to the recovered symbols and therefore increase the bit error rate.

Let us assume that the degradation caused by 1.5 degrees of phase noise is acceptable. The question is how to design the synthesizer and select the tracking loop bandwidth so that the remaining RMS phase noise not tracked does not exceed 1.5 degrees.

We begin with the power transfer function of the second order PLL in the tracking loop that can be approximated by

$$S_{Track}(f) = \frac{1}{1+\left(\frac{f}{f_C}\right)^4} \tag{1}$$

Let us further assume that the source used to generate the carrier has phase noise that follows the Lorentzian model [1], the single-sided noise density spectrum of which is defined by

$$S_d(f) = \frac{2}{\pi \cdot f_1 \cdot \left(1+\left(\frac{f}{f_1}\right)^2\right)} \tag{2}$$

where f_1 is the 3 dB linewidth of the oscillator and f is the offset from the carrier frequency at which the phase noise density function $S_d(f)$ is evaluated as depicted in Figure 7.

Then the remaining phase noise at the output of the VCO would be

$$S_{VCO}(f) = \frac{2}{\pi \cdot f_1 \cdot \left(1+\left(\frac{f}{f_1}\right)^2\right) \cdot \left(1+\left(\frac{f}{f_C}\right)^4\right)} \tag{3}$$

At the baseband output, the noise power density affecting the demodulation would then be the difference between that of the receiver input and the tracking loop output, or

$$S(f) = \frac{2}{\pi \cdot f_1 \cdot \left(1+\left(\frac{f}{f_1}\right)^2\right)} \cdot \left(1 - \frac{1}{1+\left(\frac{f}{f_C}\right)^4}\right) \tag{4}$$

and might appear as shown in Figure 8.

Figure 8 was generated with a 3 dB linewidth of 1 hertz and a tracking loop bandwidth of 1 kHz. The remaining phase noise introduces a phase error that follows a Gaussian distribution. The RMS value for small RMS phase angles φ (that is, for φ << 1 radian) (standard deviation) can be determined in radians as

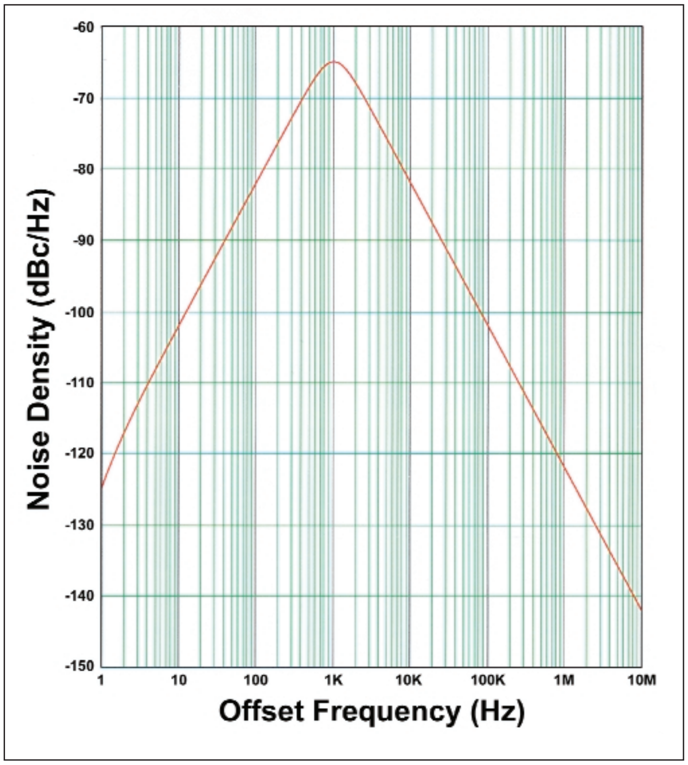

▲ **Figure 8. Phase noise power spectrum density after demodulator.**

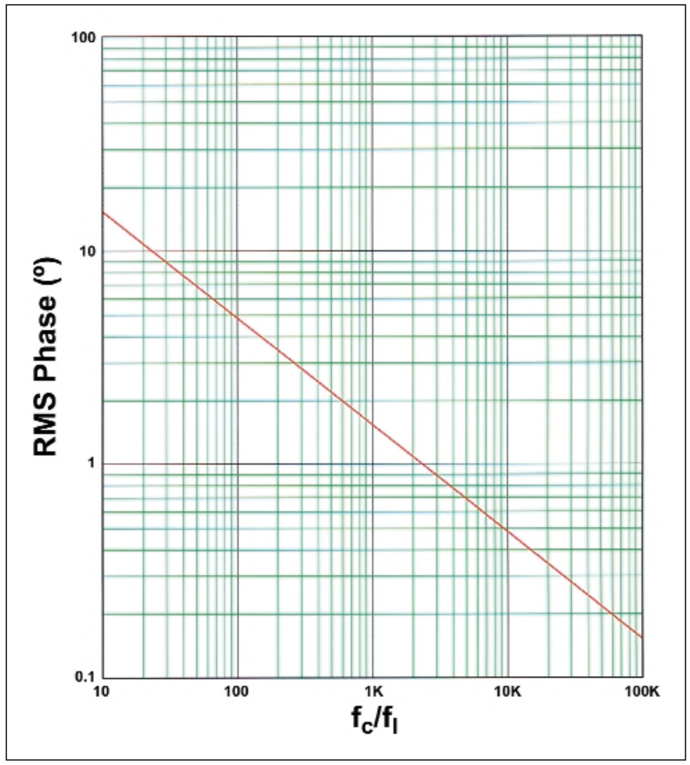

▲ **Figure 9. RMS phase noise versus f_C/f_1.**

$$\varphi_{RMS} = \sqrt{\int_0^{0.5 \cdot B} S(f)df} \tag{5}$$

where B is the channel bandwidth. Evaluation of the integral in Equation (5) is straightforward but quite tedious, and the derivation is not presented here. Although only the power within the channel bandwidth need be considered, the integral is easier to evaluate from 0 to infinity. This is a legitimate approximation because the loop bandwidth must be much smaller than the carrier bandwidth for this demodulation scheme to work, and power in the bandwidth outside of B will be very small. The so evaluated integral can be written as

$$\int_0^{\infty} S(f)df = \frac{1}{1+\left(\frac{f_C}{f_1}\right)^4} \cdot \left[1+\frac{1}{\sqrt{2}} \cdot \left(\left(\frac{f_C}{f_1}\right)^2 - 1\right) \cdot \frac{f_C}{f_1}\right] \tag{6}$$

It follows that

$$\varphi_{RMS} = \sqrt{\frac{1}{1+\left(\frac{f_C}{f_1}\right)^4} \cdot \left[1+\frac{1}{\sqrt{2}} \cdot \left(\left(\frac{f_C}{f_1}\right)^2 - 1\right) \cdot \frac{f_C}{f_1}\right]} \tag{7}$$

where φ_{RMS} is in radians. Equation (7) shows that the RMS phase noise angle is a function of the ratio of the receive tracking loop bandwidth to the Lorentzian linewidth, as shown in Figure 9.

It remains to select values for receive tracking loop bandwidth f_C and Lorentzian linewidth f_l. In general, f_C must be narrow enough so that no modulated information is lost. In the case of IEEE 802.11(a), the subchannel at the center frequency is not used and therefore contains no information. The edge of the first occupied subchannel is at 156.25 kHz. Let us assume that we want the tracking loop to suppress any modulation at that frequency by 30 dB. We determine the loop bandwidth by solving (I) for f_C, inserting 156.25 kHz for f and 10^{-3} (– 30 dB) for S_{Track}, and arrive at a value f_C = 27.8 kHz as the largest allowable tracking loop bandwidth. From Figure 9 we see that, if the phase noise into the receiver is Lorentzian, the 3 dB linewidth must be no greater than 10^{-3} the tracking bandwidth, or 27.8 Hz.

In reality, the models of the carrier phase noise and receive tracking loop are more complex. However, the following basic design rules can be used:

- Determine an acceptable tracking bandwidth for your application. Design a tracking loop with a response $S_{Track}(f)$ that provides sufficient suppression in the bands of interest.
- Design a synthesizer with a noise profile $S_d(f)$ so

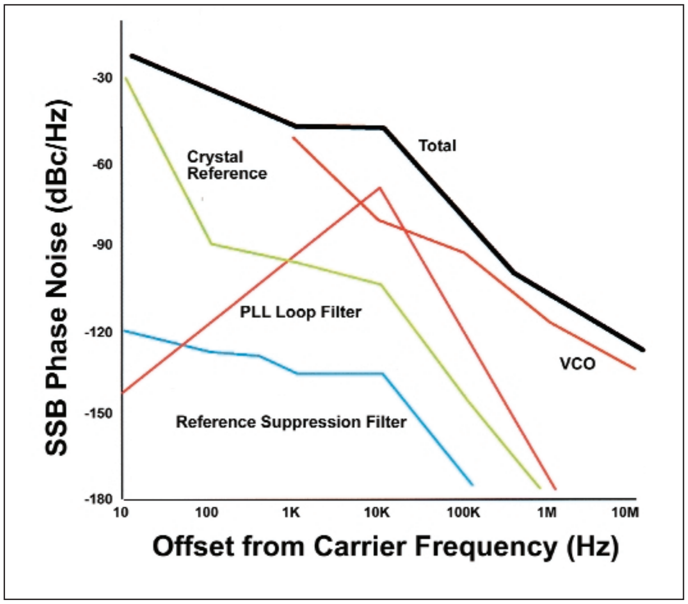

▲ **Figure 10. Phase noise profiles of frequency synthesis components.**

that the most of the phase noise spectrum falls well within the receiver tracking bandwidth.

- Determine the RMS phase noise (in radians) at the demodulator output as

$$\varphi_{RMS} = \sqrt{\int_0^{0.5 \cdot B} Sd(f) \cdot \left(1 - S_{Track}(f)\right) df} \tag{8}$$

- Determine (perhaps through simulation) whether this phase noise will degrade performance to an unacceptable level.

Equation (8) is valid only if there is no significant phase noise contribution from the receive tracking loop. This should be the case if tracking is implemented as a digital PLL.

In implementing the transmit synthesizer, care must be taken to use components that minimize phase noise far removed from the carrier. An actual synthesizer constructed of real components tends to have a phase noise spectrum dominated by VCO noise from the edge of the loop filter bandwidth to the edge of the system noise bandwidth, encompassing frequencies far from the carrier frequency. Reference frequency generators such as crystals generally possess phase noise spectra concentrated close to the carrier. While we have established that this noise will be negated by the tracking loop in the receiver, a low phase noise crystal is important for other system considerations such as meeting the 802.11a transmitter constellation accuracy requirements. The crystal frequency is generally much lower than the LO frequencies, and must be multiplied by n inside the PLL. The crystal phase noise is correspondingly multiplied by 20 log (n) where n is the multiplication factor required to arrive at the output frequency from the reference frequency. Therefore, a better approximation of a phase noise spectrum for an OFDM system would include the crystal phase noise, PLL loop response, and VCO phase noise as well as other, less prominent contributors, such as reference supression filters, summed together to generate a total SSB phase noise spectrum. A sample of such a summation is presented in Figure 10.

Care should be taken to specify components (or analog cells and processes, if synthesizer components are to be integrated into larger functional blocks) that are capable of meeting the phase noise requirements for a given system BER.

Conclusion

It has been established that phase noise causes an uncancellable and detrimental effect on the accuracy of an OFDM system, measurable as an increase in bit error rate. Simulation results depicting the effect of various levels of phase noise upon the BER of a 64-QAM system have been presented. An example of how to define the tracking filter bandwidth necessary to comply with the chosen level of residual phase noise has been presented.

Once the tracking filter bandwidth is defined, some guidelines for designing a real-world synthesizer with a noise spectrum largely within the stated tracking bandwidth are presented and some general design guidelines for choosing real-world components capable of suiting the system phase noise specifications have been established. ■

References

1. Richard Van Nee and Ramjee Prasad, *OFDM for Wireless Multimedia Communications*, New York: Artech House, 2000.

Author information

John R. Pelliccio is a senior engineer with Raytheon's RF Networking business unit in Marlborough, MA. His primary focus since joining Raytheon in 1995 has been the design, manufacture and implementation of WLAN systems. He has lectured and presented on the emergence of WLAN standards to a variety of industry and educational groups across the United States and abroad. He is the lead RF engineer for Raytheon's recently announced "Tondelayo" 54 Mbps WLAN chipset. He holds a BSEE degree with distinction from Worcester Polytechnic Institute. He may be contacted via e-mail at John_R_Pelliccio@raytheon.com.

Heinz Bachmann is a senior principal engineer with Raytheon Company in Marlborough, MA, where he has worked on systems engineering and RF design of commercial wireless products since January 2000. His interests include mathematical modeling and computer simulation of communications systems and circuits. Previously, he has worked in satellite communications for Raytheon Company (1996 to 2000) for GTE (1986 to 1993) and in commercial fixed-service point-to-point radio communications for Advanced TechCom, Inc. (1993 to 1996). He received his diploma in electrical engineering from the Federal Institute of Technology, Zürich, Switzerland in 1986. He may be contacted via e-mail at Heinz_Bachmann@res.raytheon.com

Bruce W. Myers has 23 years of experience in engineering and management of communications and radar systems. He is the director of the RF Networking business unit within Raytheon Commercial Electronics. He previously held positions in Raytheon RF Components including deputy director of engineering and manager of the Wireless Systems Department. He received a BSEE from Michigan State University in 1977 and an MSEE from Worcester Polytechnic Institute in 1983. He is co-author of five papers and a winner of the IEEE's Wheeler Award and Raytheon's Thomas L. Phillips Award for Excellence in Technology. He may be contacted via e-mail at Bruce_W_Myers@raytheon.com.

Circle 39