

Ultra-Wideband Technology for Short- or Medium-Range Wireless Communications

Jeff Foerster, Intel Architecture Labs, Intel Corp.
Evan Green, Intel Architecture Labs, Intel Corp.
Srinivasa Somayazulu, Intel Architecture Labs, Intel Corp.
David Leeper, Intel Connected Products Division, Intel Corp.

Index words: UWB, wireless, communications, LAN, PAN

ABSTRACT

Ultra-Wideband (UWB) technology is loosely defined as any wireless transmission scheme that occupies a bandwidth of more than 25% of a center frequency, or more than 1.5GHz. The Federal Communications Commission (FCC) is currently working on setting emissions limits that would allow UWB communication systems to be deployed on an unlicensed basis following the Part 15.209 rules for radiated emissions of intentional radiators, the same rules governing the radiated emissions from home computers, for example. This rule change would allow UWB-enabled devices to overlay existing narrowband systems, which is currently not allowed, and result in a much more efficient use of the available spectrum. Devices could, in essence, fill in the unused portions of the frequency spectrum in any particular location.

These recent developments by the FCC give Intel a unique opportunity to develop equipment that could potentially take advantage of the vast amount of usable spectrum that exists in the wireless space, and that could provide an engine to drive the future high-rate applications that are being conceived throughout this industry.

Intel® Architecture Labs (IAL) is currently researching UWB technology in order to better understand its benefits, limitations, and technical challenges when used for high-rate communications. This paper introduces the reader to this technology, from potential applications to regulatory hurdles, to possible implementations and future challenges.

INTRODUCTION

Ultra-Wideband (UWB) technology has been around since the 1980s, but it has been mainly used for radar-based applications until now (see [1] and the references

therein), because of the wideband nature of the signal that results in very accurate timing information. However, due to recent developments in high-speed switching technology, UWB is becoming more attractive for low-cost consumer communications applications (as detailed in the “Implementation Advantages” section of this paper). Intel Architecture Labs (IAL) is currently working on an internally funded research project whose intent is to further explore the potential benefits and future challenges for extending UWB technology into the high-rate communications arena.

Although the term *Ultra-Wideband* (UWB) is not very descriptive, it does help to separate this technology from more traditional “narrowband” systems as well as newer “wideband” systems typically referred to in the literature describing the future 3G cellular technology. There are two main differences between UWB and other “narrowband” or “wideband” systems. First, the bandwidth of UWB systems, as defined by the Federal Communications Commission (FCC) in [2], is more than 25% of a center frequency or more than 1.5GHz. Clearly, this bandwidth is much greater than the bandwidth used by any current technology for communication. Second, UWB is typically implemented in a carrierless fashion. Conventional “narrowband” and “wideband” systems use Radio Frequency (RF) carriers to move the signal in the frequency domain from baseband to the actual carrier frequency where the system is allowed to operate. Conversely, UWB implementations can directly modulate an “impulse” that has a very sharp rise and fall time, thus resulting in a waveform that occupies several GHz of bandwidth. Although there are other methods for generating a UWB waveform (using a chirped signal, for example), in this paper, we focus on the impulse-based UWB waveform—due to its simplicity. But, first, a breakdown of how this paper is organized.

The first section looks at UWB technology from the high-level perspective of how this technology compares with other current and future wireless alternatives. Next, we describe the current state of the regulatory process, where UWB transmissions are under consideration for being made legal on an unlicensed basis. Then, some implementation advantages of UWB systems are discussed that distinguish UWB transceiver architectures from more conventional “narrowband” systems. After this, we illustrate the throughput vs. distance characteristics for an example UWB system.

The high data rates afforded by UWB systems will tend to favor applications such as video distribution and/or video teleconferencing for which Quality of Service (QoS) will be very important. So, in addition to describing the physical layer attributes of UWB systems, it’s important to keep in mind the Medium Access Control (MAC) layer as well. Therefore, we have also devoted a section to describing the current mechanisms that exist to support the required QoS for these high-rate applications. Finally, we conclude with a summary of the benefits of UWB and suggest some future challenges that are currently being investigated by IAL.

WIRELESS ALTERNATIVES

In order to understand where UWB fits in with the current trends in wireless communications, we need to consider the general problem that communications systems try to solve. Specifically, if wireless were an ideal medium, we could use it to send

1. a lot of data,
2. very far,
3. very fast,
4. for many users,
5. all at once.

Unfortunately, it is impossible to achieve all five attributes simultaneously for systems supporting unique, private, two-way communication streams; one or more have to be given up if the others are to do well. Original wireless systems were built to bridge large distances in order to link two parties together. However, recent history of radio shows a clear trend toward improving on the *other four attributes* at the expense of distance. Cellular telephony is the most obvious example, covering distances of 30 kilometers to as little as 300 meters. Shorter distances allow for spectrum reuse, thereby serving more users, and the systems are practical because they are supported by an underlying *wired* infrastructure—the telephone network in the case of cellular. In the past few years, even shorter range systems, from 10 to 100 meters, have begun emerging, driven primarily by data applications. Here, the Internet is the underlying wired

infrastructure, rather than the telephone network. Many expect the combination of short-range wireless and wired Internet to become a fast-growing complement to next-generation cellular systems for data, voice, audio, and video. Four trends are driving short-range wireless in general and ultra-wideband in particular:

1. The growing demand for wireless data capability in portable devices at higher bandwidth but lower in cost and power consumption than currently available.
2. Crowding in the spectrum that is segmented and licensed by regulatory authorities in traditional ways.
3. The growth of high-speed *wired* access to the Internet in enterprises, homes, and public spaces.
4. Shrinking semiconductor cost and power consumption for signal processing.

Trends 1 and 2 favor systems that offer not just high-peak bit rates, but high *spatial capacity*¹ as well, where spatial capacity is defined as *bits/sec/square-meter*. Just as the telephone network enabled cellular telephony, Trend 3 makes possible high-bandwidth, in-building service provision to low-power portable devices using short-range wireless standards like Bluetooth* (<http://www.bluetooth.com>) and IEEE 802.11 (<http://grouper.ieee.org/groups/802>). Finally, Trend 4 makes possible the use of signal processing techniques that would have been impractical only a few years ago. It is this final trend that makes Ultra-Wideband (UWB) technology practical.

When used as intended, the emerging short- and medium-range wireless standards vary widely in their implicit spatial capacities. For example:

- IEEE 802.11b has a rated operating range of 100 meters. In the 2.4GHz ISM band, there is about 80MHz of useable spectrum. Hence, in a circle with a radius of 100 meters, three 22MHz IEEE 802.11b

¹ The term *spatial capacity* has been used by many, including Prof. Jan Rabaey at the University of California, Berkeley. An equivalent and more descriptive term might be *spatial efficiency*. The late Marc Weiser, Chief Technologist of Xerox PARC, lectured on the importance of spatial capacity in 1996 (<http://www.ubiq.com/hypertext/weiser/NomadicInteractive/>), though at the time he focused on infrared as the medium and bits/sec/*cubic-meter* as the metric. We will use *square-meter* in this paper since the relevant coverage area is usually two-dimensional rather than three-dimensional.

* Other names and brands may be claimed as the property of others.

systems can operate on a non-interfering basis, each offering a peak over-the-air speed of 11Mbps. The total aggregate speed of 33Mbps, divided by the area of the circle, yields a spatial capacity of approximately 1,000 bits/sec/square-meter.

- Bluetooth, in its low-power mode, has a rated 10-meter range and a peak over-the-air speed of 1Mbps. Studies have shown that approximately 10 Bluetooth “piconets” can operate simultaneously in the same 10-meter circle with minimal degradation yielding an aggregate speed of 10Mbps [3]. Dividing this speed by the area of the circle produces a spatial capacity of approximately 30,000 bits/sec/square-meter.
- IEEE 802.11a is projected to have an operating range of 50 meters and a peak speed of 54Mbps. Given the 200MHz of available spectrum within the lower part of the 5GHz U-NII band, 12 such systems can operate simultaneously within a 50-meter circle with minimal degradation, for an aggregate speed of 648Mbps. The projected spatial capacity of this system is therefore approximately 83,000 bits/sec/square-meter.
- UWB systems vary widely in their projected capabilities, but one UWB technology developer has measured peak speeds of over 50Mbps at a range of 10 meters and projects that six such systems could operate within the same 10-meter radius circle with only minimal degradation. Following the same procedure, the projected spatial capacity for this system would be over 1,000,000 bits/sec/square-meter.

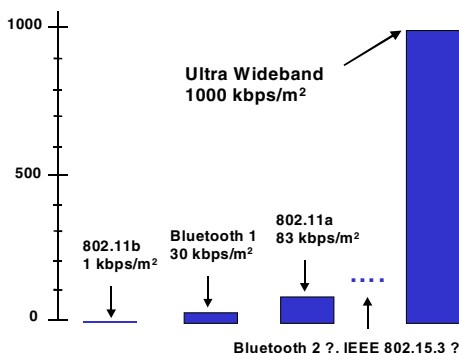


Figure 1: Spatial capacity comparison between IEEE 802.11, Bluetooth*, and UWB

As shown in Figure 1, other standards now under development in the Bluetooth Special Interest Group and IEEE 802 working groups would boost the peak speeds and spatial capacities of their respective systems still further, but none appear capable of reaching that of UWB. A plausible reason is that all systems are bound

by the channel capacity theorem [4], as shown in Figure 2. Because the upper bound on the capacity of a channel grows linearly with total available bandwidth, UWB systems, occupying 2GHz or more, have greater room for expansion than systems that are more constrained by bandwidth.

$$C = B \log_2 \left(1 + \frac{S}{N} \right)$$

Where:
 C = Maximum Channel Capacity (bits/sec)
 B = Channel Bandwidth (Hz)
 S = Signal Power (watts)
 N = Noise Power (watts)

**C grows linearly with B,
 but only logarithmically with S/N**

Figure 2: Channel capacity for additive, white Gaussian noise

Thus, UWB systems appear to have great potential for support of future high-capacity wireless systems. However, there are still several important challenges ahead for this technology before it can be realized. Not the least of these challenges is finding a way to make the technology legal without causing unacceptable interference to other users that share the same frequency space. This is addressed in the next section.

REGULATORY AND STANDARDS ISSUES

The Federal Communications Commission (FCC) is in the process of determining the legality of Ultra-Wideband (UWB) transmissions. Due to the wideband nature of UWB emissions, it could potentially interfere with other licensed bands in the frequency domain if left unregulated. It’s a fine line that the FCC must walk in order to satisfy the need for more efficient methods of utilizing the available spectrum, as represented by UWB, while not causing undo interference to those currently occupying the spectrum, as represented by those users owning licenses to certain frequency bands. In general, the FCC is interested in making the most of the available spectrum as well as trying to foster competition among different technologies.

The FCC first initiated a Notice of Inquiry (NOI) in September of 1998, which solicited feedback from the industry regarding the possibility of allowing UWB emissions on an unlicensed basis following power restrictions described in the FCC Part 15 rules. The FCC Part 15 rules place emission limits on intentional and unintentional radiators in unlicensed bands. These emission limits are defined in terms of microvolts per

meter (uV/m), which represent the electric field strength of the radiator. In order to express this in terms of radiated power (terms that are better understood by communications engineers), the following formula can be used. The emitted power from a radiator is given by the following:

$$P = E_0^2 4\pi R^2 / \eta \quad (1)$$

where E_0 represents the electric field strength in terms of V/m, R is the radius of the sphere at which the field strength is measured, and η is the characteristic impedance of a vacuum where $\eta = 377$ ohms. For example, the FCC Part 15.209 rules limit the emissions for intentional radiators to 500uV/m measured at a distance of 3 meters in a 1MHz bandwidth for frequencies greater than 960MHz. This corresponds to an emitted power spectral density of -41.3dBm/MHz.

In May of 2000, the FCC issued a Notice of Proposed Rule Making (NPRM), which solicited feedback from the industry on specific rule changes that could allow UWB emitters under the Part 15 rules. More than 500 comments have been filed since the first NOI, which shows significant industry interest in this rule-making process. Figure 3 below shows how the current NPRM rules would limit UWB transmitted power spectral density for frequencies greater than 2GHz.




Figure 3: Power spectral density limits in current NPRM

The FCC is considering even lower spectral density limits below 2GHz in order to protect the critical Global Positioning System (GPS) even more, but currently no upper boundary has been defined. Results of a National Telecommunications and Information Administration (NTIA) report analyzing the impact of UWB emissions on GPS, which operate at 1.2 and 1.5GHz, was recently published and suggests that an additional 20-35dB greater attenuation, beyond the power limits described in the FCC Part 15.209, may be needed to protect the GPS band (see www.ntia.doc.gov). However, placing proper spectral density emission limits in the bands that may

need additional protection will still allow UWB systems to be deployed in a competitive and useful manner while not causing an unacceptable amount of interference on other useful services sharing the same frequency space. This report, and others, will be carefully considered by the FCC prior to a final ruling.

The main concern regarding UWB emissions is the potential interference that they could cause to the “incumbents” in the frequency domain as well as to specific critical wireless systems that provide an important public service (for example, GPS). There are many factors which affect how UWB impacts other “narrowband” systems, including the separation between the devices, the channel propagation losses, the modulation technique, the Pulse Repetition Frequency (PRF) employed by the UWB system, and the receiver antenna gain of the “narrowband” receiver in the direction of the UWB transmitter. For example, a UWB system that sends impulses at a constant rate (the PRF) with no modulation causes spikes in the frequency domain that are separated by the PRF. Adding either amplitude modulation or time dithering (i.e., slightly changing the time the impulses are transmitted) results in spreading the spectrum of the UWB emission to look more flat. As a result, the interference caused by a UWB transmitter can be viewed as a wideband interferer, and it has the effect of raising the noise floor of the “narrowband” receiver.

There are three main points to consider when looking at this type of interference. First, if UWB follows the Part 15 power spectral density requirements, its emissions are no worse than other devices regulated by this same standard, which include computers and other electronic devices. Second, interference studies need to consider “typical usage scenarios” for the interaction between UWB and other devices. Using a “worst case” analysis may result in too great a restriction on UWB and could prevent a promising new technology from becoming viable. Third, FCC restrictions are only a beginning. Further coordination through standards participation may be necessary to come up with coexistence methods for operational scenarios that are important for the industry. For example, if UWB is to be used as a Personal Area Network (PAN) technology in close proximity to an 802.11a Local Area Network (LAN), then the UWB system must be designed in such a manner as to peacefully coexist with the LAN. This can be achieved through industry involvement and standards participation, as well as careful designs.

Figure 3 illustrates two other important considerations for UWB systems. First, UWB emissions will be allowed only at a much lower transmit power spectral density compared to other “narrowband” services. This low power can be seen as both a limitation and a benefit. It

restricts UWB emissions to relatively short distances, but results in a very power-efficient and low-cost implementation, which preserves battery life. Second, Figure 3 also shows that UWB systems will most likely suffer from interference from other “narrowband” users. For the most flexible solution, these interferers should be suppressed only on an as-needed basis, thus requiring some sort of adaptive interference suppression technique, which is the subject of research currently within the Intel® Architecture Labs (IAL).

People familiar with the FCC process suggest that rules governing UWB emissions could be finalized as soon as June or July or as late as December of 2001.

IMPLEMENTATION ADVANTAGES

As compared with traditional radio transceiver architectures, the relative simplicity of Ultra-Wideband (UWB) transceivers could yield important benefits. To explore these advantages, consider the following traditional radio architecture, which will be contrasted with an example UWB architecture. In 1918, Howard Armstrong invented the venerable super-heterodyne circuit, which, to this day, is the dominant radio architecture². A contemporary example of a low-cost, short-range wireless architecture is the Bluetooth* radio, an example of which is shown in Figure 4.

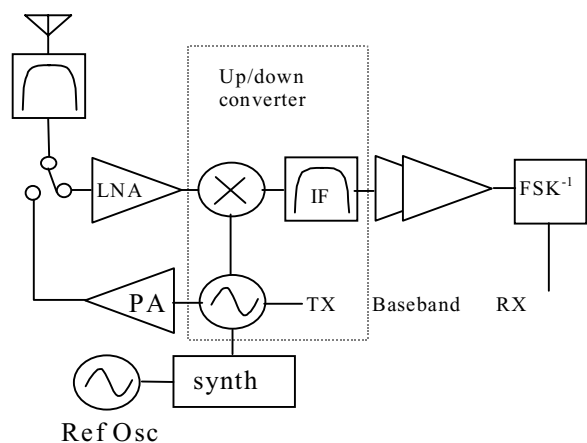


Figure 4: Example Bluetooth* transceiver

Bluetooth uses a form of Frequency Shift Keying (FSK) where information is sent by shifting the carrier

² It can be argued that the super-heterodyne was so effective at processing narrowband RF signals that it accelerated the plan to divide the radio spectrum into successively narrower bands.

* Other names and brands may be claimed as the property of others.

frequency high or low. In Figure 4, this is accomplished by applying the information bits (identified as “TX” in Figure 4) to a Voltage-Controlled Oscillator (VCO). A Phase-Lock Loop (PLL) synthesizer with a crystal reference oscillator is required to keep this oscillator’s average frequency within spec. This 1MHz-wide signal is then spread to 79MHz by a frequency-hopping technique where the synthesizer is tuned to pseudo random channels spaced at 1MHz. The resulting emitted signal is centered at 2.45GHz with a bandwidth of 79MHz.

In receive mode, the extremely weak signal from the antenna is first amplified and then down-converted to an Intermediate Frequency (IF). In this example, IF = 120MHz. The down-converter uses a heterodyne [5] technique where a non-linear “mixer” is fed both the desired signal at ~2.45GHz and a synthesized local oscillator that operates at a frequency of 120MHz either above or below the desired signal. The mixer produces a plethora of images of the desired signal where each image is centered at the sum and difference terms of the desired signal and the local oscillator (and harmonics of both). The image that falls at the desired IF frequency then passes through the IF filter, while the other images are rejected. At this low frequency, it is relatively easy to provide the stable high-gain (~90dB) circuits needed to demodulate the signal and recover the original information. Note that in higher performance radio systems, such as cellular phones, two or even three down conversion stages may be employed.

Most Bluetooth designs are based on variants of this super-heterodyne architecture with an emphasis on integrating as many functions as possible onto a single chip. In some designs, this includes the IF filters which make even Bluetooth’s relatively relaxed channel selectivity requirements very difficult to realize over operating temperature.




Figure 5: Example UWB transceiver architecture

We can now look at a prototypical UWB transceiver as shown in Figure 5. This transceiver could be used for the

same applications targeted for use with Bluetooth, but at higher data rates and lower emitted Radio Frequency (RF) power. The information could be modulated using several different techniques: the pulse amplitude could be modulated with +/-1 variations (bipolar signaling) or +/-M variations (M-ary Pulse Amplitude Modulation), turning the pulse on and off (known as On/Off Keying or OOK), or dithering the pulse position (known as Pulse Position Modulation or PPM). The pulse has a duration on the order of 200ps and, in this example, its shape is designed to concentrate energy over the broad range of 2-6GHz. A power amplifier may not be required in this case because the pulse generator need only produce a voltage swing on the order of 100mV. As with the super-heterodyne radio, a bandpass filter is used before the antenna to constrain the emissions within the desired frequency band except, in this case, the filter would have a bandwidth on the order of 4GHz.

During continuous transmission, the Bluetooth transmitter is rated to deliver about 1Mbps at an average of 1mW of RF power to the antenna, and it provides an operating range of about 10 meters. Extrapolating from the results shown in the next section, a 2.5GHz wide UWB transmitter operating at < 10uW of average power could provide the same throughput and estimated coverage range. This could translate into a significant battery life extension for portable devices. Alternately, more UWB signal power could be used to increase range or data rate.

In receive mode, the energy collected by the antenna is amplified and passed through either a matched filter or a correlation-type receiver. A matched filter has an impulse response matched to the received pulse shape and will produce an impulse at its output when presented with RF energy which has the correct (matching) pulse shape. The original information is then recovered with an adjustable high-gain threshold circuit.

Notice the relative simplicity of this implementation compared to the super-heterodyne architecture. This transceiver has no reference oscillator, Phase-Lock Loop (PLL) synthesizer, VCO, mixer, or power amplifier. This simplicity translates to lower material costs and lower assembly costs. For example, the inexpensive reference oscillators used in the typical Bluetooth radio require a center frequency adjustment lengthening the test time and hence, increasing the cost of goods sold.

Low-cost Digital Signal Processing (DSP) hardware is often used in modern digital radios to generate several modulation methods. These systems can step down the information density in their signal to serve users at greater distances (range). An advantage of UWB is that even simple implementations can provide this adaptation. For example, as the range increases, a UWB radio can use several pulses to send one information bit thereby

increasing the Signal-to-Noise Ratio (SNR) in the receiver. Since the average power consumption of a UWB transmitter grows linearly with Pulse Repetition Frequency (PRF), it is easy to envision a relatively simple UWB radio that, under software control, can dynamically trade data rate, power consumption, and range. This type of flexibility is what is needed to enable the power-constrained portable computing applications of the future.

However, there are still some design challenges for UWB systems. There is a concern that such a wideband receiver will be susceptible to being unintentionally jammed by traditional narrowband transmitters that operate within the UWB receiver's passband. Also yet to be resolved are issues such as filter matching accuracy and the extreme antenna bandwidth requirements, which can often be difficult to achieve. For a correlator-based receiver, the timing needs to be very accurate in order to properly detect the received pulse due to the short pulse durations. In addition, there appears to be a significant amount of energy in the multipath components caused by reflections in the channel, which suggests that a RAKE-type receiver [6] would significantly improve performance. Lastly, noise from an on-board microcontroller could be an issue. A common trick in narrow band radio systems is to move the noise just out of band rather than suppressing it. This trick may prove elusive given the bandwidth of a UWB receiver.

THROUGHPUT ANALYSIS

As mentioned in the previous section, there are many different modulation methods that could be applied to Ultra-Wideband (UWB) systems. The purpose of this section is to quantify the distance vs. throughput relationship for an example Pulse Amplitude Modulated (PAM) UWB system in order to highlight some of the advantages and constraints of UWB. The results here use the following system assumptions:

- Noise is Additive White Gaussian Noise (AWGN) only (multi-path will be discussed later).
- A target BER of 10^{-3} uncoded is used, which, when combined with coding, should be able to be reduced to $10^{-5} - 10^{-9}$. Note that coding will also have the effect of reducing the overall throughput.
- Transmit power spectral density is limited to -41 dBm/MHz (as specified by Part 15.209).
- An antenna gain of 0dBi is assumed.
- A 5dB link margin is assumed.
- A 6dB noise figure is assumed.

- Operating bandwidth is 2.5GHz for this example (from 2.5GHz to 5GHz to operate between the 2.4GHz ISM band and the 5GHz U-NII band).
- A center frequency of 3.75GHz is assumed (used for computing the distance loss function).
- Channel model³: Free space propagation (i.e., path loss is proportional to the square of the propagation distance), which results in a path loss given by $L(d) = 20 \log(4\pi / \lambda) + 20 \log(d)$, where λ is the carrier wavelength.

The probability of symbol error for an M-PAM system is given by (assuming coherent detection) [6]:

$$P_M = \frac{M-1}{M} \operatorname{erfc} \left(\sqrt{\frac{3k\gamma_b}{M^2-1}} \right) \quad (2)$$

and the probability of a bit error is estimated as the following:

$$P_b = \frac{1}{k} P_M \quad (3)$$

where $M = 2^k$ and γ_b is the Signal-to-Noise Ratio (SNR) per bit. Note that the SNR per symbol is $E_s / \eta_0 = k\gamma_b$, since each symbol carries k bits of information. To get a better understanding of the relative trade-offs that can be made in UWB systems by varying the pulse bandwidth and pulse repetition period (defined as T_p , which is the time between transmitted pulses), consider the SNR per symbol as the following:

$$E_s / \eta_0 = P_{ave} T_p / \eta_0 = [P_{sd} / \eta_0] \times [B_s / B_p] \quad (4)$$

where $P_{ave} = B_s P_{sd}$ is the average transmitted power, P_{sd} is the average power spectral density limited by the FCC, B_s is the equivalent occupied bandwidth of the transmitted pulse, η_0 is the noise spectral density, and $B_p = 1/T_p$ is referred to as the pulse repetition

³ The indoor channel model, which is the most suitable for UWB operation due to expected limited transmit power by the FCC, is very complicated and is a function of many factors including the availability of a LOS component, the size of the room, the distance between the transmitter and receiver, the materials of the walls, and the presence of, and the materials of equipment/furniture in the room. Free space propagation is used here for illustrative purposes.

frequency. Therefore, we can view the ratio $N_s = B_s / B_p$ as the “pulse processing gain.” Thus, increasing the occupied bandwidth of the pulse or reducing the Pulse Repetition Frequency (PRF), and equivalently, the overall throughput, the distance achieved by the UWB system can be increased for a fixed average transmit power spectral density. Note that this has the effect of increasing the peak transmit power. This factor is what allows UWB to operate at a very low average transmit power spectral density, while still achieving useful throughput and range.

Using the above equations yields the following required E_s / η_0 (SNR per symbol) for an uncoded BER of 10^{-3} .

k	M	γ_b (dB)	$E_s / \eta_0 = k\gamma_b$ (dB)
1	2	7	7
2	4	10.75	13.75
3	8	15	19.77
4	16	19.5	22.5

Table 1: Required E_s/η_0 for M-PAM systems

Note that as the E_s / η_0 requirement increases, the period separation between the symbols will need to increase for a fixed average transmit power. As a result, the data rate is reduced. Using these numbers, the following graph of throughput versus distance can be plotted for the above assumptions.




Figure 6: Example throughput curves for a UWB-based M-PAM system

Clearly, the above graph shows that UWB really provides the greatest throughput at the closer distances. Of course there are other methods for improving the throughput vs. distance relationship, including increased antenna gain, improved coding gain, reduced noise figure, and greater occupied bandwidth. Also note that a more realistic channel model may have path loss exponentials on the order of 2.8-3 for typical indoor channels that must also be considered.

The results here suggest that higher order M-PAM systems do not improve the throughput as much as using lower order 2-PAM with a higher PRF. This can be understood by recalling that PAM is a very spectral efficient modulation technique, but not necessarily very power efficient. For UWB systems, the spectrum is determined by the shape of the pulse rather than the symbol rate. Therefore, for an AWGN channel, it is reasonable to expect that lower order PAM would result in the best performance. However, if a multi-path channel is considered, the 2-PAM system would potentially experience inter-symbol interference, which could potentially limit its throughput, while the higher order systems with greater pulse repetition periods would be impacted less. More information regarding the performance of M-PAM systems in a multi-path channel can be found in [7].

One of the important advantages of UWB systems is their inherent robustness to multi-path fading [8]. Heuristically, this can be explained as follows. Multi-path fading results from the destructive interference caused by the sum of several received paths that may be out of phase with each other. The very narrow pulses of UWB waveforms result in the multiple reflections caused by the channel being resolved independently rather than combining destructively at the receiver. As a result, the time-varying fading that plagues "narrowband" systems is significantly reduced by the nature of the UWB waveform.

Clearly, the overall system performance is significantly impacted by the channel propagation and multi-path model and assumptions. The researchers in the Intel® Architecture Labs (IAL) are working with university and industry partners to get a better understanding of the UWB propagation environment in order to more accurately predict the performance of UWB systems. This information can also be used to more optimally design transmitters and receivers.

MEDIA ACCESS CONTROL (MAC) ISSUES

As the evolution of wireless networks continues to offer higher and higher data rates, a similar natural evolution is occurring in the kinds of applications that are being envisioned for these networks. Current low data-rate

Wireless Local Area Networks (WLANs) and Wireless Personal Area Networks (WPANs), which have data rates of ~1-10Mbps, are typically used for applications such as packet-switched data and cordless voice telephony, using Time Division Multiple Access (TDMA) voice circuits. Example technologies supporting these applications are the IEEE 802.11b (Wi-Fi)*, Bluetooth*, and HomeRF* networking standards. As the IEEE 802.11 and ETSI BRAN HiperLAN/2* standards (the European equivalent of 802.11) have added physical layer specifications with raw data rates up to 54Mbps, the application space is enlarging to include audio/video applications that are enabled by these higher data rates. These diverse traffic types all have different requirements in terms of the service parameters that quantify the network performance for a user of each of those applications. Thus, for example, voice telephony and video teleconferencing applications place tough demands on the latency and jitter performance. Audio/video applications require large amounts of bandwidth and may need close synchronization (e.g., connecting stereo speakers in a surround sound system). Ultra-Wideband (UWB) systems, with their potential for extremely large data rates over short distances, are naturally going to be used for networking these kinds of high-bandwidth/delay-critical data sources and sinks. Hence, it would be natural to look at the approaches to the MAC design undertaken in these other standards when considering the MAC layer design for UWB systems.

The most important functions of the MAC layer for a wireless network include controlling channel access, maintaining Quality of Service (QoS), and providing security. Wireless links have characteristics that differ from those of fixed links, such as high packet loss rate, bursts of packet loss, packet re-ordering, and large packet delay and packet delay variation. Furthermore, the wireless link characteristics are not constant and may vary in time and place. The mobility of users poses additional requirements, as the end-to-end path may be changed when users change their point of attachment. Users expect to receive the same QoS after they have changed their point of attachment. This implies that the new end-to-end path should also support the existing QoS (i.e., a reservation on the new path may be required), and problems arise when the new path cannot support the required QoS. Security is obviously an important consideration in wireless networks because, unlike wired networks, the overlaps between networks cannot be controlled. In addition, unauthorized users can also eavesdrop on transmissions. Security is handled through a combination of different means at the MAC layer, and

* Other brands and names are the property of their respective owners.

also may include physical layer properties of the network. In this section, we restrict ourselves to the channel access and QoS functions, and we first look at some current approaches being considered in the standards-setting committees.

In the IEEE 802.11 TGe committee, there is an ongoing project to enhance the 802.11 MAC to provide for prioritized channel access and QoS. The basic channel access function of the 802.11 MAC is the Distributed Coordination Function (DCF), with an optional mode called the Point Coordination Function (PCF) built atop the DCF, which offers a centralized, polling-based communication between stations and a point coordinator. With the PCF, the point coordinator defines a Contention-Free Period (CFP) during which the stations are polled and a Contention Period (CP) during which the normal DCF channel access mechanism holds. A periodic beacon identifies the start of the CFP and the duration. At the current stage, different prioritized channel access mechanisms for an Enhanced DCF (EDCF) mode are being considered. The EDCF mode provides for treating the priorities of different packets (encoded according to 3-bit traffic category tags) by giving them statistically fair access to the medium. This means that packets from the same priority class contend for the medium on an equal basis according to the 802.11 MAC rules. Packets from different priority classes contend on a weighted basis, where the higher priority packets get a higher probability of success for channel access. Thus, higher priority classes cannot, in principle, choke transfer of lower priority class traffic. In addition to the EDCF modes, a type of point coordination function called the Hybrid Coordination Function (HCF) is also being proposed. The HCF mechanism provides for contention-free and controlled-contention transfers during any part of the frame (i.e., CFP or CP) by allowing the Hybrid Coordinator (HC) to generate bursts of CFPs, as opposed to a monolithic CFP. Thus, the HC can essentially create a number of “mini-CFPs” within the CP, as needed to meet traffic specs. Using this means the HCF promises to provide a flexible scheme where, for example, traffic classes that require periodic transmission opportunities can be accommodated within the CP or the CFP. Traffic that is burstier in nature is handled through the prioritized EDCF mechanism during the CP. In addition, this concept of CFP bursts is expected to mitigate the inter-cell interference that is a problem with the centrally controlled PCF mode when the cells are overlapping in extent.

HiperLAN/2 (HL2) systems, the European counterpart of 802.11, present a very different approach to the MAC and QoS design for high-data rate systems. Where the 802.11 MAC has roots in Ethernet and IP, and the QoS enhancements are seeking to maintain backward compatibility, the HL2 MAC is based on Wireless ATM

concepts and does not have these backwards compatibility requirements. HL2 differs from 802.11 fundamentally in that it uses very short, fixed-length packets, a centrally controlled random access resource reservation channel, and a TDMA kind of resource allocation that is based on successful resource reservation attempts. This kind of architecture potentially offers good QoS performance for streaming sources. However, many feel that the complexity of implementation is quite high compared with the 802.11 MAC. In addition, some studies have shown that in uncoordinated deployment scenarios, inter-cell interference from overlapping cells can be a big problem for the centrally controlled HL2 systems.

In designing a MAC for high-data rate UWB systems, some of the particular properties of the transmission system will obviously dictate many of the design choices. UWB systems offer some unique abilities such as precise position/timing location. This can be exploited at the MAC layer, for example, to synchronize the received packets at different receivers of a multi-cast network (perhaps multiple audio speakers/video displays). UWB systems are also flexible, trading off throughput for range since the Pulse Repetition Frequency (PRF) and the peak pulse power can vary inversely to provide for constant average power. This can be used at the MAC layer to provide for signaling of different data rates on a per-packet or per-link basis, depending on the range of that particular communication. Ultra-Wideband (UWB) systems could also be designed using spread spectrum codes, which may offer better coexistence with other UWB systems, so that unplanned deployment of UWB networks in homes is facilitated. At the MAC layer, this may also result in different choices being available— notably, the problems that centrally coordinated MAC schemes face with overlapping networks may be mitigated. Another choice at the MAC layer that is available is the use of Code Division Multiple Access (CDMA) in addition to TDMA or carrier sense multiple access (CSMA). CDMA also offers the possibility of using techniques such as multi-user detection to boost the system capacity.

One promising application that has been envisioned for UWB is cable replacement for audio/video devices, which would be using wired IEEE 1394 connections (see [9] for more details on IEEE 1394). This could serve two functions: it could act as a wireless “bridge” between clusters of 1394 nodes or as wireless 1394 connections to leaf nodes. Currently, there are efforts underway to enable IEEE 802.11a and HiperLAN/2 systems with this functionality, and UWB systems are a good candidate for the next generation of wireless 1394 systems as well. Another area of great interest is the coexistence of UWB with WLAN systems such as 802.11a, given that WLAN and high-rate WPAN (for which UWB is a good

candidate) systems are likely to be located in close proximity in various systems such as PCs and home network gateways. This opens the door to various solutions for the coexistence of such networks, which will include both physical and MAC layer solutions. One of the important considerations for the success of UWB systems is the compatibility and coexistence of such systems with other WLANs or WPANs, and these considerations should play a big role in the design of the MAC.

As UWB technologies move towards standards and products, one of the decisions that will have to be made is whether to adopt some of the MAC approaches already being developed for other wireless networks, or to develop entirely new approaches. It remains to be seen whether the existing approaches offer the right capabilities for UWB applications. In addition, it is likely several UWB-specific requirements would need to be added to these MACs. On the other hand, some level of compatibility with existing MACs may promote user and market acceptance.

CONCLUSIONS AND FUTURE CHALLENGES

This paper has identified several areas that show the promise of UWB for use in high-rate, short- to medium-range communications. These include potential low-cost implementations, low-power consumption due to limits on transmit power spectral density, high throughput afforded by the wide occupied bandwidth, accurate position location that can be combined with communications capabilities, and favorable multi-path fading robustness due to the nature of the short impulse.

However, there are still challenges in making this technology live up to its full potential. The regulatory process is still in motion. Intel is involved in helping the Federal Communications Commission (FCC) identify emission limits favorable to Ultra-Wideband (UWB) systems that allow them to be competitive within the marketplace, while at the same time not allowing them to cause an unacceptable level of interference for other wireless services that happen to be sharing the same frequency band. The FCC regulations are just a first step in this process, and it is anticipated that standardization will be needed in the future to help make this technology ubiquitous in the consumer market.

In addition, we have identified three main areas that are important for helping UWB make the best use of this newly available spectrum. First, as discussed previously, a reliable channel model is critical for helping to predict performance as well as for optimizing the physical layer design. In this regard, Intel is actively engaging the industry to help determine a reliable model that systems engineers can use to help study the performance of UWB

systems. Second, we are investigating several receiver designs that will help to improve the robustness and long-term viability of this technology. This includes the ability to capture the significant amount of energy that will be present in the multiple reflections caused by the channel (i.e., something analogous to a RAKE receiver often used in CDMA systems), and mechanisms for suppressing the "narrowband" interference that will typically be seen in this type of overlay environment. Finally, we are investigating the feasibility for high-level silicon integration in order to yield a very low-cost and low-power solution. Intel® Architecture Labs (IAL) is actively involved in all of these areas and hopes to advance the state-of-the-art in this technology.

REFERENCES

- [1] J. D. Taylor, *Introduction to Ultra-wideband Radar Systems*, CRC Press, Boca Raton, Florida, 1995.
- [2] FCC Notice of Proposed Rule Making, "Revision of Part 15 of the Commission's Rules Regarding Ultra-wideband Transmission Systems," ET-Docket 98-153.
- [3] Eric Mehofer, "The Performance of Bluetooth in a Densely Packed Environment," *Bluetooth Developers Conference*, December 2000.
- [4] T. M. Cover, J. A. Thomas, *Elements of Information Theory*, John Wiley & Sons, Inc., New York, 1991.
- [5] P. Nahin, *The Science of Radio*, ISBN: 1563963477, AIP Press, College Park, MD, March 1996.
- [6] J. Proakis, *Digital Communications*, 2nd Edition, McGraw-Hill, 1989.
- [7] J. Foerster, "The Effects of Multipath Interference on the Performance of UWB Systems in an Indoor Wireless Channel," submitted to VTC2001, to be presented in May 2001.
- [8] M. Z. Win and R. A. Scholtz, "On the Robustness of Ultra-Wide Bandwidth Signals in Dense Multipath Environments," *IEEE Comm., Letters*, Vol. 2, No. 2, Feb. 1998.
- [9] S. Bard, "Real Time 1394b Data Transfer for Consumer Electronics," *Intel® Developer Update Magazine*, October 2000.

LIST OF ACRONYMS

AWGN	Additive White Gaussian Noise
CDMA	Code Division Multiple Access
CFP	Contention-Free Period
CP	Contention Period
CSMA	Carrier Sense Multiple Access
DCF	Distributed Coordination Function
DSP	Digital Signal Processor
EDCF	Enhanced Distributed Coordination Function
FCC	Federal Communications Commission
FSK	Frequency Shift Keying
GPS	Global Positioning System
HC	Hybrid Coordinator
HCF	Hybrid Coordination Function
IAL	Intel Architecture Labs
IF	Intermediate Frequency
LAN	Local Area Network
MAC	Medium Access Control
NOI	Notice of Inquiry
NPRM	Notice of Proposed Rule Making
NTIA	National Telecommunications and Information Administration
OOK	On/Off Keying
PAM	Pulse Amplitude Modulated
PAN	Personal Area Network
PCF	Point Coordination Function
PLL	Phase-Lock Loop
PPM	Pulse Position Modulation
PRF	Pulse Repetition Frequency
QoS	Quality of Service
RF	Radio Frequency
SNR	Signal-to-Noise Ratio
TDMA	Time Division Multiple Access
UWB	Ultra Wideband
VCO	Voltage-Controlled Oscillator

WLAN	Wireless Local Area Network
WPAN	Wireless Personal Area Network

AUTHORS' BIOGRAPHIES

Jeff Foerster joined Intel in August 2000 as a Wireless Researcher with IAL in Hillsboro, Oregon. He is currently focusing on UWB regulations, system design, and analysis. Prior to joining Intel, he worked on Broadband Wireless Access (BWA) systems and standards. He received his B.S., M.S., and Ph.D. degrees from the University of California, San Diego. His e-mail is jeffrey.r.foerster@intel.com

Evan Green is a Sr. Staff Systems Engineer with IAL in Hillsboro, Oregon. In 20 years of wireless product design, he has contributed in the areas of RF, digital and analog circuits, software, firmware, smart antennas, and cellular systems where he holds four patents. His current research focus is UWB systems and 5GHz wireless LAN standards. His e-mail is evan.r.green@intel.com

Srinivasa Somayazulu joined Intel in July 2000 as a Wireless Architect with IAL's Communication Architecture group. His current interests include system design and algorithm development for high-rate wireless networks including UWB and 5GHz WLANs. He received a B.Tech. degree from the Indian Institute of Technology, Madras, and M.S. and Ph.D. degrees from the University of California, Santa Barbara. His e-mail is v.srinivasa.somayazulu@intel.com.

Dave Leeper joined Intel in February 2000 as CTO for an NBI startup (eWave). Before Intel, he was with AT&T Bell Labs and Bellcore for 23 years and Motorola for 8 years. His current interests are in short-range wireless technologies and related new business development. Dave holds B.S., M.S., and Ph.D. degrees in electrical engineering from Washington University, Cornell, and the University of Pennsylvania, respectively. His e-mail is david.g.leeper@intel.com.

Copyright © Intel Corporation 2001. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://developer.intel.com/sites/developer/tradmarx.htm>.