CHAPTER 1

Introduction

The idea of sending information over radio waves (i.e., *wireless communication*) is over a hundred years old. When several devices with radio transceivers share a portion of the radio spectrum to send information to each other, we say that we have a *wireless communication network,* or simply a wireless network.

In this chapter we begin by developing a three-layered view of wireless networks. We delineate the subject matter of this book—that is, wireless networking—as dealing with the problem of resource allocation when several devices share a portion of the RF spectrum allocated to them. Next, we provide a taxonomy of current wireless networks. The material in the book is organized along this taxonomy. Then, in this chapter, we identify the common basic technical elements that underlie any wireless network as being (1) physical wireless communication; (2) neighbor discovery, association, and topology formation; and (3) transmission scheduling.

Finally, we provide an overview of the contents of the remaining nine chapters of the book.

1.1 Networking as Resource Allocation

Following our viewpoint in [89] we view wireline and wireless communication networks in terms of the three-layered model shown in Figure 1.1. Networks carry the flows of information between distributed applications such as telephony, teleconferencing, media-sharing, World Wide Web access, e-commerce, and so on. The points at which distributed information applications connect to the generators and absorbers of information flows can be viewed as sources and sinks of traffic (see Figure 1.1). Examples of traffic sources are microphones in telephony devices, video cameras, and data, voice, or video files (stored on a computer disk) that are being transmitted to another location. Examples of traffic sinks are telephony loudspeakers, television monitors, or computer storage devices.

As shown in Figure 1.1, the sources and sinks of information and the distributed applications connect to the communication network via common *information services*. The information services layer comprises all the hardware and software required to facilitate the necessary transport services, and to attach the sources or sinks to the wireless network; for example, voice coding, packet buffering and playout, and voice decoding, for packet telephony; or similar

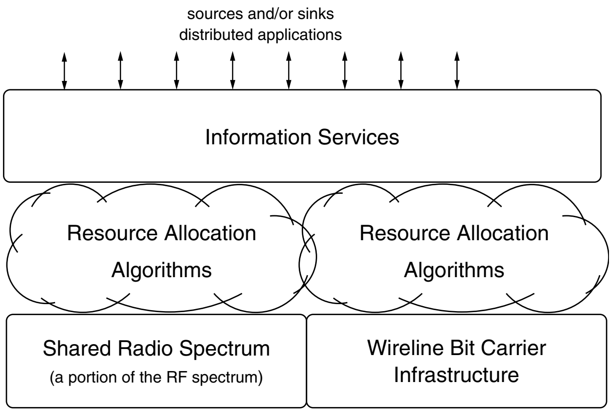

Figure 1.1 A conceptual view of distributed applications utilizing wireline and wireless networks. Wireless networking is concerned with algorithms for resource allocation between devices sharing a portion of the radio spectrum. On the other hand, in wireline networks the resource allocation algorithms are concerned with sharing the fixed resources of a bit transport infrastructure.

facilities for video telephony or for streaming video playout; or mail preparation and forwarding software for electronic mail; or a browser for the World Wide Web.

We turn now to the bottom layer in Figure 1.1. In wireline networks the information to be transported between the endpoints of applications is carried over a static bit-carrier infrastructure. These networks typically comprise high-quality digital transmission systems over copper or optical media. Once such links are properly designed and configured, they can be viewed as "bit pipes," each with a certain bit rate, and usually a very small bit error rate. The bit carrier infrastructure can be dynamically reconfigured on the basis of traffic demands, and such actions are a part of the cloud labeled "resource allocation algorithms" in the figure.

The left side of the bottom layer in Figure 1.1 corresponds to wireless networks. Typically, each wireless network system is constrained to operate in some portion of the RF spectrum. For example, a cellular telephony system may be assigned 5 MHz of spectrum in the 900 MHz band. Information bits are transported between devices in the wireless network by means of some physical wireless communication technique (i.e., a PHY layer technique, in terms of the ISO-OSI model) operating in the portion of the RF spectrum that is assigned to the network. It is well known, however, that unguided RF communication between mobile wireless devices poses challenging problems. Unlike wireline communication, or even point-to-point, high-power microwave links between dish antennas mounted on tall towers, digital wireless communication between mobile devices has to deal with a variety of time-varying channel impairments such as obstructions by steel and concrete buildings, absorption in partition walls or in foliage, and interference between copies of a signal that traverse multiple paths. In order to combat these problems, it is imperative that in a mobile or ad hoc wireless network the PHY layer should be adaptable. In fact, in some systems multiple modulation schemes are available, and each of these may have variable parameters such as the error control codes, and the transmitter powers. Hence, unlike a wired communication network, where we can view networking as being concerned with the problems of resource sharing over a static bit carrier infrastructure, in wireless networking, the resource allocation mechanisms would include these adaptations of the PHY layer. Thus, in Figure 1.1 we have actually "absorbed" the physical wireless communication mechanisms into the resource allocation layer. Hence, we can define our view of *wireless networking* as being concerned with all the mechanisms, procedures, or algorithms for efficient sharing of a portion of the radio spectrum so that all the instances of communication between the various devices obtain their desired *quality of service (QoS)*.

1.2 A Taxonomy of Current Practice

In this book, instead of pursuing an abstract, technology agnostic approach, we will develop an understanding of the various wireless networking techniques in the context of certain classes of wireless networks as they exist today. Thus we begin our treatment by taking a look at a taxonomy of the current practice of wireless networks. Figure 1.2 provides such a taxonomy. Several commonly used terms of the technology will arise as we discuss this taxonomy. These will be highlighted by the *italic* font, and their meanings will be clear from the context. Of course, the attendant engineering issues will be dealt with at length in the remainder of the book.

Fixed wireless networks include line-of-sight microwave links, which until recently were very popular for long distance transmission. Such networks basically comprise point-to-point line-of-sight digital radio links. When such links are set up, with properly aligned high gain antennas on tall masts, the links can be viewed as point-to-point bit pipes, albeit with a higher bit error rate than wired links. Thus in such fixed wireless networks no essentially new issues arise than in a network of wired links.

On the other hand the second and third categories shown in the first level of the taxonomy (i.e., access networks and ad hoc networks) involve *multiple access* where, in the same geographical region, several devices share a radio spectrum to communicate among themselves (see Figure 1.3). Currently, the most important role of wireless communications technology is in mobile access to wired networks. We can further classify such access networks into two categories: one in which resource allocation is more or less static (akin to circuit multiplexing), and the other in which the traffic is statistically multiplexed, either in a centralized manner or by distributed mechanisms.

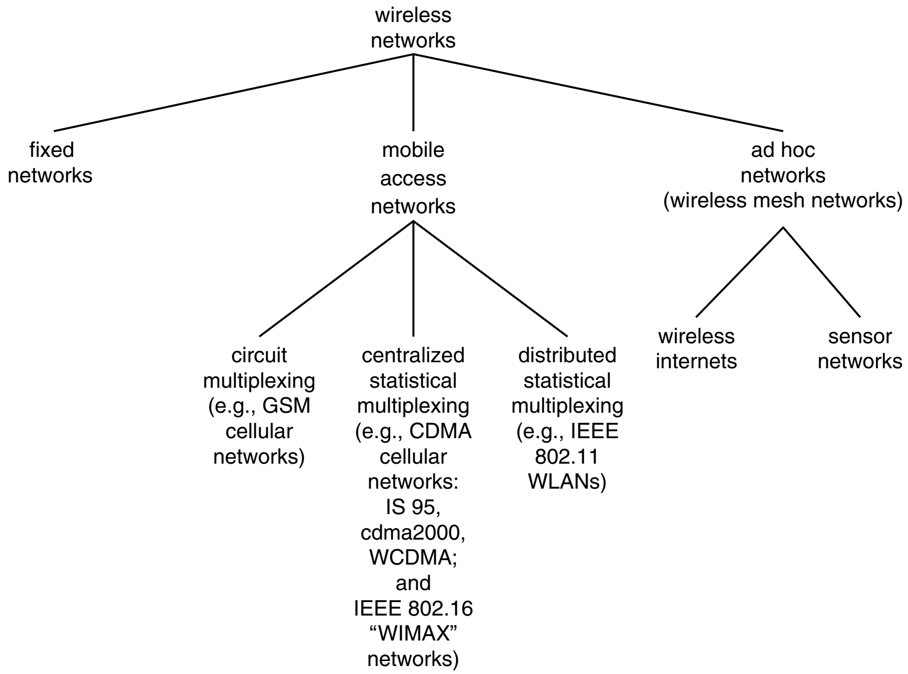

Figure 1.2 A taxonomy of wireless networks.

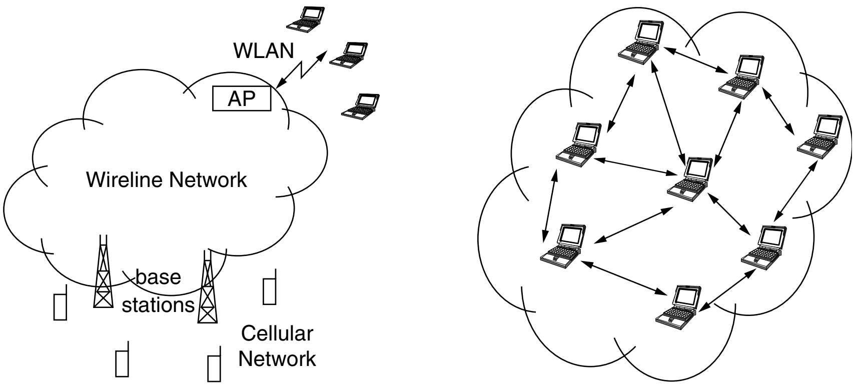

Figure 1.3 The left panel shows some access networks (a cellular telephony network, and a wireless local area network (WLAN), where the access is via an AP (access point)), and the right panel shows a mesh wireless network of portable computers.

Cellular wireless networks were introduced in the early 1980s as a technology for providing access to the wired phone network to mobile users. The network coverage area is partitioned into regions (with diameters ranging from hundreds of meters to a few kilometers) called *cells*, hence the term “cellular.” In each cell there is a *base station* (BS), which is connected to the wired network, and through which the mobile devices in the cell communicate over a one hop wireless link. The cellular systems that have the most widespread deployment are the ones that share the available spectrum using *frequency division multiplexed time division multiple access* (FDM-TDMA) technology. Among such systems by far the most commercially successful has been the GSM system, developed by a European consortium. The available spectrum is first partitioned into a contiguous *up-link* band and another contiguous *down-link* band. Each of these bands is statically or dynamically partitioned into reuse subbands, with each cell being allocated such a subband (this is the FDM aspect). The partitioning of the up-link and down-link bands is done in a paired manner so that each cell is actually assigned a pair of subbands. Each subband is further partitioned into channels or carriers (also an FDM aspect), each of which is digitally modulated and then slotted in such a way that a channel can carry up to a fixed number of calls (e.g., 8 calls) in a TDM fashion. Each arriving call request in a cell is then assigned a slot in one of the carriers in that cell; of course, a pair of slots is assigned in paired up-link and down-link channels in that cell. Thus, since each call is assigned dedicated resources, the system is said to be *circuit multiplexed*, just like the wireline phone network. These are *narrowband systems* (i.e., users’ bit streams occupy frequency bands just sufficient to carry them), and the radio links operate at a high signal-to-interference-plus-noise-ratio (SINR), and hence careful *frequency planning* (i.e., partitioning of the spectrum into reuse subbands, and allocation of the subbands to the cells) is needed to avoid *cochannel interference*. The need for allocation of frequency bands over the network coverage area (perhaps even dynamic allocation over a slow timescale), and the grant and release of individual channels as individual calls arrive and complete, requires the control of such systems to be highly centralized. Note that *call admission control*, that is, call blocking, is a natural requirement in an FDM-TDMA system, since the resources are partitioned and each connection is assigned one resource unit.

Another cellular technology that has developed over the past 10 to 15 years is the one based on *code division multiple access* (CDMA). In these networks, the entire available spectrum is reused in every cell. These are *broadband* systems, which means that each user’s bit stream (a few kilobits per second) occupies the entire available radio spectrum (a few megahertz). This is done by *spreading* each user’s signal over the entire spectrum by multiplying it by a *pseudorandom sequence*, which is allocated to the user. This makes each user’s signal appear like noise to other users. The knowledge of the spreading sequences permits the receivers to separate the users’ signals, by means of *correlation receivers*. Although no frequency planning is required for CDMA systems, the performance is *interference limited* as every transmitted signal is potentially an interferer for every other signal. Thus at any point of time there is an allocation of powers to all the transmitters sharing the spectrum, such that their desired receivers can decode their transmissions, in the presence of all the cross interferences. These desired power levels need to be set depending on the locations of the users, and the consequent channel conditions between the users and the base stations, and need to be dynamically controlled as users move about and channel conditions change. Hence tight control of transmitter power levels is necessary. Further, of course, the allocation of spreading codes, and management of movement between cells needs to done. We note that, unlike the FDM-TDMA system described earlier, there is no dedicated allocation of resources (frequency and time-slot) to each call. Indeed, during periods when a call is inactive no radio resources are utilized, and the interference to other calls is reduced. Thus, we can say that the traffic is *statistically multiplexed*. If there are several calls in the system, each needing certain *quality of service (QoS)* (bit rate, maximum bit error rate), then the number of calls in the system needs to be controlled so that the probability of QoS violation of the calls is kept small. This requires call admission control, which is an essential mechanism in CDMA systems, in order that QoS objectives can be achieved. Evidently, these are all centrally coordinated activities, and hence even CDMA cellular systems depend on central intelligence that resides in the *base station controllers* (BSCs).

Until recently, cellular networks were driven primarily by the needs of circuit multiplexed voice telephony; on demand, a mobile phone user is provided a wireless digital communication channel on which is carried compressed telephone quality (though not "toll" quality) speech. Earlier, we have described two technologies for *second generation (2G)* cellular wireless telephony. Recently, with the growing need for mobile Internet access, there have been efforts to provide packetized data access on these networks as well. In the FDM-TDMA systems, low bit rate data can be carried on the digital channel assigned to a user. As is always the case in circuit multiplexed networks, flexibility in the allocation of bandwidth is limited to assigning multiple channels to each user. Such an approach is followed in the GSM-GPRS (General Packet Radio Service) system, where, by combining multiple TDM slots on an FDM carrier, shared packet switched access is provided to mobile users. A further evolution is the EDGE (Enhanced Data rates for GSM Evolution) system, where, in addition to combining TDM slots, higher order modulation schemes, with adaptive modulation, are utilized to obtain shared packet switched links with speeds up to 474 Kbps. These two systems often are viewed, respectively, as 2.5G and 2.75G evolutions of the GSM system. These are data evolutions of an intrinsically circuit switched system that was developed for mobile telephony. On the other hand there is considerable flexibility in CDMA systems where there is no dedicated allocation of resources (spectrum or power). In fact, both voice and data can be carried in the packet mode, with the user bit rate, the amount of spreading, and the allocated power changing on a packet-by-packet basis. This is the approach taken for the *third generation (3G)* cellular systems, which are based entirely on CDMA technology, and are meant to carry *multimedia traffic* (i.e., store and forward data, packetized telephony, interactive video, and streaming video). The most widely adopted standard for 3G systems is WCDMA (wideband CDMA), which was created by the 3G Partnership Project (3GPP), a consortium of standardization organizations from the United States, Europe, China, Japan, and Korea.

Cellular networks were developed with the primary objective of providing wireless access for mobile users. With the growth of the Internet as the de facto network for information dissemination, access to the Internet has become an increasingly important requirement in most countries. In large congested cities, and in developing countries without a good wireline infrastructure, *fixed wireless access* to the Internet is seen as a significant market. It is with such an application in mind that the IEEE 802.16 standards were developed, and are known in the industry as WiMAX. The major technical advance in WiMAX is in the adoption of several high performance physical layer (PHY) technologies to provide several tens of Mbps between a base station (BS) and fixed subscriber stations (SS) over distances of several kilometers. The PHY technologies that have been utilized are *orthogonal frequency division multiple access (OFDMA)* and multiple antennas at the transmitters and the receivers. The latter are commonly referred to as *MIMO (multiple-input-multiple-output)* systems. In an OFDMA system, several subchannels are statically defined in the system bandwidth, and these subchannels are digitally modulated. In order to permit up-link and down-link transmissions, time is divided into frames and each frame is further partitioned into an up-link and a down-link part (this is called *time division duplexing (TDD)*). The BS allocates time on the various subchannels to various down-link flows in the down-link part of the frame and, based on SS requests, in the up-link part of the frame. This kind of TDD MAC structure has been used in several earlier systems; for example, satellite networks involving very small aperture satellite terminals (VSATs), and even in wireline systems such as those used for the transmission of digital data over cable television networks. WiMAX specifications now have been extended to include broadband access to mobile users.

We now discuss the third class of networks in the mobile access category in the first level of the taxonomy shown in Figure 1.2—distributed packet scheduling. Cellular networks have emerged from centrally managed point-to-point radio links, but another class of wireless networks has emerged from the idea of *random access,* whose prototypical example is the Aloha network. Spurred by advances in digital communication over radio channels, random access networks can now support bit rates close to desktop wired Ethernet access. Hence random access wireless networks are now rapidly proliferating as the technology of choice for wireless Internet access with limited mobility. The most important standards for such applications are the ones in the IEEE 802.11 series. Networks based on this standard now support physical transmission speeds from a few Mbps (over 100s of meters) up to 100 Mbps (over a few meters). The spectrum is shared in a *statistical TDMA* fashion (as opposed to slotted TDMA, as discussed, earlier, in the context of first generation FDM-TDMA systems). Nodes contend for the channel, and possibly collide. In the event of a collision, the colliding nodes *back* *off* for independently sampled random time durations, and then reattempt. When a node is able to acquire the channel, it can send at the highest of the standard bit rates that can be decoded, given the channel condition between it and its receiver. This technology is predominantly deployed for creating *wireless local area networks* (WLANs) in campuses and enterprise buildings, thus basically providing a one hop untethered access to a building's Ethernet network. In the latest enhancements to the IEEE 802.11 standards, MIMO-OFDM physical layer technologies are being employed in order to obtain up to 100 Mbps transmission speeds in indoor environments.

With the widespread deployment of IEEE 802.11 WLANs in buildings, and even public spaces (such as shopping malls and airports), an emerging possibility is that of carrying interactive voice and streaming video traffic over these networks. The emerging concept of fourth-generation wireless access networks envisions mobile devices that can support multiple technologies for physical digital radio communication, along with the resource management algorithms that would permit a device to seamlessly move between 3G cellular networks, IEEE 802.16 access networks and IEEE 802.11 WLANs, while supporting a variety of packet mode services, each with its own QoS requirements.

With reference to the taxonomy in Figure 1.2, we now turn to the category labeled "ad hoc networks" or "wireless mesh networks." Wireless access networks provide mobile devices with one-hop wireless access to a wired network. Thus, in such networks, in the path between two user devices there is only one or at most two wireless links. On the other hand a wireless ad hoc network comprises several devices arbitrarily located in a space (e.g., a line segment, or a two-dimensional field). Each device is equipped with a radio transceiver, all of which typically share the same radio frequency band. In this situation, the problem is to communicate between the various devices. Nodes need to discover neighbors in order to form a topology, good paths need to be found, and then some form of time scheduling of transmissions needs to be employed in order to send packets between the devices. Packets going from one node to another may need to be forwarded by other nodes. Thus, these are *multihop* wireless packet radio networks, and they have been studied as such over several years. Interest in such networks has again been revived in the context of *multihop wireless internets* and *wireless sensor networks*. We discuss these briefly in the following two paragraphs.

In some situations it becomes necessary for several mobile devices (such as portable computers) to organize themselves into a multihop wireless packet network. Such a situation could arise in the aftermath of a major natural disaster such as an earthquake, when emergency management teams need to coordinate their activities and all the wired infrastructure has been damaged. Notice that the kind of communication that such a network would be required to support would be similar to what is carried by regular public networks; that is, point-to-point store and forward traffic such as electronic mails and file transfers, and low bit rate voice and video communication. Thus, we can call such a network a *multihop wireless internet*. In general, such a network could attach at some point to the wired Internet.

Whereas multihop wireless internets have the service objective of supporting instances of point-to-point communication, an ad hoc wireless sensor network has a global objective. The nodes in such a network are miniature devices, each of which carries a microprocessor (with an energy efficient operating system); one or more sensors (e.g., light, acoustic, or chemical sensors); a low power, low bit rate digital radio transceiver; and a small battery. Each sensor monitors its environment and the objective of the network is to deliver some global information or an inference about the environment to an operator who could be located at the periphery of the network, or could be remotely connected to the sensor network. An example is the deployment of such a network in the border areas of a country to monitor intrusions. Another example is to equip a large building with a sensor network comprising devices with strain sensors in order to monitor the building's structural integrity after an earthquake. Yet another example is the use of such sensor networks in monitoring and control systems such as those for the environment of an office building or hotel, or a large chemical factory.

1.3 Technical Elements

In the previous section we provided an overview of the current practice of wireless networks. We organized our presentation around a taxonomy of wireless networks shown in Figure 1.2. Although the technologies that we discussed may appear to be disparate, there are certain common technical elements that constitute these wireless networks. The efficient realization of these elements constitutes the area of wireless networking.

The following is an enumeration and preliminary discussion of the technical elements.

1. *Transport of the users' bits over the shared radio spectrum.* There is, of course, no communication network unless bits can be transported between users. Digital communication over mobile wireless links has evolved rapidly over the past two decades. Several approaches are now available, with various tradeoffs and areas of applicability. Even in a given system, the digital communication mechanisms can be adaptive. First, for a given digital modulation scheme the parameters can be adapted (e.g., the transmit power, or the amount of error protection), and, second, sophisticated physical layers actually permit the modulation itself to be changed even at the packet or burst timescale (e.g., if the channel quality improves during a call then a higher order modulation can be used, thus helping in store and forward applications that can utilize such time varying capacity). This adaptivity is very useful in the mobile access situation where the channels and interference levels are rapidly changing.

2. *Neighbor discovery, association and topology formation, routing*. Except in the case of fixed wireless networks, we typically do not "force" the formation of specific links in a wireless network. For example, in an access network each mobile device could be in the vicinity of more than one BS or *access point* (AP). To simplify our writing, we will refer to a BS or an AP as an access device. It is a nontrivial issue as to which access device a mobile device connects through. First, each mobile needs to determine which access devices are in its vicinity, and through which it can potentially communicate. Then each mobile should associate with an access device such that certain overall communication objectives are satisfied. For example, if a mobile is in the vicinity of two BSs and needs certain quality of service, then its assignment to only a particular one of the two BSs may result in satisfaction of the new requirement, and all the existing ones.

   In the case of an access network the problem of routing is trivial; a mobile associates with a BS and all its packets need to be routed through that BS. On the other hand, in an ad hoc network, after the associations are made and a topology is determined, good routes need to be determined. A mobile would have several neighbors in the discovered topology. In order to send a packet to a destination, an appropriate neighbor would need to be chosen, and this neighbor would further need to forward the packet toward the destination. The choice of the route would depend on factors such as the bit rate achievable on the hops of the route, the number of hops on the route, the congestion along the route, and the residual battery energies in devices along the route.

   We note that association and topology formation is a procedure whose timescale will depend on how rapidly the relative locations of the network nodes is changing. However, one would typically not expect to associate and reassociate a mobile device, form a new topology, or recalculate routing at the packet timescale.

   If mobility is low, for example in wireless LANs and static sensor networks, one could consider each fixed association, topology, and routing, and compute the performance measures at the user level. Note that this step requires a scheduling mechanism, discussed as the next element. Then that association, topology, and routing would be chosen that optimizes, in some sense, the performance measures. In the formulation of such a problem, first we need to identify one or more performance objectives (e.g., the sum of the user utilities for the transfer rates they get). Then we need to specify whether we seek a cooperative optimum (e.g., the network operator might seek the global objective of maximizing revenue) or a noncooperative equilibrium. The latter might model the more practical situation, since users would tend to act selfishly, attempting to maximize their performance while reducing their costs. Finally, whatever the solution of the problem, we need an algorithm (centralized or distributed) to compute it online.

If the mobility is high, however, the association problem would need to be dynamically solved as the devices move around. Such a problem may be relatively simple in a wireless access network, and, indeed, necessary since cellular networks are supposed to handle high mobility users. On the other hand such a problem would be hard for a general mesh network; highly mobile wireless mesh networks, however, are not expected to be "high performance" networks.

3. *Transmission scheduling.* Given an association, a topology, and the routes, and the various possibilities of adaptation at the physical layer, the problem is to schedule transmissions between the various devices so that the users' QoS objectives are met. In its most general form, the schedule dynamically needs to determine which transceivers should transmit, how much they should transmit, and which physical layer (including its parameters, e.g., transmit power) should be used between each transceiver pair. Such a scheduler would be said to be *cross-layer* if it took into account state information at multiple layers; for example, channel state information, as well as higher layer state information, such as link buffer queue lengths. Note that a scheduling mechanism will determine the schedulable region for the network; that is, the set of user flow rates of each type that can be carried so that each flow's QoS is met.

In general, these three technical elements are interdependent and the most general approach would be to jointly optimize them. For example, in a mobile Internet access network the mobile devices are associated with base stations. The channel qualities between the base stations and the mobile devices determine the bit rates that can be sustained, the transmission powers required, and transmission schedule required to achieve the desired QoS for the various connections. Thus, the overall problem involves a joint optimization of the association, the physical layer parameters, and the transmission schedule.

In addition to the preceding elements that provide the basic communication functionality, some wireless networks require other functional elements that could be key to the networks' overall utility. The following are two important ones, which are of special relevance to ad hoc wireless sensor networks.

- *Location determination.* In an ad hoc wireless sensor network the nodes make measurements on their environment, and then these measurements are used to carry out some global computation. Often, in this process it becomes necessary to determine from which location a measurement came. Sensor network nodes may be too small (in terms of size and available energy) to carry a GPS (global positioning system) receiver. Some applications may require the nodes to be placed indoors, where GPS signals may not penetrate. Hence GPS-free techniques for location determination become important. Even in cellular networks, there is a requirement in some countries that, if needed, a mobile device should be geographically locatable. Such a feature can be used to locate someone who is stranded in an emergency situation and is unaware of the exact location.

- *Distributed computation.* This issue is specific to wireless sensor networks. It may be necessary to compute some function of the values measured by sensors (e.g., the maximum or the average). Such a computation may involve some statistical signal processing functions such as data compression, detection, or estimation. Since these networks operate with very simple digital radios and processors, and have only small amounts of battery energy, the design of efficient self-organizing wireless ad hoc networks and distributed computation schemes on them is an important emerging area. In such networks there is communication delay and also data loss; hence existing algorithms may need to be redesigned to be robust to information delay and loss.

1.4 Summary and Our Way Forward

We began with a discussion of our view of networking as resource allocation. Figure 1.1 summarizes our view. This was followed by a taxonomy of current wireless practice in Section 1.2. Next, the common technical elements that underlie the apparently disparate technologies were abstracted and discussed in Section 1.3.

Before we can proceed to the core topic of this book—resource allocation to meet specified QoS objectives—we will need to understand basic models of, and notions associated with, the wireless channel. Along with this, the important techniques employed in digital communication will be covered in Chapter 2. These concepts will be like the building blocks in terms of which our resource allocation problems will be posed, and answers sought. Essentially, in Chapter 2, our discussion will be confined to the so-called PHY layer.

However, before commencing our study of resource allocation problems, we will pause and take a look at the applications that usually are carried on communication networks. Our objectives will be to understand the characteristics of the bit streams or the packet streams generated by various applications (the top layer of Figure 1.1), as well as the performance requirements the streams demand. This will be the topic of Chapter 3.

Beginning with Chapter 4, we will consider, one by one, the different wireless networks shown at the second level of our taxonomy (Figure 1.2). In each case, the emphasis will be on posing and solving resource allocation problems specific to that type of network. In Chapter 4, narrowband cellular systems will be studied. Power, bandwidth, and time are the resources here, and the principal objective is to maintain the signal-to-interference ratio (SIR) at an adequately high level. Our discussion will give rise to several important concepts, including frequency reuse, sectorization, spectrum efficiency, handover blocking, and channel reservation.

Continuing with cellular access networks, we will focus on CDMA systems in Chapter 5. The distinguishing feature here is that of universal frequency reuse. As before, the main theme is to assign power so as to ensure that the signal-to-interference-plus-noise ratio (SINR) is adequately high. We will see how the notions of other-cell interference, power control, and hard as well as soft handover arise in this context.

In Chapter 6, we will turn to OFDMA-TDMA systems, where power, frequency, and time constitute the basic resources to be allocated. Unlike FDM-TDMA and CDMA systems, where to each flow a fixed bit rate is assigned, in OFDMA-TDMA systems, the resources are assigned dynamically over time, depending on time varying user requirements and channel conditions. Generally speaking, the objective is to maximize the aggregate bit capacity of a time-varying channel, subject to a constraint on the average power. The important notion of the water-filling power allocation will emerge from our discussions.

In Chapter 7, the focus shifts to random-access systems and, in particular, IEEE 802.11 WLANs. The principal resource here is channel time, and distributed control of access to the channel is of interest. In a system of n colocated WLAN nodes, what is the saturation throughput that each can achieve? We will analyze this important question. Various issues pertaining to the transport of voice and data traffic over WLANs will also be discussed.

Continuing with our discussion of the various networks according to our taxonomy, we will study multihop wireless mesh networks in Chapters 8 and 9. In Chapter 8, we assume that a wireless mesh network is given. On this network, we will address the fundamental question of optimal routing and link scheduling of packet flows for a given set of source-destination pairs. Again, the basic resources here are bandwidth, time, and power, and it is of interest to know which nodes should get access to the bandwidth at what times so as to achieve the objective of maximizing throughput. Our analysis will lead to the notions of optimal scheduling and routing. We first consider open loop flows. The flow rates may be given or they may be unknown. For the latter case, the important maximum weight scheduling is described in detail. We also consider routing and scheduling for elastic flows so as to maximize a network utility function.

In Chapter 9, we will address some fundamental questions that arise in the context of wireless mesh networks. First, we ask, what is the minimum power level that nodes can use while ensuring that the network of nodes remains connected? After a suitable definition of the network capacity we also obtain the capacity of arbitrary and random networks. Although asymptotic analyses provide interesting insights, wherever possible, we also consider finite networks.

Finally, in Chapter 10, we will turn to wireless sensor networks. Apart from power and bandwidth, each sensor itself can be considered as a resource now.

A variety of new problems arise; for example, if sensors are deployed in a random manner over a given area, how many of them are required so that every point in the area is sensed by not less than *k* sensors? As mentioned before in Section 1.3, wireless sensor networks often have special needs; for example, localization and distributed computation. Resource allocation problems for meeting such objectives will also be discussed.