Volume 07 Issue 03 Published, August 19, 2003 ISSN 1535-864X

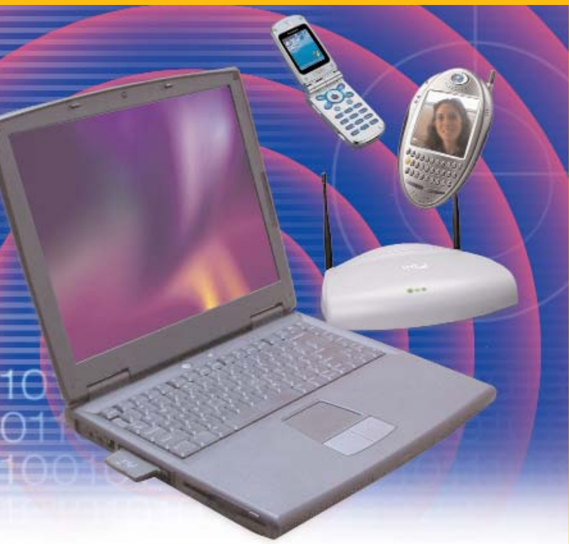

Intel® Technology Journal

Wireless Technologies

High-Throughput Wireless LAN Air Interface

High-Throughput Wireless LAN Air Interface

Boyd Bangerter, Corporate Technology Group, Intel Corporation
Eric Jacobsen, Corporate Technology Group, Intel Corporation
Minnie Ho, Corporate Technology Group, Intel Corporation
Adrian Stephens, Corporate Technology Group, Intel Corporation
Alexander Maltsev, Intel Communications Group, Intel Corporation
Alexey Rubtsov, Intel Communications Group, Intel Corporation
Ali Sadri, Intel Communications Group, Intel Corporation

Index words: High-Throughput WLAN, 802.11n, OFDM, MIMO, Adaptive Bit Loading

ABSTRACT

This paper highlights the major candidate technologies for unlicensed band Wireless Local Area Network (WLAN) air interface performance improvements and discusses the advantages and disadvantages associated with each technology. We also explore the improvements in digital and radio frequency process technology that will contribute to the commercialization of high-throughput wireless LAN systems.

INTRODUCTION

Wireless Local Area Network (WLAN) technology is a fast-growing segment in the computing and communications equipment market, and the technology is evolving at a rapid pace. This evolution is happening as a result of coordinated standards efforts as well as innovation by leading-edge companies. It is taking place on many fronts: improvements in antenna systems, radio frequency (RF) components, modulation schemes, medium access control (MAC) mechanisms, and security mechanisms. Today, the air interface of standards-based WLAN equipment supports data rates as high as 54 Mbps. Designers of future WLAN equipment are attempting to specify an air interface that will achieve data rates in excess of 250 Mbps.

A variety of challenging tasks face designers of future air interfaces for WLAN systems. First, in order to serve a worldwide market, high-performance modes must be designed for more frequency bands than previous generations of WLAN equipment. Second, future designs must provide robust methods for co-existence with legacy equipment, and in many cases must also provide methods for backwards compatibility. Third, designers must select from a wide variety of performance-improving candidate technologies. The stand-alone performance characteristics of many candidate technologies are well understood; however, it is a challenge to understand how these technologies interact with each other as part of a complete system. Fourth, many of these candidate technologies require computationally intensive signal-processing algorithms and more complex radio frequency (RF) systems. Designers are faced with a large array of tradeoffs in order to find the optimum mix of performance, price, and power consumption.

THE CHALLENGES OF HIGH-SPEED DESIGN IN UNLICENSED BAND WIRELESS CHANNELS

Wireless channels present some challenging engineering problems that are not found in wireline systems. Most Wireless Local Area Network (WLAN) systems use omnidirectional antennas that provide good coverage but don't concentrate the transmitted power to the intended user. This also means that because signal energy is scattered and reflected from objects in the environment, components of the signal arriving at the receiver are spread out over a longer period of time than is desirable. Since the frequency spectrum used by WLAN devices is unlicensed, other devices may be attempting to use the same channel resources and thus create interference. The challenge is then to provide a high-performance, reliable data link that can operate with restricted receiver power levels, severe channel fading due to multipath reflections, and interfering energy from other devices.

As with many technologies, it is useful to compare achieved performance with theoretical limits. In this case the theoretical limit is Shannon's capacity, which is usually expressed as $C = B\log_2(1+SNR)$

where B is the occupied bandwidth and SNR is the signal-to-noise ratio at the receiver. C is then the theoretical maximum rate in bits-per-second that can be achieved in that channel with no transmission reception errors.

This calculation has some important limitations, however. It is developed only for a single use of a memoryless channel impaired only by Additive White Gaussian Noise (AWGN). The result does not hold for receivers experiencing multipath fading or interference from other devices, which are both impairments that are often greater than the noise contribution in many wireless systems.

HIGH-THROUGHPUT WIRELESS LAN CANDIDATE TECHNOLOGIES

Fortunately, there are technologies that can take advantage of multipath in wireless channels. One such set of technologies involves the use of multiple antennas for both transmitting and receiving. Another such set of technologies involves wideband adaptive Orthogonal Frequency Division Multiplexing (OFDM). We describe both sets of technologies in the following sections.

MIMO Systems for High Throughput

Systems employing multiple antennas for both transmitting and receiving are often called multiple-input multiple-output (MIMO) systems. The primary advantage of MIMO systems over single-input single-output (SISO) systems is that in the presence of a "rich" multipath, a MIMO system with M antennas at the transmitter and M antennas at the receiver provides M times the peak throughput of an SISO system without increasing the frequency bandwidth. This is performed by dividing the channel into multiple "spatial channels" through which independent data streams can be transmitted. This technique is known as "spatial multiplexing."

Additionally, a MIMO system can be used to increase diversity gain, where the slope of the Bit Error Rate (BER) curve when plotted against the Signal-to-Noise Ratio (SNR) can be dramatically increased, up to a maximum of MxMxL over a SISO system, where L is the number of effective, independent multipath components, or taps. This increase in diversity gain translates to an increase in range at a given BER or an increase in data rate at a given range up to the maximum rate of the original SISO link. Other possible gains, such as array gain or coding gain, are not discussed here.

Since our focus is on MIMO systems for high throughput, we first consider spatial multiplexing techniques and then consider diversity techniques to improve performance. This is analogous to a wideband, adaptive OFDM system, where the wide bandwidths are used to increase the peak throughput of the original link, and adaptive bit loading (a diversity technique) is added to improve performance. We categorize diversity techniques for MIMO systems into the following two categories. In the first category, the transmitter has no knowledge of the channel and uses a coding technique to achieve diversity over an ensemble average of channel realizations. In the second category, the transmitter has partial or full knowledge of the channel and uses this knowledge to increase diversity gain. In both categories, we presume the receiver has knowledge of the channel, derived from training sequences, in order to separate data from the multiple spatial channels.

The first category is conceptually the simplest. A transmitter simply encodes the bits over space and over frequency and transmits these bits over multiple spatial channels. The receiver then separates the symbols from the multiple spatial channels and decodes the bits. An example of an architecture belonging to the first category is depicted in Figure 1.

The second category of diversity techniques is more complex because in order for a MIMO transmitter to gain full or partial knowledge of the channel, one of the following two possibilities must occur: (1) knowledge of the channel at the receiver is simply turned around and used at the transmitter, and no feedback is necessary or (2) knowledge of the channel at the receiver is fed back to the transmitter using a handshaking protocol. An example of (1) is depicted in Figure 2, and an example of (2) includes adaptive bit-loading techniques, described in later sections of this paper.

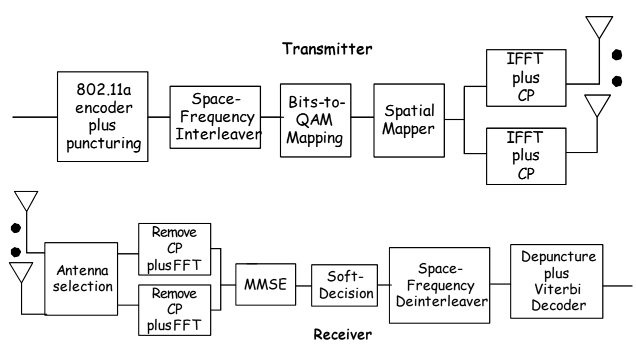

Figure 1: MIMO system without knowledge of the channel at the transmitter

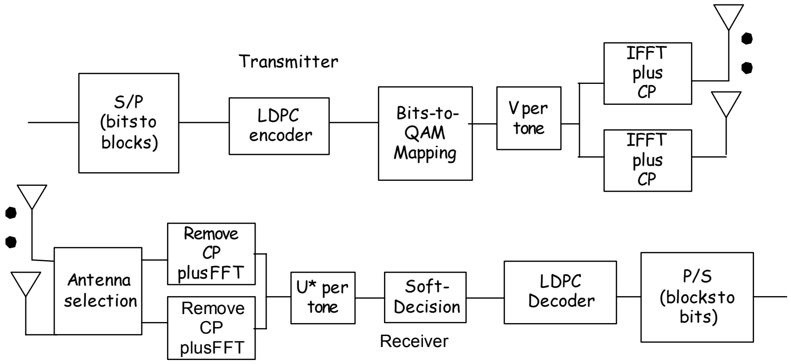

Figure 2: MIMO system with knowledge of the channel at the transmitter, but without feedback

MIMO Simulation Results

Figure 3 depicts the mean capacity vs. SNR (Shannon limits) of MIMO systems with different numbers of transmit and receive antenna elements. This simulation uses a random Rayleigh channel model. From this figure, we see that the capacity of the MIMO system with equal numbers of transmit and receive antenna elements increases almost *directly in proportion to the number of transceiver antenna elements* in comparison with usual SISO communication systems *without increasing the total transmitted power or frequency bandwidth*.

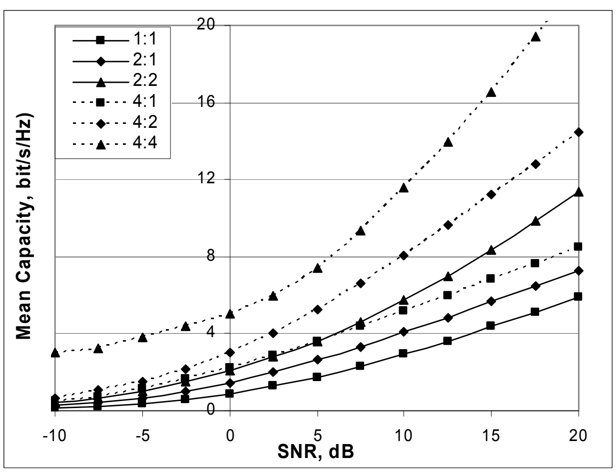

Figure 3: Mean capacity vs. SNR for MIMO systems

In Figure 4, the system depicted in Figure 1 was simulated with M_t transmit antennas and M_r receive antennas. Of these M_r receive antennas, M_c of the "best" antennas were chosen, where a Shannon capacity metric was used to determine the best antennas. This is denoted in the legend using the following notation: (M_r) M_c x M_t. We can see from this plot that we can achieve twice the throughput of a switched-diversity SISO (2) 1x1 system with a (4) 2x2 system, with a loss in SNR of 1 dB.

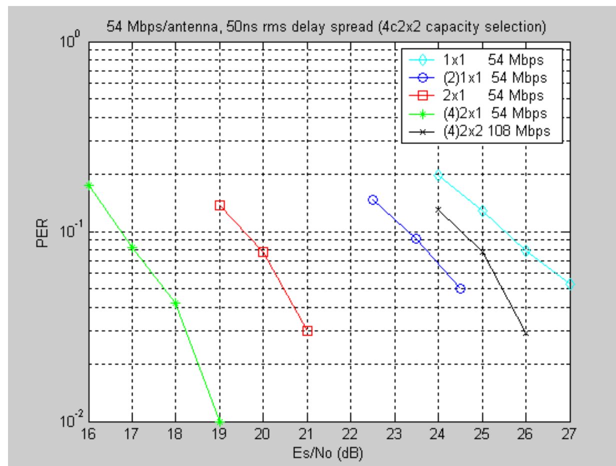

Figure 4: Simulation results for MIMO system depicted in Figure 1

In Figure 5, two comparisons are shown. The first comparison is between the techniques with and without channel knowledge at the transmitter: open loop minimum mean squared error (MMSE) scheme and closed loop singular value decomposition (SVD) scheme. The second comparison is between two forward error correction codes: (1) the 802.11a convolutional code and (2) a low-density parity check (LDPC) code described in a later section. Receive diversity selection is applied in all curves, and the number of antennas available to select are specified in the parentheses in the legend. In the SVD scheme, the channel matrix H is decomposed into the multiple spatial channels using a Singular Value Decomposition (H=UDV*). An LDPC code is applied to each of the spatial channels. We can see that the "SVD+LDPC+SVD" packet-error-rate (PER) curves gain about 3.56 dB in SNR over the "MMSE+CC" PER curve at PER 10^{-1}. The curves also demonstrate that the substitution of convolutional code with LDPC code by itself delivers 2 and 1.5 dB for open and closed loops, respectively. We also observe that we can achieve twice the throughput of a switched-diversity SISO (802.11a) system at a gain of 1 dB in SNR. It should be noted that the LDPC block size is of 4 μs duration, which is similar in latency to the convolutional code. In the figure, the upper solid lines are for PER and the lower dash lines are for BER performance.

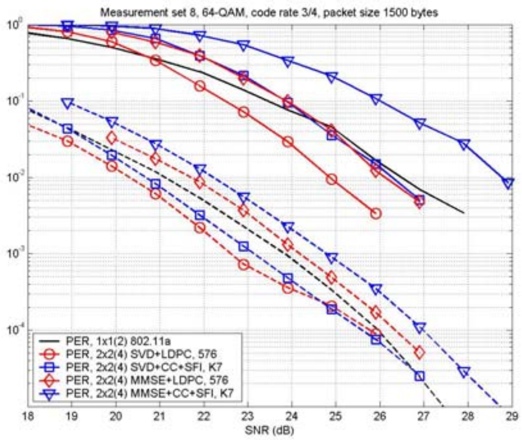

Figure 5: Simulation results for the MIMO system depicted in Figure 2

Wideband Adaptive OFDM

Another technology that may extend the capabilities of WLAN systems is adaptive OFDM. OFDM is the modulation currently used by 802.11a and 802.11g systems. OFDM works by dividing up a wideband channel into a larger number of sub-channels. By placing a subcarrier in each sub-channel, each subcarrier may be modulated separately depending on the SNR characteristics in that particular narrow portion of the band. As the channel varies over time, further adaptations can be made on each subcarrier in order to continually optimize the data-carrying capacity of the channel.

One possible approach to achieving higher WLAN data rates is to apply adaptive bit and power-loading techniques while simultaneously increasing the frequency bandwidth of signals from 20 MHz (as in 802.11a) to multiples of 20 MHz up to 80 MHz. Such an approach will result in capacity gains from both the adaptive techniques within the band as well as the raw linear increase in data rate expected with the increase in bandwidth.

Fast Link Adaptation With Adaptive Bit and Power Loading

Much research has gone into achieving the potential channel capacity using adaptive modulation, subcarrier power allocation, and coding techniques for OFDM systems. The underlying idea behind these methods is to exploit "good" and "bad" subcarriers in such a way as to maximize the data-carrying ability of the channel. This approach is often referred to as the "water-filling" approach. In the water-filling approach, more power and higher order modulation can be put onto subcarriers with larger SNRs, and lower SNR subcarriers will receive less power and lower order modulation up to a certain threshold, after which the subcarrier should simply be turned off, or "punctured."

A few different types of adaptive bit and power-loading algorithms have been proposed and analyzed. Some algorithms, such as those proposed by Chow *et al* [1] and Fischer with Huber [2], are suboptimal in the sense of the maximization of the SNR margin for a target data rate. Other versions of a finite granularity loading algorithm were proposed by Cioffi *et al*. [3] for transmission over asymmetric digital subscriber lines (ADSL). This algorithm maximizes the data rate at a given SNR margin that guarantees the BER is less than a certain target level. This algorithm turns off some of the subcarriers and divides the remaining (active) subcarriers into fixed subsets in accordance with either water-filling or uniform power loading. In each subset, subcarriers are given the same combination of "modulation+encoding," and then rescaling of subcarrier powers is performed. Hanzo *et al.* [4] made a significant contribution to refining adaptive bit and power-loading algorithms for duplex OFDM wireless links.

We propose and investigate different types of pragmatic adaptive bit and power-loading schemes that guarantee a target PER for wireless OFDM systems. They are suitable for duplex time division multiplexing (TDM) communication between two wireless stations, when they are acting over a slowly varying frequency-selective channel. Suppose that a first OFDM station (STA1) transmits data to a second OFDM station (STA2). During an initial handshake period (RTS/CTS exchange), STA2 measures channel characteristics observed when receiving the RTS frame and calculates an appropriate bit and power-loading allocation. STA2 transfers the results of this calculation to STA1 in the response (CTS) packet. STA1 will subsequently modulate subcarriers in accordance with the received modulation parameters. STA2, knowing these parameters in advance, will be able to receive and process the OFDM symbols. This closed-loop adaptation mechanism provides fast link adaptation that can follow time-varying channel characteristics.

We have developed a number of pragmatic finite-granularity adaptive bit and power-loading schemes that differ in throughput performance, complexity, and the amount of feedback service information. Three schemes are described below.

The Adaptive Bit and Power-Loading (ABPL) Scheme

This scheme exploits the water-filling principle for initial power allocation and determination of "no transmission" (non-active) subcarriers. The number of subcarriers in each subset with a uniform modulation plus encoding type is recursively determined. Then SNR pre-equalization is performed within each subset in order to guarantee a target fixed BER for each subcarrier. During this procedure, surplus power removed from higher-order modulation subsets is added to lower-order modulation subsets and is also used for turning on subcarriers that were initially turned off. At all stages this scheme is constrained by FCC requirements on peak power spectral density (PSD). It is the most computationally complicated scheme and has the best throughput performance. At the same time, however, it requires the largest amount of feedback service information, and it is very sensitive to channel state estimation errors.

The Adaptive Bit-Loading (ABL) Scheme

The ABL uses the same transmit power for all active subcarriers. It guarantees a BER less than some target threshold for each subcarrier and uses a "no transmission" mode for very bad subcarriers. The ABL scheme has some throughput degradation compared to the ABPL scheme, but it needs less feedback service information. Simulations have shown that the ABL scheme is more robust to impairments such as inaccurate channel state estimation and to interference and time variation of the wireless channel caused by Doppler spread.

To decrease the amount of feedback required, we investigated modifications of these schemes that allocate subcarriers the same parameters in groups of size 2, 4, or 6. In each group, subcarriers have the same power and modulation+coding type that are determined on the basis of effective noise power per subcarrier in the group.

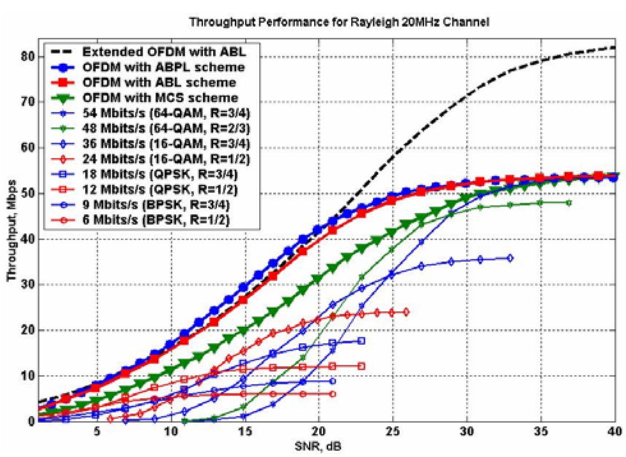

Figure 6: Simulation results for a static random Rayleigh 20 MHz channel model

Simulations show that a 2-subcarrier grouping has negligible throughput performance degradation (less than 0.3 dB in equivalent SNR); a 4-subcarrier grouping leads to performance degradation of about 1 – 1.5 dB, and a 6-subcarrier grouping gives substantial degradation in throughput performance (about 2 – 3 dB). A 2-subcarrier grouping halves the amount of feedback information required and is recommended for a high-throughput WLAN operating in a typical indoor environment.

Optimized Modulation and Coding Selection (MCS)

Another fast adaptation scheme is when all subcarriers are given the same modulation and power, based on channel state feedback. The optimal selection of modulation and coding type is made on the basis of the "momentary link performance" i.e., short-term SNR or PER estimates, and also a priori knowledge of momentary PER vs. momentary SNR performance curves [5]. We have investigated one algorithm for a modulation and coding selection (MCS) scheme. It shows a throughput degradation compared to ABPL or ABL. However, it requires a minimum amount of channel state feedback and gives a 1.5 – 2.5 dB throughput gain over a conventional long-term MCS strategy (see Figure 6).

Extended Bandwidth OFDM Simulation Results and Backwards Compatibility

Throughput and PER performances for the adaptive loading schemes were investigated for a 20 MHz OFDM system for different channel models, including models with Doppler spread. Results of the simulations for the static random Rayleigh channel model [6] are shown in Figure 2. Four modulations (BPSK, QPSK, 16-QAM, and 64-QAM) and three convolutional code rates (802.11a convolutional code with R=1/2, 2/3, 3/4) were considered. A 3-bit soft decision Viterbi decoder was built into the OFDM system model. The length of the data frame was 1000 bytes, and the target PER was 1%. Ideal carrier frequency estimation, timing synchronization, and channel state estimation were assumed.

It can be seen from Figure 6 that the ABPL scheme has about a 6 – 7 dB throughput improvement over standard OFDM. The ABL scheme gives about a 5 – 6 dB throughput gain over standard OFDM. The MCS scheme has a 1.5 – 2.5 dB throughput gain compared to a standard OFDM.

Extending Modulations, Coding Rates, and Frequency Bandwidth for HT WLAN

It is clear from Figure 2 that for a frequency-selective Rayleigh channel using any adaptive bit and power-loading scheme, a significant performance increase over conventional modulation is achieved, but there is no increase in the maximum achievable data rate (54 Mbps). We now consider two ways to enhance the maximum data rate in the context of single-antenna systems.

The first way is to use higher order modulations and/or higher coding rates to improve spectral efficiency. We have investigated adaptive loading schemes in conjunction with an extended set of modulation types (including 256-QAM) and coding rates (R = 7/8). This can achieve a data rate of up to 84 Mbps within a 20 MHz channel. The throughput results for the ABL scheme using the extended modulation and code-rate set is shown in Figure 6 by the black dashed curve.

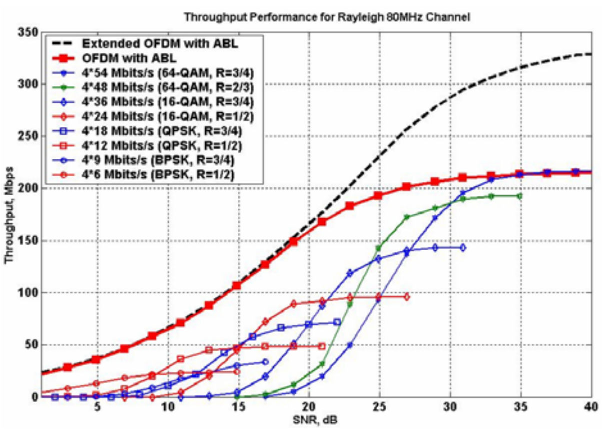

Figure 7: Simulation results for static random Rayleigh 80 MHz channel model

The second way is to use increased channel widths. We investigated using up to four 20 MHz channels in one U-NII band in conjunction with the adaptive loading schemes. For example, simulation results for an OFDM system with ABL, operating in a static random Rayleigh 80 MHz channel model [5], are shown in Figure 7. It can be seen that the ABL scheme gives about a 6 dB throughput gain (see bold red curve with squares) over conventional OFDM. In an extended bandwidth mode, the OFDM system throughput can achieve 300 Mbps (see black dashed curve).

Throughput Versus Range Performance for Extended Bandwidth OFDM Systems With an ABL Scheme

Extended bandwidth OFDM leads to decreased transmit power spectral density because of FCC limits on the total transmit power of wireless devices. Increasing frequency bandwidth by a factor of four results in a decreased SNR per subcarrier by the same factor. This causes range performance degradation in extended bandwidth OFDM systems.

At the same time, using these adaptive bit and power-loading schemes in extended bandwidth systems leads to economic use of available transmit power because only good subcarriers are allocated power, and bad subcarriers can be turned off.

We investigated throughput versus range for extended bandwidth OFDM systems using adaptive loading schemes. Figure 8 shows simulation results for an indoor environment with an exponential channel path loss model and one additional wall (exponent equal to 2.15) [6]. We used the ABL scheme in the extended mode (256-QAM and R = 7/8) for OFDM systems with various frequency bandwidths (20, 40, 60, 80 MHz) and different total transmit power. OFDM systems with ABL schemes demonstrate both maximum throughput improvement and range extension.

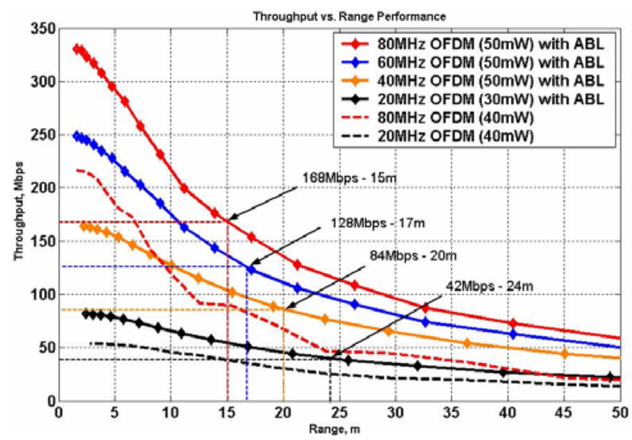

Figure 8: Throughput versus range performance for extended bandwidth OFDM systems with ABL scheme

Degradation of ABL Throughput Performance Due to Time Delay of Data Transmission

To study ABL throughput performance degradation in a non-stationary environment we used an indoor channel model with Doppler spread. The impulse response of a frequency-selective channel with Doppler spread was modeled as a tapped delay line where tap coefficients were independent Rayleigh fading quantities (Jakes model) [7]. In this channel model, the tap coefficients vary slowly because of mobile and scatterer motion. To illustrate this, Figure 9 shows variations of the channel transfer function for one random 20 MHz channel realization during a 10 ms time interval.

A variation in the channel transfer function during packet exchange may lead to performance degradation for any adaptive bit and power-loading scheme.

This channel model was used to investigate Doppler spread affects on the performance of an OFDM system with ABL. The channel is measured by STA2 receiving (for example) an RTS frame from STA1, and the optimal bit and power loading is calculated at the end of the RTS frame. These results are sent to STA1 in some response frame (for example, CTS) and applied by STA1 during transmission of the subsequent data packet to STA2. This data packet from STA1 is sent through a channel that is changed due to Doppler spread and exchange time delay.

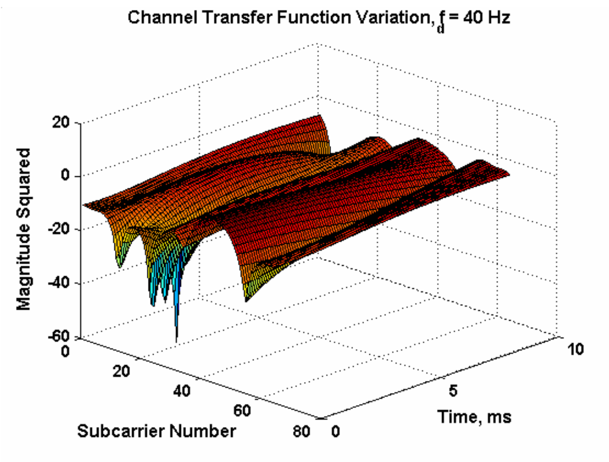

Figure 9: Channel transfer function variation due to Doppler spread with f_{dmax}= 40 Hz

We have investigated the throughput performance degradation of an OFDM system with ABL versus time delay between channel state measurement (by STA2) and ABL data packet transmission (by STA1). The length of the data packet was 1000 bytes.

Simulations of the OFDM system with ABL have shown that even for the modeled severe Doppler spread (f_{dmax} = 40 Hz), throughput degradation no worse than 1 dB is seen when time delay is about 2 – 2.5 ms.

Channel Feedback Information on Backwards Compatibility Issues

Increasing the frequency bandwidth of OFDM signals and exploiting adaptive bit and power-loading techniques requires development of a packet structure for the extended bandwidth OFDM physical layer (PHY). A new packet structure is required that can overcome the frequency selectivity and non-stationary behavior of the wider channel, and that can provide backwards compatibility with existing IEEE 802.11a systems operating in the same frequency band.

We propose to respect the existing time- and frequency-related parameters associated with the OFDM PHY layer convergence protocol (PLCP) of IEEE 802.11a. For example, all simulations for extended bandwidth OFDM systems were done assuming that OFDM data symbols are generated by replication in the frequency domain of conventional 20 MHz OFDM symbols.

To reduce interference from legacy IEEE 802.11a devices, we suggest the use of a “compatibility preamble” mechanism. The compatibility preamble may be constructed from one to four IEEE 802.11a PLCP headers (including preamble and signal field) spread in the frequency domain. This provides support for the coexistence of extended bandwidth OFDM systems with legacy wireless devices. The compatibility preamble allows legacy devices to receive the beginning of an extended bandwidth OFDM transmitted packet. Although they cannot decode the remainder of the packet, they honor the duration specified in the legacy PLCP header, thereby providing protection to the extended bandwidth OFDM packet from legacy interference. The compatibility preamble can also be used for the conventional receiver tasks of channel estimation, timing synchronization, and frequency offset estimation.

In order to exchange feedback service information between extended bandwidth OFDM stations needed for ABL, we propose to insert into the OFDM packet an additional OFDM PLCP header that can be understood only by next-generation WLAN devices. It contains channel state feedback and is modulated with the most robust modulation and encoded with the slowest code. 3 to 5 OFDM symbols of feedback (BPSK with R=1/2) should be enough for successful operation of extended bandwidth OFDM systems with ABL. The OFDM PLCP headers also contain any additional parameters that are required to manage the packet exchange.

ADVANCED CODING

Nearly all wireless systems employ Forward Error Correction (FEC) techniques to correct the numerous transmission errors that occur in wireless channels. As shown previously, Shannon provided a means to calculate a theoretical maximum transmission rate for Additive White Gaussian Noise (AWGN) channels at which error-free transmission could be achieved. It took almost fifty years from the development of Shannon’s theoretical limits until FEC codes were developed that could reasonably approach those performance levels. Only recently are these types of codes making it into the commercial market, and research continues into developing more efficient solutions.

The existing IEEE 802.11 wireless standards all specify a well-known convolutional FEC code, and the Viterbi algorithm is universally used for decoding these codes. At the time of the development of the current standards this was the most practical solution considering cost, complexity, power consumption, and decoding latency. Unfortunately convolutional codes and the Viterbi algorithm leave a significant amount of power efficiency unclaimed with respect to the theoretical capacity limits.

About ten years ago the FEC community experienced a revolution with the discovery of Turbo codes, which were the first practical codes developed that could come reasonably close to theoretical capacity. Even after ten years of additional research, Turbo codes have seen limited acceptance in certain applications due to their complexity and decoding latency. One of the difficulties with Turbo codes is a loss of performance when the size of the code block is reduced. For streaming applications like broadcast video, this is not a problem, but WLAN packet sizes are often very small and transmissions as simple as packet acknowledgements require link reliabilities comparable to long packets. This degradation of performance for short blocks as well as their implementation complexity makes it difficult to successfully apply Turbo codes to WLAN applications.

The discovery of Turbo codes, which use iterative decoding to achieve much of their performance, renewed research interest in another family of FEC codes known as Low Density Parity Check (LDPC) codes. Although this type of code was initially investigated in 1963 by Gallagher [8] little research was devoted to them until the publication of Turbo codes sparked interest in iterative techniques. Recent research has shown that carefully designed LDPC codes do not suffer the same performance degradation for short data blocks as Turbo codes. The underlying codes in LDPC codes are very simple parity-check relationships, so the potential for complexity reduction compared to Turbo codes is significant.

Figure 10 shows a performance comparison of a type of LDPC code developed by Intel for possible use in WLAN systems against the Viterbi decoder currently used in 802.11 systems. In this case, blocks of 2193 data bits are used to generate 731 parity bits that are appended to make a 2924-bit codeword (i.e., a (2924, 2193) code). This provides a code of the same rate as the R = ¾ convolutional code used for comparison. Longer data packets would be coded by concatenating successive code words. Shorter packets would be accommodated by using code shortening, where the encoder and decoder assume the remainder of the data block is zeros. Since the assumed zero-padding values are known, the unused portion of the shortened code word does not need to be transmitted.

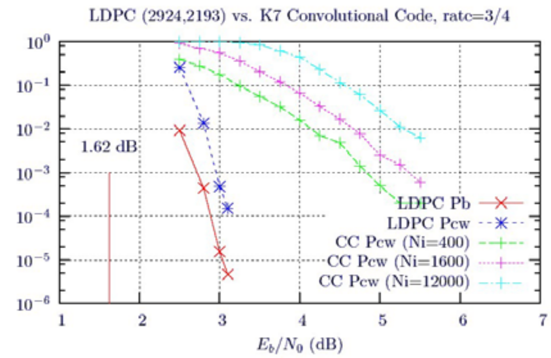

Figure 10: Comparison of Bit Error Probabilities (Pb) and Code Word Error Probabilities (Pcw) for the K = 7 convolutional code used in the current 802.11 standards and a proposed (2924, 2193) LDPC code

Shortening the LDPC code for smaller blocks maintains a significant improvement over the current system. Capacity for an R = ¾ code is shown at Eb/No = 1.62 dB. It can be seen that the LDPC code is operating less than 1.5 dB from capacity. Performance for the convolutional code is shown for data blocks of 400, 1600, and 12000 bits.

The performance increase shown in Figure 10, where the LDPC codes require about 3 dB less transmit power to achieve Pcw = 10^{-5}, can be used to increase range or to increase the selected modulation order for higher throughput.

HIGH-THROUGHPUT MAC CONSIDERATIONS

Why do we need to modify the MAC at all? Isn't getting high throughput largely a PHY issue? The answer is “no” because while the PHY data rates are increasing significantly, the PHY overhead is not decreasing by the same factor. Therefore, throughput becomes increasingly dominated by overhead (radio turnaround times, modem pipeline delay, preamble, PLCP headers, etc.).

In order to make the most efficient use of the wireless medium, we need to take a system view and design the operation of the MAC so that it effectively uses and manages the services of the PHY while providing an unmodified interface to the higher layers.

Interface to Higher Layers

The MAC provides a connectionless (so-called UNITDATA) interface to higher layers. This means that each packet sent down from the higher layers is treated by the MAC independently of all others. The MAC preserves order between any source/destination address pair.

While there is an increasing amount of knowledge in the higher layers used to manage the MAC through its management interfaces, from the data point of view, it is still a "wireless Ethernet."

MAC Improvements

The improved MAC has two main jobs to do: (1) it manages the new features of the PHY and (2) it improves throughput by aggregation and scheduling.

In conventional 802.11 systems, the MAC can manage the PHY by adapting its transmit rate to observed conditions (you might see this as an "automatic" rate setting on your control panel). There are a number of things it can attempt to observe to make these decisions, including packet error rate. However, without explicit communication from its peer about receive conditions, its decisions are based on possibly flawed assumptions. Likewise (but less frequently attempted), conventional systems may try to automatically adapt fragment size.

The WLAN system described here gains significant throughput improvements by adapting its transmit parameters to closely match the properties of the link between two devices. The more capacity the link has, the more throughput the system will obtain. It does this through a training process that can be repeated as often as necessary to keep training data fresh.

The other new main job the MAC has is to create bursts of packets addressed to the same destination. It sends these using the 802.11e Block Ack mechanism thereby significantly reducing overhead.

MAC Tradeoffs

In order to create bursts, the MAC has to buffer traffic by destination address until some criterion (e.g., burst length or age) makes it eligible for transmission. This creates a scheduling problem. We have to find a balance between fairness, delay, and throughput. This is an ongoing topic of research.

A next-generation MAC will also support 802.11e Quality of Service (QoS) enhancements (this is necessary in order to use 802.11e Block Ack). 802.11e supports negotiation of the QoS attributes of a traffic stream through a "traffic specification" (TSPEC). The TSPEC allows a next-generation WLAN device to respect traffic-specific delay limits in its formation of bursts.

Another tradeoff the MAC can manage is the liveness of training data. The MAC obtains training data through an exchange of training frames with a particular destination. This training exchange is an overhead. The MAC can track the age of training data and decide when it is likely to be out of date.

When the MAC decides to transmit to a destination, it can optimize throughput based on the age of the training data and the amount of data to be sent. For example, it is not worth the overhead of re-training just to send small amounts of data.

Coexistence Options

It is the primary job of the MAC to manage access to the wireless medium.

We want to make efficient use of the wireless medium, but we want to be fair to existing users.

In most deployment scenarios, next-generation access points (APs) will be added to create a new network or to overlay an existing network. In these cases, the best coexistence option is for the new APs to select an operating channel that avoids creating any overlap with legacy 802.11 devices sharing the same band. In the case of enterprise deployments, frequency planning would make this happen. In the home environment, the lower installed density makes it very likely the new APs can automatically find a clear operating channel.

If the next-generation device cannot find a channel clear of legacy devices it can attempt to negotiate with legacy devices (e.g., using 802.11h signaling) to get them to relocate.

If this fails, the network has to operate in a so-called "mixed" mode that requires the use of a protection mechanism. It also requires that such devices be capable of sending and receiving legacy frames so that they respect medium reservation of the legacy devices. The protection mechanism uses the transmission of legacy frames (or compatibility preambles, as described above) so that the legacy devices respect either MAC-level or PHY-level medium "reservations."

A next-generation AP controls whether its stations operate in mixed mode. An alternative to mixed mode operation is for the AP to manage access to the wireless medium by preventing groups of stations from contending for access by setting their Network Allocation Vector (NAV), for example using beacons. This AP avoids the overhead of the protection mechanism and can control the amount of time the medium is used by legacy or high-throughput devices according to its own local policy.

CONCLUSION

This paper has shown that there are techniques that can improve throughput and range for next-generation wireless LAN devices. MIMO-OFDM techniques, fast link adaptation, and advanced coding schemes are likely to be used in next-generation WLAN products. These techniques neatly dovetail together to provide a system including both MAC and PHY layers that increases channel capacity and makes the most of this capacity.

Intel is contributing to and committed to work towards the development of this standard using these kinds of techniques.

ACKNOWLEDGMENTS

The authors acknowledge Alexei Davydov, Andrey Pudeyev, Vadim Sergeyev, and Sergey Tiraspolsky, who provided the simulation results for WB OFDM with fast link adaptation techniques. They also acknowledge Bo Xia for providing the LDPC simulation results and Sumeet Sandhu and Qinghua Li for contributing to the MIMO simulations.

REFERENCES

[1]Chow, P.S., Cioffi, J.M., and Bingham, J.A.C., "A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels," *IEEE Transactions on communications*, V. 43, p. 773, 1995.

[2] Fisher, R.F.H. and Huber, J.B., "A new loading algorithm for discrete multitone transmission," *Proc. IEEE GLOBECOM'96*, p. 724.

[3] Leke, A. and Cioffi, J.M., "Transmit optimization for time-invariant wireless channel utilizing a discrete multitone approach," *Proc. of IEEE ICC'97*, V.2, pp. 954 – 958.

[4] Keller, T. and Hanzo, L., "Adaptive modulation techniques for duplex OFDM transmission," *IEEE Transactions on Vehicular Technology*, Vol. 49, Issue 5, Sept. 2000, pp. 1893 – 1906.

[5] Simoens, S. and Bartolome, D., "Optimum performance of link adaptation in HIPERLAN/2 networks," *Vehicular Technology Conference*, Vol. 2, May 2001, pp. 1129 – 1133.

[6] Rhodes, V.J. and Jacobsen, E.A., "Static propagation characteristics for 5 GHz Wireless LAN applications," (to be published).

[7] Jakes, W.C., *Microwave Mobile Communications*, John Wiley & Sons Inc., New York, 1974.

[8] R.G. Gallager, "Low Density Parity Check Codes," *IRE Trans. Information Theory*, Vol. IT-8, pp. 21-28, Jan. 1962.

AUTHORS' BIOGRAPHIES

Boyd Bangerter manages the System Architecture and Standards group in Intel's Corporate Technology Group. He joined Intel in 1993 and has managed various programs related to wired and wireless broadband access technologies for the past ten years. He holds an MBA degree from Brigham Young University and a BSEE degree from the California Institute of Technology. His e-mail is boyd.r.bangerter@intel.com.

Eric Jacobsen is a wireless systems engineer in Intel's Corporate Technology Group. He is Principal Investigator in the Advanced OFDM research group where his current focus is on low-cost high-speed WLAN solutions. He is a past member of the IEEE 802.16 Working Group and is an active member of the 802.11 High Throughput group.

Eric has 18 years of experience in research and development in signal processing and communications in radar, avionics, satellite, and network applications. He has BSEE and MSEE degrees from South Dakota School of Mines and Technology. His e-mail is eric.a.jacobsen@intel.com

Minnie Ho leads a group conducting Smart Antenna Research in Intel's Corporate Technology Group. She joined Intel in 2001 as part of the Ultra-Wideband Research group. Prior to joining Intel, she worked at four start-up companies on projects involving ADSL (technology standardized by ITU), broadband wireless access, fiber-optic, and ultra-wideband communications. She holds M.S. and Ph.D. degrees from Stanford University in Electrical Engineering, and a B.S. degree in Electrical Engineering from Princeton University. Her e-mail is minnie.ho@intel.com

Alexander Maltsev is a staff research scientist in Intel's Wireless Networking Group. He leads the WPD Advanced Development Russian team, which works on research and development of signal processing algorithms for wireless communication systems.

Alexander has over 30 years experience in the development of analog and digital signal-processing algorithms, synchronization systems, adaptive antenna arrays, adaptive active vibration and noise control systems; and space-time signal processing in non-stationary environments. He received the Candidate of Science (Ph.D.) and Doctor of Science degrees from Nizhny Novgorod State University (NNSU), Russia, in 1975 and 1990, respectively, all in Radiophysics. Since 1994 he holds the chair in Statistical Radiophysics at NNSU. His e-mail is Alexander.Maltsev@intel.com

Alexey Rubtsov is a researcher in Intel's Wireless Networking Group. He works in the WPD Advanced Development team where his current focus is on adaptive bit and power-loading algorithms for OFDM systems. He is a student member of the IEEE.

He received his B.S. and M.S. degrees with honors from Nizhny Novgorod State University, Russia in 1999 and 2001, respectively. He is currently pursuing a Ph.D. degree in Radiophysics at NNSU. His e-mail is Alexey.Rubtsov@intel.com

Ali Sadri received his Bachelors, Masters, and Ph.D. degrees in Electrical Engineering in 1989, 1991, and 2000, respectively from North Carolina State University. From 1990 through 2000 he was with IBM in the Research Triangle Park, NC, where he was initially a member of the voice band modem and multimedia team. During the last seven years of Dr. Sadri's work at IBM, he became the program manager for IBM telecommunications standards where he was leading and coordinating IBM contribution and strategy at the International Telecommunications Union (ITU) and Telecommunications Industry Association (TIA). Dr. Sadri later joined BOPS Inc. During the years of 2000-02 he was the director of Telecommunications Development where he led the effort to design a low-power programmable wireless LAN DSP core. In the past two years he has been an adjunct associate professor at Duke University and a lecturer at Oregon State University. He is currently working at Intel Corporation and leading the advanced WLAN development effort. His areas of research are adaptive modulation, Adaptive TPC, WCDMA and OFDM. His e-mail is ali.s.sadri@intel.com.

Adrian Stephens is a wireless MAC architect in Intel's Corporate Technology Group. His focus is on developing next-generation IEEE 802.11 standards. He is an active member of the 802.11 High Throughput group (currently chair of their usage model committee).

Adrian has over 20 years of product development and research experience working for industry and government. Adrian received a B.A. degree and a Ph.D. degree from Cambridge University. He is a member of the IEEE, the IEE, and is a chartered engineer. His e-mail is adrian.p.stephens@intel.com.

For further information visit:

developer.intel.com/technology/itj/index.htm