1 Introducing ultra-wideband (UWB)

If you are interested in a deep theoretical treatise on ultra-wideband, there are several excellent texts, which are listed at the end of this chapter, that we recommend [1, 2]. *Essentials of UWB* will definitely not fill that need. It is far too concise and practical and it fails to take up the requisite three inches of shelf space that are required to fill that niche in the literature.

If you are an engineer, business professional, regulator or marketing person who needs enough technical information to build, sell or regulate products that include a UWB radio, but don't aspire to become a radio frequency (RF) deity in your own right, this is the text that you are looking for. Our objective in writing this book is to provide a dependable overview of the data that you need to know to understand the technology and the industry. This includes technical overviews, industry organization, intellectual property overview, standardization and regulatory discussions. We will also attempt to provide pointers to source documents for deeper investigation for those who are so inclined. We know where the good data are buried because in many cases we had a hand in putting it there. Dr Aiello founded two UWB start-ups, contributed actively to the US regulatory processes, participated in the IEEE standardization wars and performed much of the early development of UWB modulation schemes and radio designs. He has also been a board member in the WiMedia Alliance for a number of years. Mr Wood has been the President of the WiMedia Alliance for several years, has participated in regulatory proceedings in the ITU (International Telecommunications Union) and in Europe, has been a principal architect for the industry structure, has been heavily involved in WiMedia's relationship with Ecma, ISO and ETSI and has played a major role in the development of WiMedia's intellectual-property strategy.

So, with the preliminaries out of the way, it is only appropriate to begin with a description of ultra-wideband. Describing UWB is slightly more difficult than one would immediately suspect. While the press may speak of UWB as a single technology, it is not. Ultra-wideband is best described as a class of radios that use unusually wide bandwidth signals to achieve their application objectives. These radios can operate over a wide range of frequencies, and with different signal characteristics. By way of example to demonstrate the variability in UWB: automotive UWB radars are allowed to operate in the 24 GHz band and require licences; US data communication applications operate over a 3.1–10.6 GHz range and are unlicensed; ground-penetrating radars in the USA operate over the same frequencies as data-communication devices and at the same power levels, but require licensing and have additional restrictions about their use; and the same data-communication applications that operate over 3.1–10.6 GHz in the USA are allowed to operate in a more limited range of frequencies in other parts of the world and with additional restrictions to reduce the probability of generating harmful interference. As one can see, for a single technology, there are very few common elements that can readily be used to define it.

The principal point that is common is that relatively low power is being traded off against wide bandwidths to obtain enough performance to be economically interesting. The very low power levels of UWB would appear to be a major shortcoming, but instead they are offset by the exceedingly wide available bandwidth, which allows the technology to operate at extremely high data rates. Ultra-wideband is distinct from most other commercial technologies because of its underlay philosophy. Figure 1.1 demonstrates the very wideband nature of UWB as compared with more traditional narrowband signals.

Before UWB, the spectrum was divided up by frequency (principally) and only a very restricted overlap of authorized services was allowed. By contrast, UWB is intentionally designed to overlap a broad swath of other services, as is depicted in Figure 1.1. The regulators who initially authorized UWB believe that a higher-priority service would tend to overpower UWB when it needs the spectrum and that UWB would have additional frequencies available that would allow it to move out of the way when this occurs. By implementing an underlay, regulators were attempting to increase the amount of spectrum that was

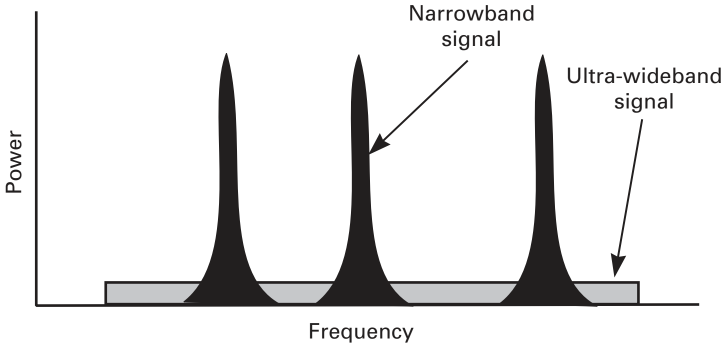

Figure 1.1 Narrowband vs. wideband signals

available for an increasing volume of communication needs. This logic has since been proven to be somewhat less than perfect, but it was the prevailing opinion when the Federal Communication Commission (FCC) made its first rulings in 2002.

1.1 Ultra-wideband application classes

With this description of UWB in hand, the best way to proceed is to provide an abbreviated overview of the applications that employ UWB. If one were to comb through the FCC documents as well as those generated in the ITU and European and Asian proceedings, one would find a more exhaustive list than is being provided here, but also one that tends toward the esoteric. The more practical applications are first grouped into four broad categories including high-data-rate communications, low-data-rate communications, general imaging and automotive radar. The applications are grouped this way because the regulatory proceedings have been crafted around these groupings and will figure in the discussion on regulation during Chapter 8. This structure is also reflective of the market. There are very separate market sectors promoting these products and developing standards.

1.1.1 High-data-rate communications

High-data-rate (HDR) communication is the first application class and the one which will be the principal focus of this text. High-data-rate applications are emerging in the personal computer (PC), consumer electronic (CE) and mobile-handset sectors. These applications share a common need for a very high-data-rate radio that can be built at low cost and that operates within a single room (usually less than 10m). Data rates for HDR designs range from roughly 110 Mbps (set in IEEE 802.15.3a proceedings) to levels exceeding 1 Gbps (the current state of development). High-data-rate applications are clustered across the 3.1–10.6 GHz frequency range. This spectrum was part of the original FCC allocation for UWB and is the basis that the rest of the world used as a starting point in their regulatory proceedings.

To be a little more specific, HDR applications can be broken down to include file transfer, asynchronous communications, streaming video and streaming audio. File transfers may be an exchange that is point-to-point, such as moving music files from two mobile players when schoolchildren wish to share music. Likewise, transferring pictures from a digital camera to a printer, loading movie files to a portable video player or downloading a game from a kiosk would all be considered point-to-point applications.

A point-to-point application is a special-purpose exchange. There are only two parties in the communication. The link is set up exclusively to perform the transaction and is usually taken down when the transaction is concluded. This is a rather special-purpose approach to communication. By contrast, networked communications tend to be a more generalized superset. Instead of assuming that only two devices are active, they are provisioned to be able to handle multiple concurrent conversations across a shared media.

A computer sending a word-processing document to a printer would be an example of a point-to-point exchange, but one that would be likely to be transacted across the more general networked environment. This would allow a scanner or a remote hard drive to share the link with the printer instead of having separate purpose-built radio. Because of the more versatile nature of a networked architecture, the designs incorporated in all HDR implementations deployed today use that approach.

The second group of HDR applications is asynchronous communications. The term 'asynchronous communications' is intended to describe an ongoing, intermittent stream of exchange or conversation. A wireless keyboard or monitor would be an example. In asynchronous communications, the link may remain persistent, and the exchanges are generally thought of less as files and more as individual data blocks.

The third group of HDR applications is streaming video. In this group, a continuous stream of video data is carried by the radio. Usually, this involves feeding a video display or transferring video between two devices for eventual display. Streaming has a normal assumption that the data are being consumed in real time by a person. The term streaming implies that the data are being fed at a consistent rate, which is gauged to be the normal rate of consumption by a person. Streaming video may be moved between a DVD or a game console and a television. As a general rule, streaming video is moved between storage devices and displays. If a human being is not in the direct path, it is frequently more effective to perform a file transfer instead.

Like streaming video, streaming audio has the same assumption of being transferred at a data rate synchronized to the consumption rates that people expect. Unlike streaming video, streaming audio operates at much slower data rates. Ultra-wideband is frequently overkill for this traffic load and would generally only be employed if higher fidelity audio quality is required. Audio distribution also requires synchronization and quality of service. Streaming audio would also be used for wireless speakers in a stereo system or in a home theatre.

Within the high-data-rate category of applications, there will be a noticeable evolutionary shift, which will take place over two generations. The first generation of UWB applications will be all about removal of the cable and replacing it with a high-speed wireless connection. Customers are interested in doing this principally to get rid of the cable mess that hides behind most PCs and televisions. By itself, this has the potential to be a market of satisfactory size and value. It would be reasonable to say that this is what UWB will become. But this isn’t the end of the story.

In the second generation, UWB has the potential to become a short-range, high-speed wireless infrastructure, which connects mobile platforms to stationary platforms and to each other. It enables a wide range of new-use models that were not practical without it. Just as the emergence of the Ethernet enabled email by connecting PCs together, UWB performs the same function for mobile devices. In the PC world, dial-up lines and RS232 cables provided connectivity, but not under terms which were sufficient for the application to operate effectively. In the mobile environment, the WAN (wide area network) and low-rate Bluetooth connections are equally inefficient at connecting the mobile platform to various stationary exchange partners.

As an example, there is a great deal of discussion in the consumer electronics community around the deployment of kiosk networks, which will allow mobile devices to engage a kiosk at a grocery store or train station in order to purchase MP3 files, maps, electronic postcards or other services. Evolving UWB technology could allow you to put the moral equivalent of a kiosk connection into your car or home. Ultra-wideband connections could provide the common touch point that allows mobile platforms to transfer data with stationary systems and the Internet.

1.1.2 Low-data-rate communications

The second major application class is that of low-data-rate communications (LDR). Low-data-rate applications are synonymous with sensor networks for all practical purposes. Inexpensive, low-power sensors are deployed within a building, in a factory, in agricultural fields and a variety of other places. Sensor networks are used for intelligent lighting and energy control within buildings, factory automation and warehousing applications. Typically, these applications involve the transfer of very small volumes of data between battery-powered transceivers. Some applications have a requirement for tracking the physical location of goods and use the unique characteristics of the UWB signal to establish a precise physical location of the transmitter. Sensor networks trade off peak throughput in favour of extended range. The increased range and link margin that they gain by this trade-off allows them to operate a network with low infrastructure costs using very small batteries.

Unlike the HDR market, which uses UWB almost exclusively, there are a number of proprietary UWB radio designs that are employed for LDR applications. While the total sensor market is expected to be quite large, the total market volume will be split between these various radio technologies. Because of this it is unclear at this time whether the UWB contribution to the LDR market will be particularly significant.

1.1.3 Imaging

The third major application category is general imaging. Imaging is a category that includes ground-penetrating radar (GPR), through-wall imaging, in-wall imaging and security perimeters. Ground-penetrating radar is used by utility companies, construction companies and archeologists to look for objects that are below the surface of the ground. Through-wall imaging is used by police and the military to look into adjacent rooms for people, obstacles and hazards. In-wall imaging is used by building construction workers to search for hidden elements such as pipes, wiring and studs in a wall. Security perimeters use the radar properties of UWB to establish a virtual 'fence'. Intruders crossing the fence are detected and an appropriate alarm is issued. Of these general imaging applications, GPR is referenced most often in the regulatory processes and other literature. The volume of GPR units is expected to be very much lower than either HDR or LDR.

1.1.4 Automotive radar

The fourth and final major category of applications is automotive radar for collision avoidance. Figure 1.2 is a graphical representation of the various collision radars that are planned for cars. In these applications, UWB is used by the car as a radar to trigger automated braking when a collision is judged to be imminent. By forcing the car to brake involuntarily, the force of the impact is reduced substantially. Automotive manufacturers expect that this new capability will save a significant

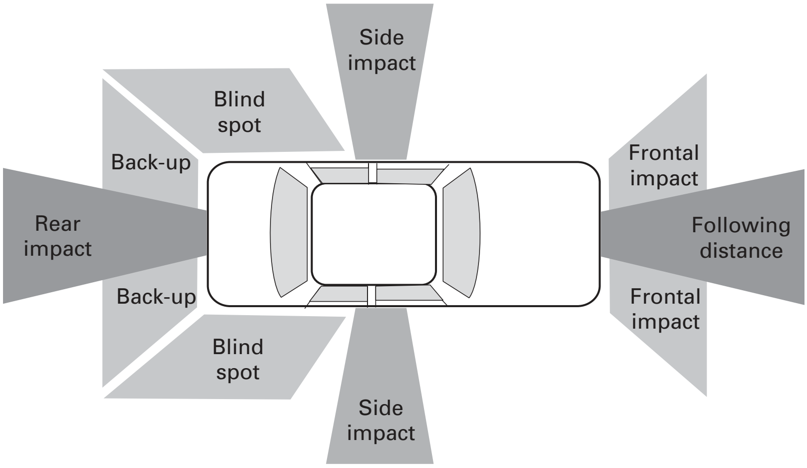

Figure 1.2 Automotive radar

number of lives. Like ground-penetrating radar, automotive radar is developed for its imaging properties and is not used for communications. Additionally, automotive radar is the only form of UWB that presently uses the 24 GHz band. It does not use the 3.1–10.6 GHz range that is employed by HDR, LDR and GPR applications.

1.2 Next-generation HDR applications

The HDR applications described in the prior section are based upon a very simple view of the world. Conceptually, devices were connected by a wire. First-generation applications are based upon a logic wherein the wire is removed and is replaced with a radio in order to improve the aesthetics of an installation and to prevent customers from having to keep up with the wires for their mobile devices. In the second generation applications, connectivity is all about connecting mobile devices while making the process somewhere between simple and invisible from the user's perspective.

To begin this topic, consider a trend that has been evident in computers for the last 50 years or so. Computational platforms are becoming smaller, cheaper and more capable with every generation. So far, the industry has moved through the mainframe, mini-computer, desktop and laptop generations. The next generation, which is emerging now, is the handheld computer–phone. It will be a converged product that includes the intelligence of a computer, communications capability of a phone and the entertainment capabilities of a music player, video player and gaming device all rolled into one.

Each generational step brings new capabilities to the consumer that were not possible before. In the handheld generation, the new capability is full mobility. Consumers will soon be able to carry huge volumes of data, programming and entertainment content with them wherever they go.

Look at the same picture from another perspective. The existing generations of PCs, set-top boxes and game consoles have been bound together by the Internet into a very large communications web. As the handheld computer platform emerges, it needs to be bound into that network as well to perform the functions that it wishes to execute. The cellular community imagines that this activity will take place across the wide area network (WAN). Unfortunately, the performance limits of that network make the WAN insufficient to handle the full job. For instance, transferring a large volume of high-resolution photographs, a full-length video or a state-of-the-art game program would take an extremely long time across the cellular network. Additional connectivity is required to do the job. The PAN (personal area network) is the principal path through which much of this work gets done.

To understand in greater detail about the applications that will emerge from connecting the mobile generation, one needs to get a feel for what consumers wish to do with the device. To begin, consumers wish to be able to work cooperatively with stationary devices such as the STB, media centre and the PC to archive data that is not being used immediately. Most individuals have far more photographs than they carry around in their wallets. Likewise, people will archive most photographs on a stationary device and only carry the most meaningful photos with them. In the case of content, only that material which the user is inclined to consume in the near future will be carried on the mobile platform. The remainder will be archived on stationary devices. The present wire-replacement strategy would seem to work reasonably well in this case.

Looking closer at the mobile device, it doesn't take long to realize that the user interface is pretty meagre. The keypad is inconveniently small, as is the display. Asking the consumer to enter file names and directory addresses will quickly become intolerable. To compensate for this, there are several things that can be done. First, if the mobile device is in range of a more capable stationary device, like a high-definition display or a computer, it will be desirable for the mobile device to be able to use them. In the case of the display, imagine the consumer touching the display on a phone (using near-field communications to establish the link) and then displaying photographs on the large screen. Or, imagine the same consumer sitting down at an Internet café. His hard drive is linked to the stationary PC. From his perspective, he has effectively turned the rented PC into his own system.

Both of these applications require the high-speed link and short range for which a PAN is optimized. But there is more. In many instances, there simply isn't an available keyboard to use to make transactions. Other methods to minimize the amount of interaction required between the user and the keyboard are necessary.

Imagine if a consumer were to walk up to a kiosk or an STB and wishes to download a video onto a mobile device to watch on the train into work. The movie is selected manually, but instead of entering information about credit-card numbers or file-storage locations, the transaction simply occurs. This would be an example of using intelligence tightly coupled with the communication links to make use of the handheld device convenient. While this type of capability could be implemented at an application level, it would also be possible to implement it as a standardized set of primitives that operates at a lower level.

With this type of approach a number of new applications become convenient. Imagine a mobile device subscribing to a music service. The mobile device is placed next to a PC each morning and its music is updated. The device is then carried out into the car where the music is then copied onto a car radio. If a song is skipped or repeated, these actions are recorded and transferred back to the mobile, the PC and ultimately to the music service, where a preference profile is developed