SPLY003 – September 2003
White Paper

WLAN Channel Bonding: Causing Greater Problems Than It Solves

Several enhancements have been proposed for increasing the throughput of 802.11 wireless local area networks (WLANs). Some of these suggested enhancements, such as increasing the modulation density, improvements to the media access (MAC) layer and techniques like multiple in, multiple out (MIMO), show real promise. But the use of channel bonding is short-sighted because it could sacrifice long-term WLAN market growth for short-term increases in throughput.

Overview

Unlike standard 802.11a/g Wi-Fi® networks which implement three separate 20 MHz wireless carrier channels, a WLAN using channel bonding would combine two of these carrier channels into one 40 MHz channel to increase throughput. Unfortunately, this type of arrangement leaves clients and access points (AP) extremely susceptible to adjacent channel interference. More traditional 802.11 implementations with super heterodyne and direct conversion radios perform better in this regard.

With channel bonded implementations, the amount of WLAN throughput lost to the effects of adjacent channel interference may not be significant during the early phases of 802.11 WLAN market growth, but as the market grows, WLAN implementations will become denser and interference problems will become apparent. Even in traditional 802.11 implementations interference generated by an AP can affect other APs close by. Channel bonding exacerbates this problem to the point where it could have detrimental effects on market growth.

A Brief Review of Standard 802.11 Implementations

The IEEE 802.11a/g standard specifies three separate 20 MHz channels and Orthogonal Frequency Division Multiplexing (OFDM) waveforms (Figure 1).

At the peak processing rate of 54 megabits per second (Mbps), the data stream is block encoded with a Viterbi trellis foreword error encoder and multiplexed into 48 sub-channels. The sub-channels are modulated using 64 QAM. The modulated sub-channels are multiplexed with four pilot tones to aid demodulation and 12 zero-filled guard band sub-channels. These sub-channels are then block processed using a 64-point Fast Fourier Transform (FFT), and they then are serially presented as samples to the RF modulator after passing through the baseband digital-to-analog converter (DAC). 802.11a/g waveforms occupy the standard 20 MHz WLAN or Wi-Fi channel. The products of the FFT modulation occupy approximately 17 MHz of this channel because of the 1.5 MHz guard band at either end of the modulated waveform.

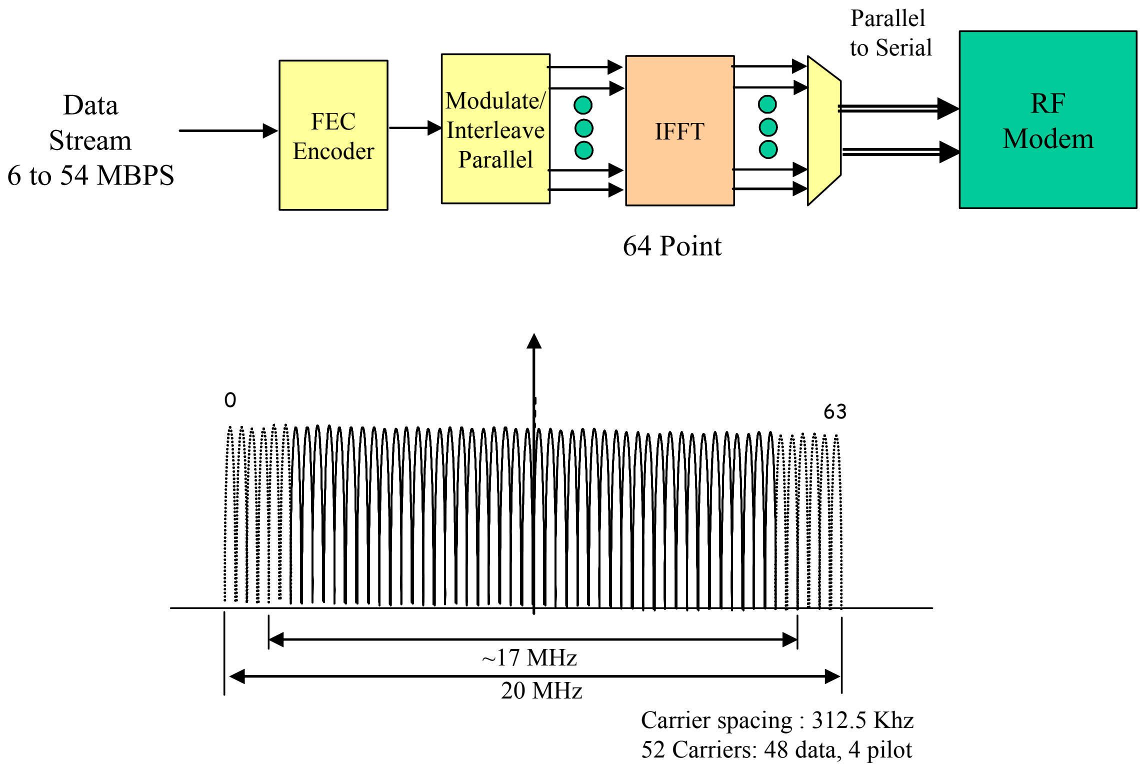

Figure 1

What Is Channel Bonding?

WLAN channel bonding, as the name implies, bonds or combines two of the standard 802.11 OFDM channels together and uses them simultaneously to achieve greater data throughput. The channels that are bound together are adjacent or contiguous to one another (Figure 2). Channel bonding also doubles the size of the FFT, allowing the FFT to multiplex twice the amount of data. To maintain multi-path performance of the 802.11 waveform, channel bonding only has been applied to OFDM 802.11 modulations and is not being applied to the older CCK and direct sequence 802.11b modulation formats.

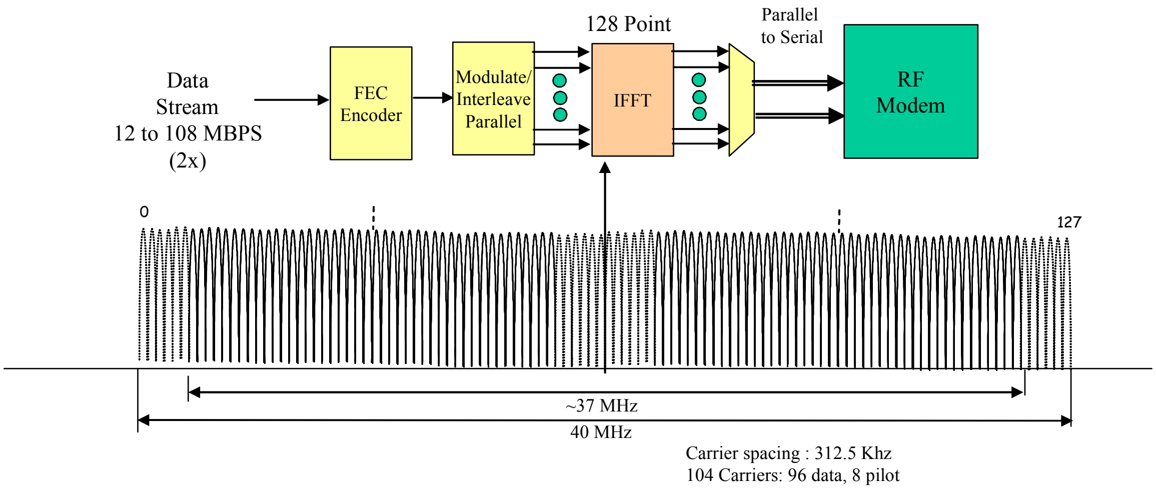

Figure 2

Channel bonding follows all of the same processes of the original 802.11a/g standard, except that a 128-point FFT is implemented. In addition, the sampling and clock rates must be doubled to maintain the standard 802.11 symbol period over a 40 MHz channel (37 MHz effective signal bandwidth) that is double the bandwidth of a standard 802.11 channel.

This type of channel bonding is spectrum compatible, which means that 802.11a/g OFDM demodulator hardware can be re-used, but the de-multiplexer and demodulator must be run at twice their normal speed since they are receiving output from a 128-point FFT rather than a 64-point FFT. With a channel bonding WLAN implementation, the IFFT's output is treated as two separate blocks of 64 points each.

Variations to this channel bonding scheme have been suggested, including the elimination of the 1.5 MHz guard band tones at the center of the 40 MHz bonded channel to increase throughput even more. However, this greatly impedes the ability of standard 802.11 modems to recognize and properly perform clear channel assessments for the 802.11 DCF and EDCF MAC protocols.

Implementation Problems with Channel Bonding

In real-world implementation of WLANs, channel bonding greatly complicates the RF filtering that takes place in 802.11 RF receivers.

To properly accomplish RF filtering in a WLAN using channel bonding, the receivers should deploy two sets of filters, one for the standard 20 MHz 802.11 channel to guard against adjacent channel interference from nearby 802.11 WLANs and another for the 40 MHz bonded channel. For example, a super-heterodyne radio architecture would provide two switched filters at the super-het intermediate frequencies (IF) and an anti-aliasing filter (un-switched) as a baseband.

Unfortunately, the high cost of such an architecture dictates the deployment of only the 40 MHz of filtering for anti-aliasing over the bonded channel and eliminates the filtering for 802.11

adjacent channel interference. Eliminating these filters increases the interference that a WLAN's AP and client receivers must contend with (Figure 3).

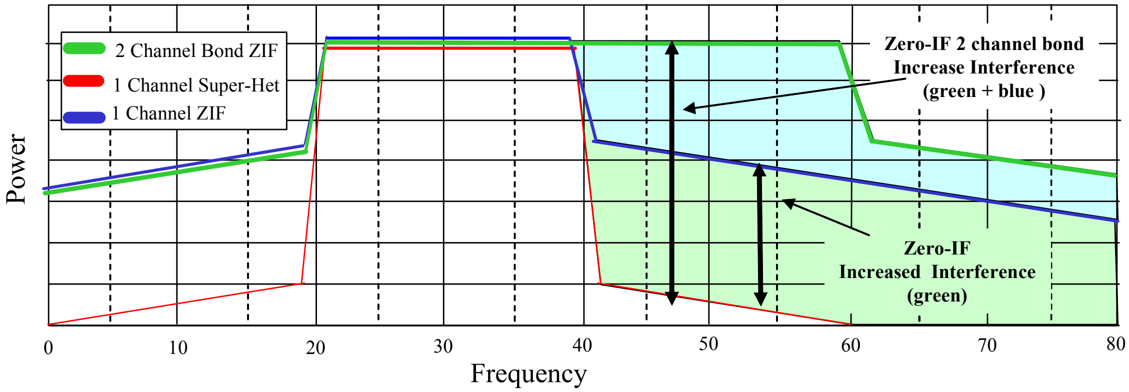

Figure 3

In Figure 3, the receiver channel response is illustrated for a bonded 40 MHz channel with a direct conversion (DC) modem (green line), for a single 20 MHz channel with a super-het receiver (red line) and for a single 20 MHz channel with a DC modem. The graph shows that channel bonding severely limits the performance of the 40 MHz channel in terms of adjacent channel interference (ACI). Comparing the performance of the bonded channel to that of the 20 MHz channel with a super-het receiver, the bonded channel's receiver experiences a nearly 60 dB loss of ACI protection from the nearest adjacent 20 MHz of spectrum and a 40 dB loss of ACI protection in the 20 to 40 MHz range at the edge of the band.

This loss of this protection from ACI is a significant compromise for a modem in an unlicensed band of the spectrum. Obviously, it can adversely affect the modem's performance by making it more susceptible to interference, but it also has a deleterious effect on frequency re-use in WLAN deployments involving multiple 802.11 access points.

Some might assert that the FFT in a WLAN with channel bonding provides filtering naturally and an analog filter is not needed. Figures 4 and 5 show that this is not true. Figure 4 illustrates how a standard 20 MHz 802.11 channel suffers from the leakage of ACI signal energy across the received signal because of the composite SinX/X response of the sum of the receiver FFT bins.

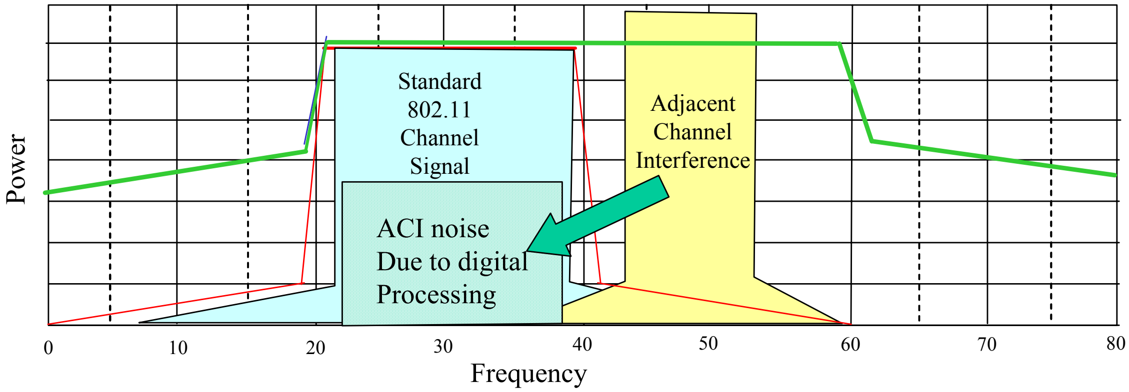

Figure 4

Figure 5 shows the FFT bin leakage in detail. Combined with fixed point precision limitations, it is apparent that the bin leakage results in a maximum adjacent channel rejection floor for the receiver of 25 to 30 dB. This is particularly significant in denser deployments of WLAN technology, such as a heavily populated urban setting or an apartment building.

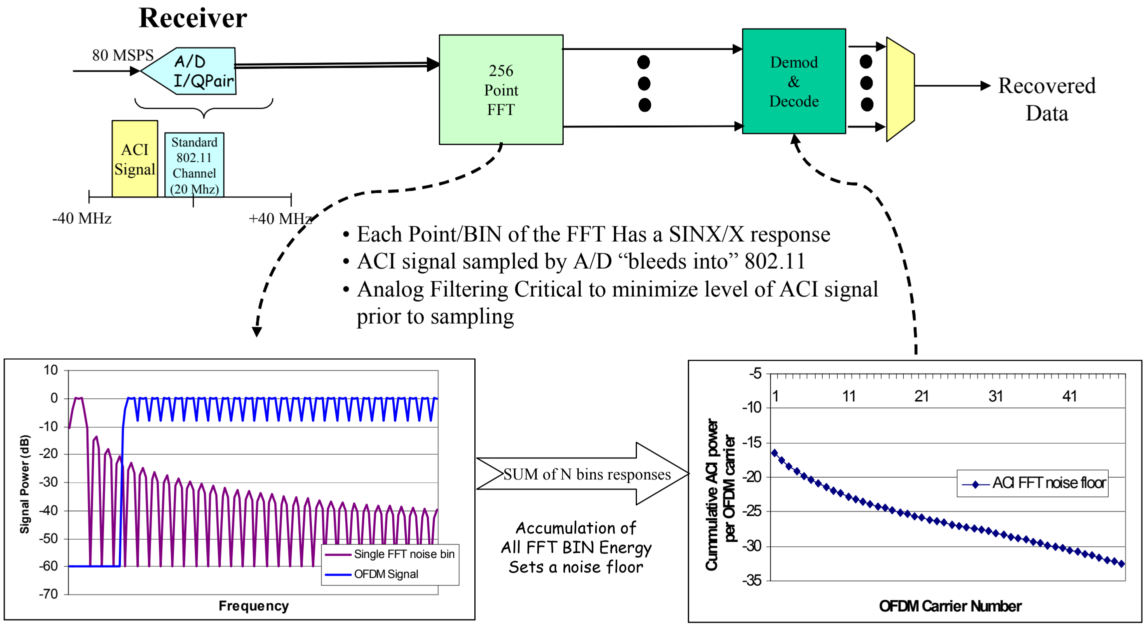

Figure 5

An environment much like the one illustrated in Figure 6 is quite conceivable. This drawing shows one AP (AP1) with a channel bonding receiver that combines two standard 802.11 channels. Another AP (AP2) 100 feet away is using an adjacent 802.11 channel.

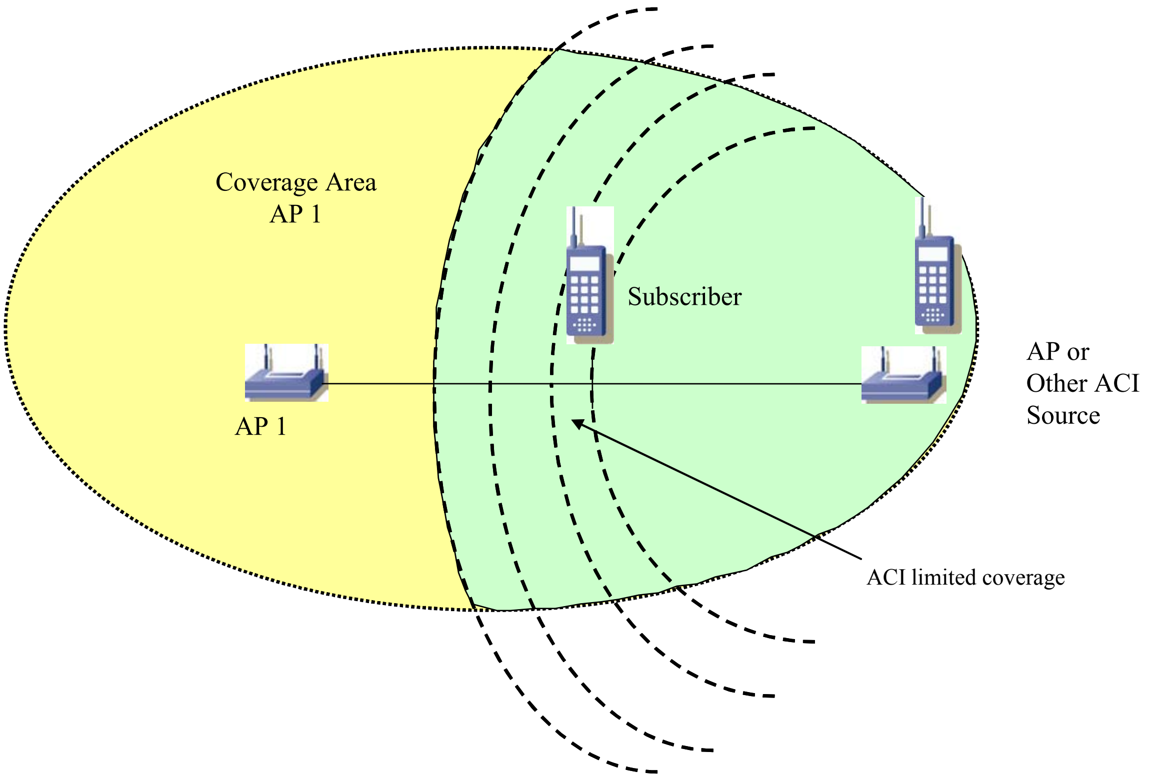

Figure 6

Given the type of environment portrayed in Figure 6, the data throughput figures and effective range for three different types of receivers are shown in Table 1 below. The performance of a super-het receiver for one 20 MHz non-bonded channel is compared with a direct conversion receiver for one 20 MHz non-bonded channel and a direct conversion receiver for a 40 MHz bonded channel. The effects of adjacent channel interference on performance as the distance from the access point increases are apparent.

	Distance from Main Access Point (meters): standard (20 MHz) super-het receiver									
	3.5	7	10.5	14	17.5	21	24.5	28	31.5	34
RX power AP1 (dbm)	-36.3	-45.4	-50.6	-54.4	-57.3	-59.7	-61.7	-63.4	-64.9	-65.9
RX power AP2 (dbm)	-64.9	-63.4	-61.7	-59.7	-57.3	-54.4	-50.6	-45.4	-36.3	-20.0
SIR (db)	113.6	103.1	96.0	90.3	85.0	79.7	74.0	66.9	56.4	39.1
Modulation support	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM

	Distance from Main Access Point (meters): standard (20 Mhz) direct conversion receiver									
	3.5	7	10.5	14	17.5	21	24.5	28	31.5	34
RX power AP1 (dbm)	-36.3	-45.4	-50.6	-54.4	-57.3	-59.7	-61.7	-63.4	-64.9	-65.9
RX power AP2 (dbm)	-64.9	-63.4	-61.7	-59.7	-57.3	-54.4	-50.6	-45.4	-36.3	-20.0
SIR (db)	78.6	68.1	61.0	55.3	50.0	44.7	39.0	31.9	21.4	4.1
Modulation support	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	54 OFDM	36 OFDM	NO LINK

	Distance from Main Access Point (meters): bonded (40 MHz) direct conversion receiver									
	3.5	7	10.5	14	17.5	21	24.5	28	31.5	34
RX power AP1 (dbm)	-36.3	-45.4	-50.6	-54.4	-57.3	-59.7	-61.7	-63.4	-64.9	-65.9
RX power AP2 (dbm)	-64.9	-63.4	-61.7	-59.7	-57.3	-54.4	-50.6	-45.4	-36.3	-20.0
SIR (db)	53.6	43.1	36.0	30.3	25.0	19.7	14.0	6.9	-3.6	-20.9
Modulation support	54 OFDM	54 OFDM	54 OFDM	54 OFDM	36 OFDM	PBCC 22	PBCC 22	NO LINK	NO LINK	NO LINK

Table 1

The performance figures in Table 1 point out the significant effect ACI has on the 40 MHz bonded channel receiver. The following performance limitations in this channel can be attributed to ACI:

- Throughput reductions in 60 percent of the range
- Loss of link in 28 percent of the range

Unlike the bonded channel receiver, the super-het receiver for a single 20 MHz channel is unaffected by the ACI generated by AP2. A super-het receiver is able to maintain peak data throughput rates (54 Mbps) across the entire range of the access point.

As the performance figures in Table 1 demonstrate, the coverage range of a channel bonding implementation is limited. Part of the reason for this is the fact that for a given RF transmit power amplifier (PA) design, the channel bonded waveform will lose a minimum of approximately three dB of power per OFDM bin. This loss comes about because the transmitter's PA power output must be spread over a bonded channel that is twice the standard 802.11 bandwidth. As a result, each transmitted OFDM frequency bin will have half the transmitted energy.

Typical signal propagation losses are calculated as R^3 (radius to the third power). So, a three dB loss in the signal would cause a 21 percent reduction in the effective radius of an AP's cell and a 37 percent loss of coverage area for the cell. In addition, the signal power losses will probably be greater than three dB in many cases because of PA back-off and receiver implementation losses.

Frequency Reuse in Large WLAN Deployments

More and more these days, multiple 802.11 WLAN access points are being deployed in one network to form a wide area WLAN. This is already happening and will become more prevalent as 802.11 technology is implemented in campus and enterprise settings, in larger hotspots like airports and hotels, and in dense urban developments like apartment or condominium complexes. When wide area deployments of WLANs take place, network designers essentially are designing a micro-cellular RF network. With this type of network, and especially in an unlicensed band of the spectrum, RF planning and frequency reuse become critically important for the success of the deployment.

RF network planning begins with the frequencies that are available. For 802.11a/b/g radios, that means the following:

- 5.1 to 5.3 GHz band with eight frequencies
  - NOTE: This band will expand to include 5.4 to 5.825 GHz
- 2.4 GHz band with three frequencies

For access points with simple omni-directional antennas, Figure 7 below illustrates the seven frequency and three frequency repeat patterns with a frequency reuse of one. The seven frequency plan can be used for the 802.11a mode in the 5.x GHz range while the three frequency plan can be used for 802.11b/g networks.

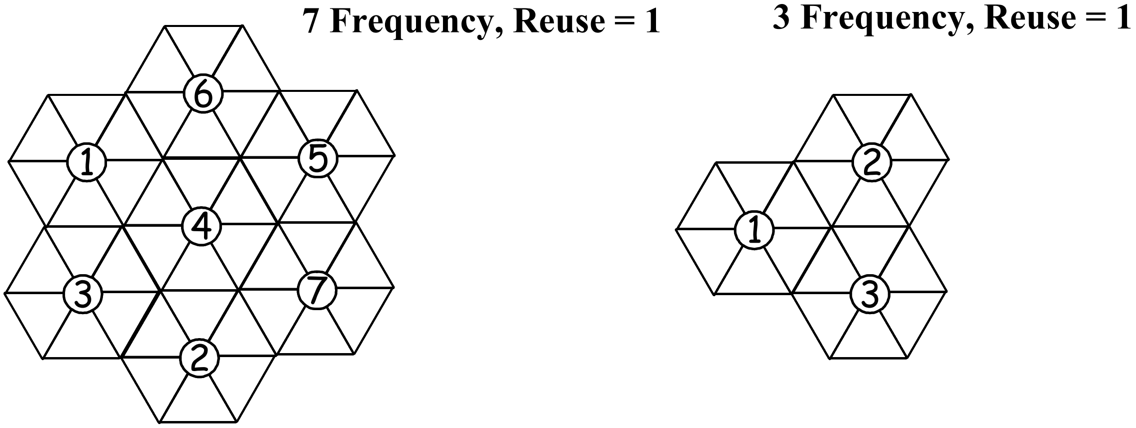
7 Frequency, Reuse = 1

3 Frequency, Reuse = 1

Figure 7

In the deployments shown above, the cell reuse distance, R_u, can be defined as the following:

- C = 7 (7 frequency): $R_u = R_{cell}$*sqrt(3C) = 4.48* R_{cell}
- C = 3 (3 frequency): $R_u = R_{cell}$ *sqrt(3C) = 3.00* R_{cell}

WLAN network planners must assume that the signals from an access point will propagate beyond the cell's radius. Therefore, RF propagation loss will not be free space (R^2) but rather R^3 to R^4. This would result in interference reduction between the cells of at least the following:

- C = 7: 19.5 dB to 26.1 dB ***(allows 36 to 54 Mbps OFDM)***
- C = 3: 14.3 dB to 19.1 dB ***(allows PBCC 22 Mbps to 36 Mbps OFDM)***

This means that in WLAN deployments using seven frequencies, OFDM data rates from 36 to 54 Mbps could be supported. Where three frequencies are implemented, PBCC and OFDM data rates from 22 to 36 Mbps would be supported. This shows that network planners of large WLANs could deploy 802.11a/b/g access points with omni-directional antennas and each AP could automatically select its frequency plan for optimum network performance up to 54 Mbps.

The Effect of Channel Bonding on Frequency Reuse

In WLANs where channel bonding is deployed, at most two RF channels with optimal frequency spacing will be possible in the 2.4 GHz ISM band. The reuse pattern for these two channels is shown in Figure 8.

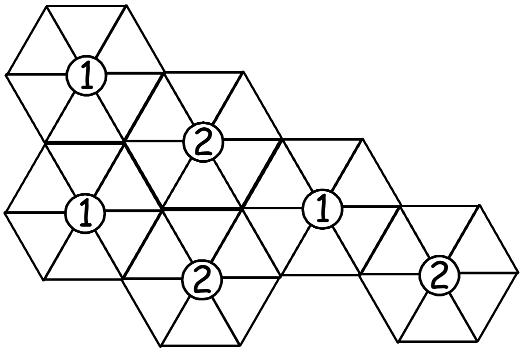

Figure 8

This reuse pattern would yield the following RF propagation losses among access points:

- Reused distance: C = 2 (2 frequencies): $R_u = R_{cell}$ *sqrt(3C) = 2.44* R_{cell}
- Interference reduction : C = 2: 11.6 db to 15 dB ***(allows 11 Mbps CCK)***

This analysis indicates that channel bonding reduces the throughput of a WLAN in a high-density deployment by a factor of two, effectively negating the reason why a network planner would consider implementing channel bonding in the first place.

Based on analyses such as these, it is easy to understand why European and Asian spectrum regulatory agencies have banned 802.11 channel bonding from their countries. If channel bonding were deployed, it would limit or eliminate 802.11 frequency reuse in most deployment scenarios and it would prove to be an ineffective way of increasing WLAN data throughput.

Viable Alternatives to Channel Bonding

Texas Instruments has a long history of providing high-performance 802.11 WLAN solutions that are bandwidth compatible with the 802.11 standard and completely interoperable with other 802.11 standard products and systems. The first example of this was TI's 802.11 "+" technology for 802.11b networks. This made use of PBCC spectrum-compatible modulation at the physical layer (PHY) to double 802.11b data rates to 22 Mbps. Soon thereafter, media access (MAC) improvements from TI resulted in another 30 percent improvement in data throughput on 802.11b WLANs. Figure 9 illustrates these TI innovations.

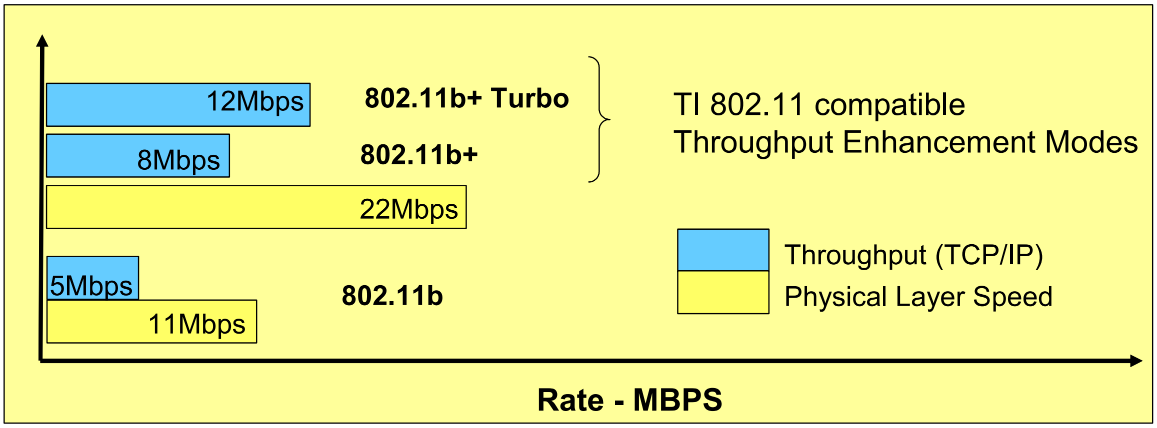

Figure 9

For 802.11a/g technology, TI's track record for improving WLAN performance is no less impressive. Systems based on the TNETW1130 and TNETW1230 single-chip MAC and baseband processors will be capable of throughput speeds 30 percent higher than comparable devices. Later, TI's "+" technology will be integrated into the TNETW1130 and TNETW1230 devices for a spectrum-compatible solution capable of speeds of 100 Mbps. When additional enhancements are made to the MAC technology of the TNETW1130, the throughput of this device will be three times that of standard 802.11g products and the TNETW1130 will provide 50 percent more throughput than a channel bonding solution. Figure 10 shows the results of these innovations.

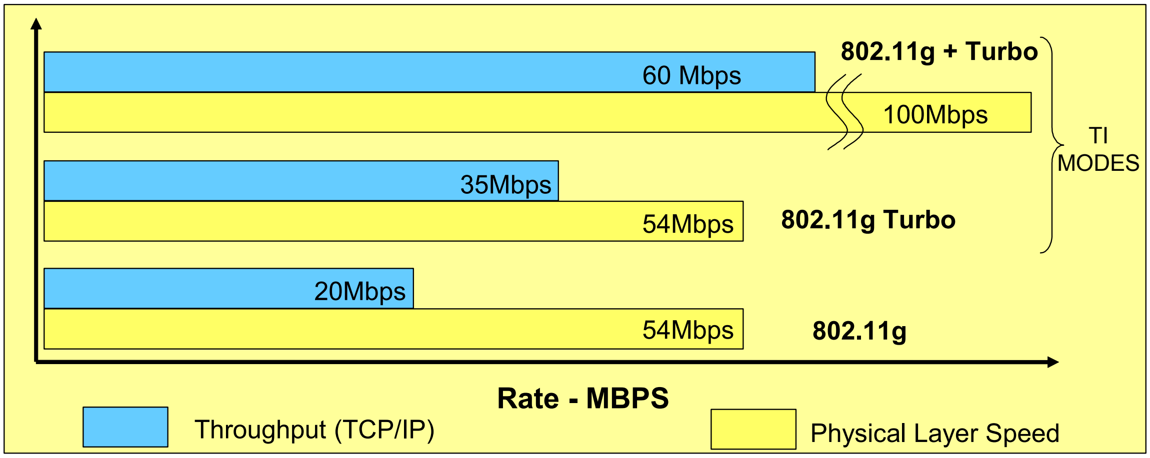

Figure 10

Conclusion

Because it makes use of the FFT engine currently implemented in 802.11 OFDM designs, channel bonding may appear on the surface to be an effective way of quickly improving the data rate and throughput of Wi-Fi-based wireless local area networks. After all, by implementing channel bonding, 802.11 chipset makers can increase the throughput of their chipsets with a bare minimum of development effort. Unfortunately, the benefits of channel bonding bring with them numerous limitations that are not associated with other alternatives for increasing throughput. The disadvantages of channel bonding include the following:

- Sensitivity to Adjacent Channel Interference is far greater with channel bonding implementations. And the ACI problem will only increase as the adoption of WLAN technology gains momentum. Problems with ACI could have severe ramifications on user satisfaction with Wi-Fi and this, in turn, could stunt market growth.
- Compatibility between systems with channel bonding and legacy 802.11a/b/g products will be an issue as strong adjacent channel energy from channel bonding systems will interfere with standard 802.11 packet detection schemes. As a result, throughput will be greatly reduced from increased packet collisions on WLANs.
- Frequency reuse becomes an issue with channel bonding systems, especially in the 2.4 GHz range of 802.11b/g WLANs. Poor frequency reuse will dampen the deployment of WLANs in many urban or high-density areas such as apartments, condominiums, campus or enterprise settings, and large hotspots like airports and hotels.
- Channel bonding is not allowed in European and Asian markets.

Because of these limitations, TI is increasing the throughput of WLANs with solutions that are spectrum compatible with the 802.11 standard and backwards interoperable with existing and legacy 802.11a/b/g/h modems. TI's current solutions, such as its "+" WLAN technology, already have provided three times the performance of standard 802.11a/g modems. In the future, TI will develop solutions that are just as innovative and include new advancements, such as MIMO technology, which could provide up to six times the throughput performance of standard 802.11a/g solutions.

For more information, visit the Texas Instruments Web site:
www.ti.com/wlan