How 802.11b/g WLAN and Bluetooth Can Play

Without standards-based solutions, ICs must referee spectrum rivalry

A WHITE PAPER

INTRODUCTION

The rapid growth of wireless networking is enabling a revolution in the way information can be delivered – literally – to a user's fingertips. This, in turn, is spawning a wide range of applications and use scenarios that were not anticipated by standards-making bodies. So it is not surprising that some use scenarios – particularly those related to handheld devices – are creating serious problems for system design engineers.

Two of the most popular wireless standards – IEEE 802.11b/g wireless LAN (WLAN) and Bluetooth – absolutely must play together because they are increasingly co-located in the same handheld device. But attaining the happy marriage of these two popular technologies is quickly becoming a problem because of four interrelated facts:

1 IEEE 802.11b/g and Bluetooth use the same 2.4 GHz ISM frequency band (although they use different access mechanisms).
2 Standards bodies did not fully anticipate the range of scenarios in which WLAN and Bluetooth would compete for the same spectrum. As a result, they did not include comprehensive, robust, and cooperative mechanisms in their respective standards to mitigate interference.
3 Intense bandwidth utilization when Bluetooth and WLAN transceivers are simultaneously in use can easily overwhelm the error corrections mechanisms implemented in the two standards to manage typical interference scenarios.
4 Co-location (the presence of Bluetooth and WLAN in the same handheld device) adds the problem of desensitized receivers to the basic spectrum access issues.

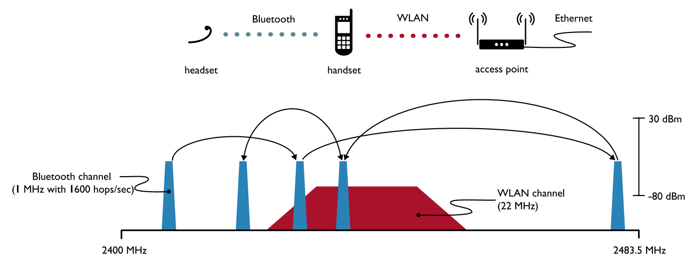

Figure 1 illustrates the overlapping signal channels when Bluetooth and WLAN are both in use.

Making Bluetooth and WLAN play together is easier said than done – but it can be done by carefully analyzing the interference scenarios and devising solutions based on the scenarios. It quickly becomes apparent, however, that the two devices (WLAN and Bluetooth) must have a means of communicating with each other in order to design a complete, effective solution.

Since no mechanism for exchanging signal status information has been built into the two standards, the job of minimizing interference must be accomplished by other means. Although techniques have been developed that do not require communication between the two devices (such as Bluetooth's Adaptive Frequency Hopping), the most efficient and comprehensive solution for the most serious problems can be accomplished by silicon vendors. They can implement information exchange capabilities within the designs of the Bluetooth and WLAN chip sets.

An important corollary to this solution is that lacking a standard protocol, only IC vendors such as Philips Semiconductors who offer both Bluetooth and WLAN chips can currently provide this solution. But before exploring the multifaceted interference problem in depth, it is important to understand the circumstances and the use environments that were, for the most part, unanticipated.

PHILIPS

Complementary use scenarios

The primary reason for the scenarios that require WLAN and Bluetooth capability in the same handheld devices is simply that the two technologies were designed for different applications. Bluetooth is most frequently used in one of three modalities: as a wireless headset; to move a data file to a Bluetooth-enabled rendering device (i.e. printer, TV, etc.); and, to exchange/synchronize data with a PC.

The rationale for creating the WLAN standard sprung from a very different environment – the networked office. In addition to being used in personal or corporate networks, WLAN is also gaining popularity as a high-speed, convenient connection to the Internet at hot spots such as coffee shops, hotel lobbies and airports.

The success of both standards makes it imperative that systems houses that design mobile phones, PDAs and portable CE devices and want to deliver the full functionality just described – short distance peer-to-peer connections and network connections – must incorporate Bluetooth and WLAN.

Additional WLAN/Bluetooth scenarios are not difficult to predict, including support for a Voice over Internet Protocol (VoIP) connection over WLAN simultaneously with a Bluetooth connection to a headset. Other scenarios would include a seamless handoff of a WLAN network connection to a Bluetooth network connection and environments that have a mixed setup.

Mitigating coexistence interference

Mutual Bluetooth/WLAN interference is the most likely interference scenario. It occurs when the two devices are attempting to transmit and receive in the same geographic area even if the transmitters are not located in the same handheld device.

Bluetooth uses frequency hopping as its access mechanism. The Bluetooth transmitter hops between 79 1-MHz wide channels with 1600 hops per second. On the other hand, 802.11b/g uses the listen-before-talk (Carrier Sensing Multiple Access) mechanism. It employs three specific, non-overlapping, 22-MHz wide channels. As a result, there is a 27.8% chance (22 divided by 79) that Bluetooth will attempt to transmit inside a WLAN channel. Depending on the relative strength of the WLAN signal, either the Bluetooth signal or both the Bluetooth and WLAN signals will be corrupted and the receiver will not be able to decode the data.

If the traffic in the cell is light enough, the problem can be corrected within the existing standards. Both standards define error correction protocols that basically require the system to fall back and retransmit the data. If, however, traffic in the cell is heavy, retransmission only makes matters worse by creating even more traffic. Valid data throughput can drop dramatically.

It's not too difficult to fashion an effective solution to the coexistence problem even in heavy traffic situations. Once Bluetooth determines (by the degree of interference it encounters) that a WLAN is using a specific 22 MHz of the ISM band, it can avoid that particular band. This technique is called Adaptive Frequency Hopping (AFH) and the Bluetooth Special Interest Group (SIG) has extended its standard to accommodate AFH in Bluetooth Version 1.2.

Collocation interference – a tougher nut to crack

An additional problem arises in the increasingly common scenario in which both WLAN and Bluetooth are integrated into the same mobile phone or PDA, which puts the two transceivers in very close proximity. Since the low-noise receiver amplifier is situated in front of the channel filter in the signal chain, it must process all incoming signals from the entire ISM spectrum band. In other words, the receiver is very sensitive to all incoming signals independent of channel.

Collocation interference arises because of the proximity of the two transceivers. Signals being transmitted from one device cause the LNA of the other device to saturate and its receiver becomes desensitized. For a short period of time, it becomes incapable of successfully receiving a legitimate signal, which by the way, can be more than 10 million times weaker than the signal of the other transmitter.

It becomes a design imperative therefore, to create a means of avoiding having one system transmitting while the other one receives. Another important problem can occur if both systems are transmitting at the same time. In this case, non-linearities in the setup could cause cross modulation of the two transmit signals. The generated spurs might cause the combined device to fail FCC certification.

The IEEE 802.15.2 recommended practice provides four mechanisms that can be used to avoid simultaneous transmission but it does not mandate that any be used. The mechanisms fall into two categories: collaborative and non-collaborative.

- Collaborative mechanisms require that Bluetooth and WLAN exchange information when accessing the medium. Packet Traffic Arbitration (PTA), Alternating Wireless Medium Access (AWMA) and Deterministic Spectral Excision (DSE) are all collaborative mechanisms proposed in the IEEE 802.25.2 recommendations.

- Non-collaborative mechanisms require that Bluetooth or WLAN take independent means to avoid interference. Adaptive Frequency Hopping (which was discussed earlier for solving the co-existence problem) is a non-collaborative mechanism. Bluetooth takes all the evasive action by itself – but alone it cannot solve the transceiver desensitization problem.

AWMA is a simple procedure that divides the time interval for transmission and reception into a Bluetooth interval and an 802.11 interval. Although simple in concept, this mechanism has the substantial disadvantage of requiring in-field upgrades to all Access Points already in use.

DSE starts from the perspective that Bluetooth be considered a narrowband interferer for the 22 MHz wide 802.11b/g signal. It is a suppression technique that puts a null in the 802.11b/g's receiver at the frequency of the Bluetooth signal. Since Bluetooth is hopping to a new frequency for each packet transmission, however, the 802.11b/g receiver needs to know the frequency hopping pattern as well as the timing of the Bluetooth transmitter. This additional signal processing adds significant complexity to existing WLAN receivers and does not solve the collocation issue.

PTA implements a control supervisor that acts like a time division traffic controller between the WLAN and Bluetooth MACs. The control supervisory circuits ensure that each MAC protocol uses a handshake mechanism to authorize its transmission before actually sending out information. A dynamic algorithm schedules the traffic. The device needs to have knowledge of time-frequency collisions. It can be implemented without encountering the problems just described for AWMA and DSE.

Philips Semiconductors has implemented its collocation solution primarily through PTA. The solution requires physical connections between the Bluetooth and 802.11b/g chips as well as software modifications. But since the problem is caused by transmitters that are co-located in the same handheld device, there is no requirement for altering equipment already in the field.

Additional dedicated logic in the form of a state machine is needed in the WLAN chip because there are only a few microseconds from the time that notice is received that a frame is being received or sent and the time that the WLAN transceiver actually sends or receives a frame. In other words, software alone cannot react quickly enough.

To accommodate all possible receive/transmit scenarios, including fragmented frames, frame bursts, and Bluetooth's master/slave modalities, a total of four signal paths are set up between the Bluetooth and WLAN chips. Figure 2 shows the chip configuration and control signals RXIND (WLAN receive indicator), WL (Wireless LAN arbitration signal), BT (Bluetooth arbitration signal) and PRI (Bluetooth priority indicator for very important traffic).

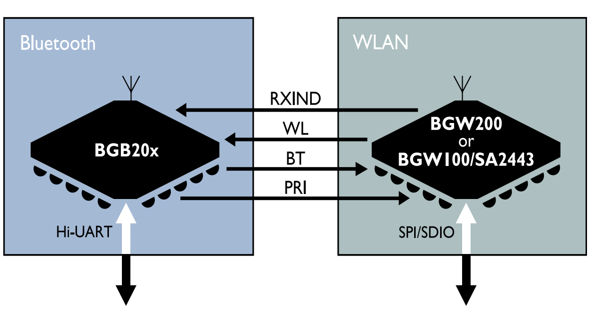

Figure 2 Packet Transmission Arbitration (PTA) is implemented by four control signals

How it works

BT and WL are the basic arbitration signals. Its simplest arbitration proceeds as follows:

If Bluetooth wants to transmit, it raises the BT signal to indicate its intentions to the WLAN chip. Once Bluetooth has raised the signal and determined that WLAN does not want to access the medium, it has ownership of the band. WLAN access to the medium works in the reverse manner. It raises the WL signal and verifies that Bluetooth does not want access at that time.

Once one system or the other has ownership of the medium, the other chip needs a means of requesting that ownership be transferred. This is accomplished by raising its signal (BT or WL) and waiting until the other system lowers its signal. The timing diagram in Figure 3 shows the states of the BT, WL and RXIND signals during a typical arbitration sequence.

- Bluetooth and WLAN reserve the medium before transmitting to avoid simultaneous transmissions
- WLAN can reserve medium during backoff sequence

WLAN transmission sequence

To react in a timely fashion to the previously noted short interval between intent and actual transmission on the WLAN side, Philips added hardware that is used in the following fashion: When WLAN wants to transmit a frame, it programs the state machine to raise the WL signal either at the beginning of the backoff sequence, sometime during backoff, or just before the WLAN transmits the frame. This scheme gives the solution maximum flexibility but does not require direct software interaction.

Once the WLAN frame has been transmitted, the WL line stays high until either an acknowledgment (ACK) frame has been received or the ACK window has timed out. If no additional frame is waiting to be transmitted, WLAN releases the WL signal. To handle cases such as fragmented frames, burst frames, or other instances in which additional fames should be sent immediately, the WL line can be programmed to stay high. This, of course, raises the possibility of shutting out Bluetooth for an extended period. To eliminate this possibility, the number of frames for which the WL signal stays high is programmable.

WLAN reception sequence

As difficult as it is to predict when a WLAN transmission will take place, it is even more difficult to predict when a reception is imminent. In addition, it is not known whether the incoming frame is addressed to the collocated WLAN device or to some other network node. Philips provides two mechanisms to facilitate this arbitration.

The first mechanism uses the RXIND signal to indicate reception is underway by setting RXIND high at the moment a WLAN preamble is detected. Bluetooth must respect this signal. The second mechanism determines whether the frame is addressed to the node co-located on the handheld device – in other words, a unicast frame addressed to the co-located MAC's address or a broadcast frame addressed to all MAC addresses.

Once this identification is made, the WL signal is automatically set high. If the received frame must send an acknowledgment of receipt, the WL signal is automatically extended until the end of the ACK frame. If a fragmented frame or a burst frame is received, WL stays high until the end of the frame exchange sequence. On the other hand, if the received frame is a broadcast frame, or, if the frame contains errors, WL is de-asserted.

Bluetooth reception and transmission

Whereas dedicated logic and revised software are needed to solve the co-location interference problems for WLANs, only software upgrades are required for Bluetooth in the Philips solution. The remaining signal – the PRI (Bluetooth priority indicator for important packets) is under Bluetooth control and is at the top of the signaling priority hierarchy. Philips has also taken steps to accommodate the dual (master/slave) personality that is possible for Bluetooth nodes.

If the co-located Bluetooth device is a master, it first checks the RXIND signal before it attempts to transmit a packet. If the RXIND signal is not set high, it asserts the BT signal. If RXIND is set high, Bluetooth waits one slot (or hop) if the RXIND signal is asserted and repeats the process until it can assert the BT signal. The delay allows the WLAN device to protect a possible receive frame by raising WL when the MAC address has been received. It also allows Bluetooth to access the medium if the WLAN frame was addressed to a different WLAN device.

If the collocated Bluetooth device is the Bluetooth slave, problems similar to the WLAN receive problems must be addressed. The Bluetooth device does not know whether it is about to receive a frame. In order to protect a possible frame, the collocated Bluetooth device asserts the BT line before the RX slot. If no packet is detected it immediately de-asserts the BT line again. If a packet is detected, it keeps the BT line asserted to protect the remainder of the packet and reserve the airwaves for the transmit slot.

The PRI signal is used to indicate very important Bluetooth packets and, as the name suggests, is accorded priority status. These packets require guaranteed delivery within a specific time frame. An example would be voice connections to a headset. Voice distortion results if all of these packets are not received and processed in a timely fashion. When PRI is set high, WLAN will stop any ongoing transmission.

Another instance in which PRI can be set high occurs when Bluetooth uses its inquiry mode to assess if other Bluetooth devices are nearby. Although Bluetooth nodes will continue to try to connect, the protocol specifies a time-out sequence in which the node simply stops trying to make the connection after a specified period of time.

The reason for the high Bluetooth priority therefore is to avoid situations in which the user knows there is another Bluetooth device in the vicinity but cannot make a connection with it. The WLAN is required to stop any ongoing transmission because in some scenarios such as a long data stream transmitted between WLAN nodes, a Bluetooth time out could occur.

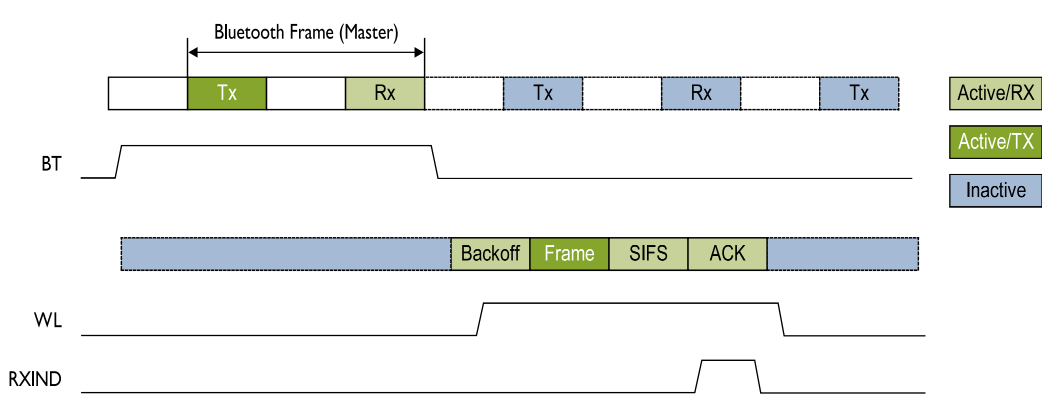

Figure 3 Timing diagram for a typical arbitration sequence

Conclusion

Bluetooth and WLAN technologies were developed to serve two quite different applications: Bluetooth offers a quick way for systems such as headsets and printers to interface with a PC or handheld device; WLANs offer wireless alternatives to local-area networks. They use the same spectrum to provide their different, highly valued applications. So it is inevitable that they will be used in close proximity and even integrated into the same handheld device or portable data terminal.

Standards-making bodies never fully envisioned the interference scenarios that have occurred as the two wireless technologies operate in a collocated, coexistence environment. Normal error-correction techniques can easily be overwhelmed in high-traffic environments.

Philips Semiconductors has provided the industry with a solution that enables communication between its WLAN and Bluetooth chips when they are collocated in the same portable equipment. By offering both Bluetooth and WLAN chips, Philips can implement effective and efficient co-existence solutions that leverage chip-to-chip communication (collaborative) in addition to independent interference avoidance mechanisms (non-collaborative).

Philips Semiconductors
Philips Semiconductors is one of the world's top semiconductor suppliers, with 20 manufacturing and assembly sites and a sales organization that delivers in 60 countries.
For a complete up-to-date list of our sales offices please visit our website http://www.semiconductors.philips.com/sales

date of release: september 2005
document order number: 9397 750 13426
Printed in the Netherlands