Bluetooth and WiFi integration: Solving co-existence challenges

As popular wireless technologies like Bluetooth and WiFi are being integrated on common semiconductor substrates and printed circuit boards, this article looks at a number of system challenges that must be overcome to ensure proper interoperability with co-located Bluetooth and WiFi functionalities.

By Jeffrey Wojtiuk

The two most popular wireless technologies are Bluetooth™ and wireless local area networks (WLAN), which is more commonly known as WiFi®. These proven technologies provide complementary means to wireless access for various applications. Bluetooth is a wireless personal area network (PAN) technology that enables short-range peer-to-peer data transfer required, for example, to enable universal serial bus (USB) cable replacement for computer peripherals, data synchronization with personal digital assistants (PDAs), and wireless headsets. By comparison, WLAN operates as an Ethernet cable replacement technology, enabling faster data transfer over longer ranges for wireless networking, Internet access, and transmission of data, video and voice over Internet protocol (VoIP). In addition to the home and office, WiFi is installed increasingly in public hotspots to enable wireless access in airports, hotels and cafes.

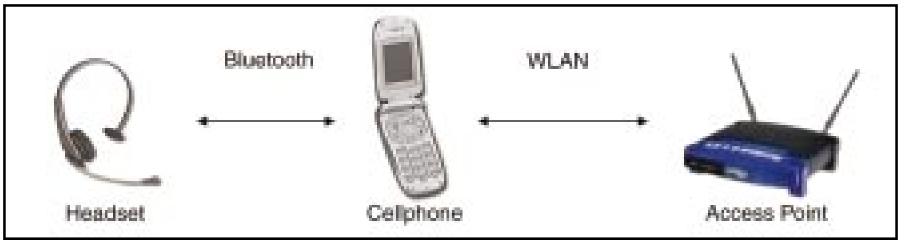

Figure 1. Bluetooth and WLAN usage example.

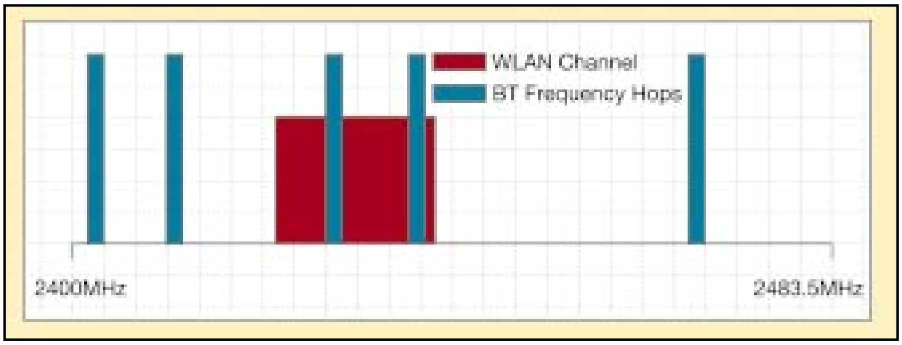

Figure 2. Spectral interference of FHSS Bluetooth on WLAN.

The growing consumer acceptance of these technologies is driving the emergence of new usage scenarios that require Bluetooth and WLAN functionality (Figure 1). These include:

- A personal computer or laptop with WLAN Internet access that also uses a Bluetooth connection to a mouse or keyboard.
- A laptop, personal digital assistant (PDA) or cellular telephone with WLAN hotspot access using VoIP and a Bluetooth-enabled headset.
- Mixed-mode devices that can support seamless handover between WLAN and Bluetooth networks.

> *Designing consumer electronics that include WLAN and Bluetooth can be quite problematic, particularly when a simultaneous mode of operation is required.*

Many of these scenarios were not foreseen by standards organizations. While the Bluetooth 1.2[1] specification includes a provision for interference mitigation, neither standard includes a provision that allows for implicit control or coordination to mitigate interference by traffic from other wireless transmissions. It is left to the designer to minimize interference and ensure that multimode end products achieve the performance, reliability and stability required by end users. Key challenges must be overcome, including real-time control of packet size, type and timing, radio and antenna isolation, optimizing error vector magnitude (EVM), amplifier linearity and efficiency.

Interference mitigation

Designing consumer electronics that include WLAN and Bluetooth can be quite problematic, particularly when a simultaneous mode of operation is required. Bluetooth and 802.11b/g WLAN systems operate in the 2.4 GHz band, and both technologies use a significant portion of the available spectrum. (Figure 2).

Interference occurs when a Bluetooth and a WLAN device are in close proximity and attempt to transmit and receive wireless signals at the same time. The two technologies use different methods for signal transmission: carrier sense multiple access (CSMA) and frequency hopping spread spectrum. The former is used by an 802.11b/g transceiver to listen for a clear channel before transmitting

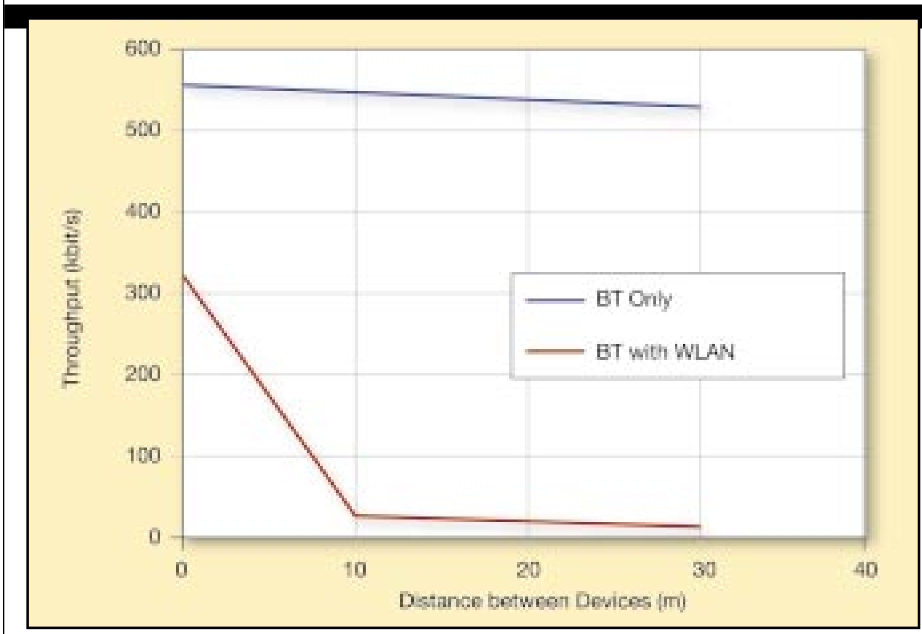

Figure 3. Bluetooth throughput, 802.11b interferer seven centimeters away.

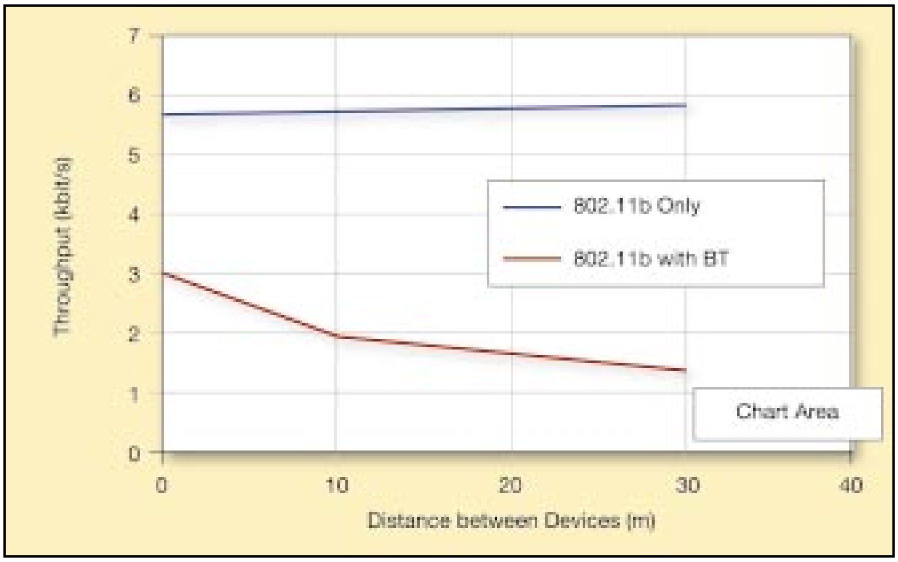

Figure 4. 802.11b throughput, Bluetooth interferer seven centimeters away.

a signal that is approximately 20 MHz wide and typically occupies one of up to three possible non-overlapping channels spaced 25 MHz apart.

In contrast, Bluetooth uses frequency hopping spread spectrum to hop between 79 1 MHz-wide channels at a rate of 1600 hops per second, delivering short time division multiplexed packets with each hop. A Bluetooth connection is made when one device initiates a connection and becomes the master of the piconet. If the address is known, the device transmits a page message. If the address is not known, an inquiry message is sent, followed by a page message. Once synchronized, two Bluetooth-enabled devices connect, with each assuming a unique media access control (MAC) address. A simple calculation demonstrates that a Bluetooth transmitter will output a signal that collides with an 802.11b/g signal approximately 25% of the time.

The effect of this co-channel interference depends on the relative strengths of both the signals and the transmit length and duty cycles of the data packets. Various analytical and simulation studies have investigated interference scenarios whereby one or both types of signals are adversely affected[2,3]. The specifications have a degree of interference resilience resulting from their signal design: direct sequence spread spectrum (DSSS) for 802.11b, orthogonal division frequency multiplexing (OFDM) for 802.11g, and frequency hopping for Bluetooth. They also use protocols based upon data packet retransmission and data rate reduction. However, these measures can result in a dramatic reduction in data throughput, which can negatively impact performance in some applications. Take, for example, Bluetooth audio transmission or VoIP over WLAN, where packet error rates greater than a few percent would cause intolerable audio delays or dropped calls.

The Bluetooth specification version 1.2 includes adaptive frequency hopping (AFH) as a way of allowing Bluetooth devices to detect and avoid interference. Using AFH, a channel can be classified as good or bad so that bad channels are avoided and replaced in the hopping sequence by pseudo-randomly selecting out of the remaining good channels. Once a Bluetooth device determines that a WLAN device is operating within the 2.4 GHz band, the frequency hopping channels that overlap are designated as bad and avoided. The use of AFH in the United States was made possible by the Federal Communications Commission (FCC) in 2002 when it began allowing frequency hopping between a minimum of 15 and maximum of 75 channels. In other regulatory domains such as the European Telecommunications Standards Institute (ETSI), frequency hopping is also allowed with a minimum of 20 channels. This is the minimum number chosen for AFH.

Co-location interference

Unfortunately, techniques such as AFH are designed specifically to detect and avoid interference for 2.4 GHz devices rather than to enable co-existence. AFH as a stand-alone technique is insufficient when Bluetooth and 802.11 devices are co-located in the same design (Figures 3 and 4). This is due in large part to the fact that WLAN devices must have higher output power to support reliable high data rate transmission of Internet, voice data and video over longer distances.

In products where WLAN is co-located with Bluetooth Class 2 or Class 3 capability, the WLAN transmitter can have as much as a +20 dB higher output power. This increases the bandwidth of appreciable interference power relative to a non-co-located scenario, thereby minimizing the number of frequency-hopping channels available, and limiting the effectiveness of spectral spreading. To put this in perspective, the sole use of AFH in a dual-mode portable terminal will not support a voice conversation using a Bluetooth headset while the 802.11b/g device is trying to upload a packet.

Another consideration is the emerging use of Class 1 Bluetooth, which also uses a power amplifier for wireless transmission over a longer 100-meter range. This specification requires +20 dBm output power, which when co-located in a dual-mode terminal, can disrupt WLAN throughput performance. In fact, even with +40 dB isolation between radios, interference from a Bluetooth Class 3 transmitter, which has 0 dBm output power can degrade the throughput of an 802.11b device.

Apart from the overlapping transmitter

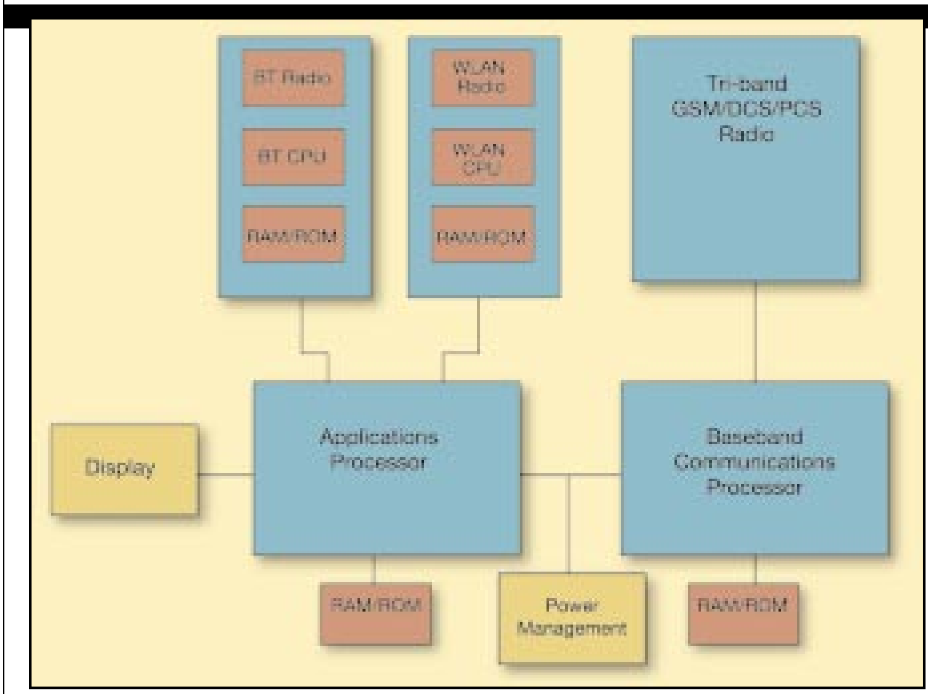

Figure 5. Typical architecture for today's multimode cellular phone integrating Bluetooth and WLAN capability.

spectrum, receiver desensitization is a serious concern. Due to the close proximity of the two transceivers, signals being transmitted from one RF chain can be large enough to saturate the receiver front-end of the other RF chain. This can occur regardless of the frequency offset, since all channel-specific filtering occurs at a lower IF, or at the baseband for the most common direct conversion receiver architectures. If transmissions between the two systems are not coordinated at a high level, then the saturation could occur anywhere in time during the reception of a data packet. During reception of a packet the receiver automatic gain control is set and the wanted signal could easily be 70 dB or 80 dB below the unwanted transmitter signal. Even if the transmitter signal itself is not a strong blocker, its wideband noise floor may be high enough to dominate the receiver noise floor.

The phase noise power of the receiver local oscillator may be strong enough at the transmitter frequency offset such that a co-channel interferer is formed on the downconverted signal by a process known as reciprocal mixing. Non-linearities in both co-located transmitters operating simultaneously can cause cross-modulation components at certain forbidden frequencies and at high enough power spectral densities to violate FCC and ETSI limits on wideband spurious generation. Similarly, transmitter harmonics can interfere with higher frequency restricted bands and 5 GHz 802.11a devices.

Isolation between co-located radios is the most effective way to reduce the level of transmitter blocker signals. This is usually limited in practice by small form factors and finite isolation effects afforded by antenna orientation and layout. Additional filtering on each transmitter is the most effective way of reducing interference levels, but this also impacts the bill of materials size and cost. The other disadvantages of extra filtering are that its losses effectively reduce transmitter efficiency and linearity at a given output power. Cross modulation components can be reduced by increasing the transmitter linearity but at the cost of efficiency, which again is undesirable for portable devices.

In many cases, the Bluetooth and WLAN functions are served with an architecture that includes physically separated baseband, RFIC, front-end and antenna blocks. The radios are on separate sides of a multilayer PCB, separated by ground plane and shielded—an approach that yields the maximum isolation. Attempts to share radio blocks such as antennas can offer a lower bill of materials but can increase the design difficulty to achieve a set amount of isolation.

This is increasingly difficult with the integration of more functionality, such as a multimode cellular phone designed to support triple-band cellular GSM/DCS/PCS operation, as well as 2.4 GHz WLAN 802.11b and Bluetooth capability. The addition of Bluetooth, WLAN and DSP-based applications can overtax the cellular baseband. To support this additional functionality while ensuring performance expectations are met, today's multimode handsets typically include an applications processor (Figure 5). The cellular, Bluetooth and WLAN radios are implemented as stand-alone modules, each controlled from either the communications or applications processor.

Recent integration efforts aim to produce a more integrated architecture, wherein the Bluetooth and WLAN radios are implemented as separate modules, each interfacing to the same central processor (Figure 6). This approach reduces the overall complexity and enhances the processing efficiency and coordination of each radio module.

In this architecture, Bluetooth and WLAN co-existence may be tackled with attention given to RF filtering shielding and isolation. To illustrate, consider Bluetooth and WLAN co-existence with cellular. The typical GSM/PCS receiver standard reference sensitivity is –102 dBm, but standalone competitive designs are often at –108 dBm or less. Receiving GSM voice at this level requires a signal to noise ratio (SNR) of about 6 dB.

Attempts to share radio blocks such as antennas can offer a lower bill of materials but can increase the design difficulty to achieve a set amount of isolation.

What is the effect of a 2.4 GHz 0 dBm Bluetooth transmitter on a 1900 MHz band PCS receiver? The interferer source will be wideband phase noise, and can be assumed to be flat due to the large carrier offset in the region of 400 MHz. Typical CMOS VCO noise floors can range from –155 dBc/Hz to –120 dBc/Hz depending on design constraints and IC process technology. It can be expected that a low-cost Bluetooth transceiver will be on the high side of this range. From this the filtering requirements can be determined. Let's assume a –130 dBc/Hz noise floor, a 180 kHz bandwidth for PCS/GSM, and a 12 dB noise contribution margin to minimize SNR degradation. The required phase noise contribution at the PCS receiver would have to be less than –102 dBm –10 log (180 kHz) – 6 dB – 12 dB = –172.5 dBm/Hz. The RF selectivity requirement to minimize Bluetooth phase noise at the PCS receiver is 42.5 dB for a 0 dBm transmitter. The selectivity requirements could be up to 20 dB tougher for a Class 1 20 dBm Bluetooth transmitter.

Now consider the effect of PCS transmit noise on a Bluetooth receiver. The Bluetooth standard reference sensitivity is –70 dBm, a typical competitive sensitivity is –80 dBm, and a SNR of 20 dB is required to receive the

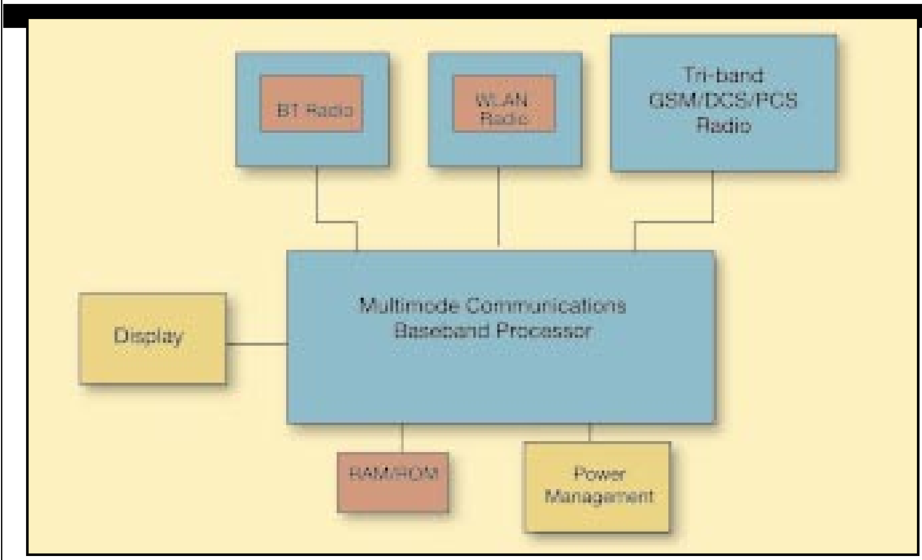

Figure 6. Integrated architecture for next-generation multimode cellular phones.

Bluetooth signal in a 1 MHz bandwidth. The PCS band VCO is toward the better end of the CMOS VCO performance due to tougher standards compliance to minimize receiver blocker interference and transmit noise in the GSM/PCS receive band.

The noise contribution to the Bluetooth receiver should be kept to less than –70 dBm—10 log (1 MHz)–20 dB – 12 dB = -162 dBm/Hz. For a 1 Watt PCS band transmitter with a VCO noise floor of –155 dBc/Hz, the RF selectivity requirement is 30 dBm–155 dBc/Hz –(–162 dBm/Hz) = 37 dB.

> *The most effective means to providing performance levels approaching no-interference scenarios is MAC-level switching.*

Other interference filtering requirements may be determined in a similar manner. However, it is important to note that designing for simultaneous operation under the cost, power consumption, and size constraints is often not possible by good RF practice alone. Because it is not accounted for in the radio standards, it has been up to the manufacturers to introduce their own proprietary for coordination and control of both radios.

Coordination measures for co-existence

Various proprietary mechanisms exist to coordinate radio activity to prevent simultaneous operation of the co-located devices[4,5]. The approaches vary in detail but essentially act to interleave operation in order to make the operation appear simultaneous. The techniques address the scheduling and priority setting of the two systems, making trade offs on transmission duty cycle, idle times, and packet type (data/beacon/paging). The packets of one system can be sent while the other is idle and vice versa; the end effect is to deliver reliable communication on both systems with a negligible loss of throughput. Three examples of coordination approaches are highlighted below.

Dual-mode radio switching is the easiest co-existence mechanism to implement. This method works by completely suspending the operation of one radio while the other is operational. There are two ways to implement this. The first approach requires turning off the non-operating radio without signalling to other nodes in the network. The drawback of this approach is that it can reduce performance in some cases below that which is achieved with no coordination. The second method signals other network nodes that the radio is suspended. The reported performance is better than shutting off the radios, but is still about 60% lower than that of unhindered radios. Neither method supports switching while Bluetooth voice links are in operation.

Driver-level transmit switching generally describes an approach in which application transmit requests are mediated at the driver level, thereby avoiding simultaneous transmission. This approach degrades throughput because only one radio is active at a time, and it causes problems with packet collisions. As a result, systems using driver-level transmit switching can suffer from dropped packets, interference and potential user difficulties when simultaneously transmitting one protocol and receiving another. As with dual-mode radio switching, this approach does not switch quickly enough to support Bluetooth SCO links, and will also have difficulties mitigating the interference from Bluetooth piconet master/slave polling activities.

The most effective means to providing performance levels approaching no-interference scenarios is MAC-level switching. MAC-level switching is a collaborative technique, which exchanges information between the two protocols at the MAC level, and manages transmit and receive operations. Since MAC-level switching is performed in the baseband, it is able to switch between protocols at a much faster rate than driver-level approaches. MAC-level switching does not suffer from transmitting signals into incoming receptions, Bluetooth polling or operating system latency.

Conclusion

The interference and co-existence issues that exist with co-located Bluetooth and WLAN devices are amongst the most difficult in radio design because of the low power, cost and size constraints. Since neither standard contains a method for coordination, a mix of RF design practices and proprietary coordination mechanisms are used to ensure interoperability with negligible performance impact to the end-user. RFD

References

1. Bluetooth SIG Core Specification version 1.2, 05 November 2003.
2. S. Shellhammer, Packet error rate of an IEEE 802.11 WLAN in the presence of Bluetooth, IEEE P802.15 Working Group Contribution, IEEE P802.15-00/133r0, Seattle, WA (May 2000).
3. O. Bar-Shalom, G. Chinn, K. Fleming, U. Gadammsetty, On the Union of WPAN and WLAN in Mobile Computers and Hand-Held Devices, INTEL Technology Journal, Volume 7, Issue 3, August 2003.
4. Bluetooth and WiFi Coexistence Overview White Paper, 2003, www.siliconwave.com.
5. How 802.11b/g WLAN and Bluetooth Can Play Together, White paper, 2004, Philips Semiconductor.

ABOUT THE AUTHOR

Jeffrey Wojtiuk is a systems architect at SiGe Semiconductor Inc. who has more than 12 years of experience in radio design for the wireless communications and semiconductor markets. Most recently, he was an RF consulting engineer with Tality and Symbionics in Cambridge, United Kingdom, where he was involved in various design services projects including multiband transceivers for GSM/EDGE, fractional-N synthesizers, and base station power control. Wojtiuk holds a PhD, master's degree and a bachelor of engineering degree with honors from the University of South Australia.