




Wireless Communications

ANDREA GOLDSMITH
Stanford University



CAMBRIDGE
UNIVERSITY PRESS

UMass Dartmouth

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
40 West 20th Street, New York, NY 10011-4211, USA

www.cambridge.org

Information on this title: www.cambridge.org/9780521837163

© Cambridge University Press 2005

This publication is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2005

Printed in the United States of America

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication data

Goldsmith, Andrea

Wireless communications / Andrea Goldsmith.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-521-83716-3

1. Wireless communication systems. I. Title.

TK5103.2.G65 2005
621.382 – dc22

2005047075

ISBN-13 978-0-521-83716-3 hardback
ISBN-10 0-521-83716-2 hardback

Cambridge University Press has no responsibility for
the persistence or accuracy of URLs for external or
third-party Internet Web sites referred to in this publication
and does not guarantee that any content on such
Web sites is, or will remain, accurate or appropriate.

from the Isle of Wight to a tugboat 18 miles away, and radio communications was born. Radio technology advanced rapidly to enable transmissions over larger distances with better quality, less power, and smaller, cheaper devices, thereby enabling public and private radio communications, television, and wireless networking.

Early radio systems transmitted analog signals. Today most radio systems transmit digital signals composed of binary bits, where the bits are obtained directly from a data signal or by digitizing an analog signal. A digital radio can transmit a continuous bit stream or it can group the bits into packets. The latter type of radio is called a *packet radio* and is often characterized by bursty transmissions: the radio is idle except when it transmits a packet, although it may transmit packets continuously. The first network based on packet radio, ALOHANET, was developed at the University of Hawaii in 1971. This network enabled computer sites at seven campuses spread out over four islands to communicate with a central computer on Oahu via radio transmission. The network architecture used a star topology with the central computer at its hub. Any two computers could establish a bi-directional communications link between them by going through the central hub. ALOHANET incorporated the first set of protocols for channel access and routing in packet radio systems, and many of the underlying principles in these protocols are still in use today. The U.S. military was extremely interested in this combination of packet data and broadcast radio. Throughout the 1970s and early 1980s the Defense Advanced Research Projects Agency (DARPA) invested significant resources to develop networks using packet radios for tactical communications in the battlefield. The nodes in these ad hoc wireless networks had the ability to self-configure (or reconfigure) into a network without the aid of any established infrastructure. DARPA's investment in ad hoc networks peaked in the mid 1980s, but the resulting systems fell far short of expectations in terms of speed and performance. These networks continue to be developed for military use. Packet radio networks also found commercial application in supporting wide area wireless data services. These services, first introduced in the early 1990s, enabled wireless data access (including email, file transfer, and Web browsing) at fairly low speeds, on the order of 20 kbps. No strong market for these wide area wireless data services ever really materialized, due mainly to their low data rates, high cost, and lack of "killer applications". These services mostly disappeared in the 1990s, supplanted by the wireless data capabilities of cellular telephones and wireless local area networks (WLANs).

The introduction of wired Ethernet technology in the 1970s steered many commercial companies away from radio-based networking. Ethernet's 10-Mbps data rate far exceeded anything available using radio, and companies did not mind running cables within and between their facilities to take advantage of these high rates. In 1985 the Federal Communications Commission (FCC) enabled the commercial development of wireless LANs by authorizing the public use of the Industrial, Scientific, and Medical (ISM) frequency bands for wireless LAN products. The ISM band was attractive to wireless LAN vendors because they did not need to obtain an FCC license to operate in this band. However, the wireless LAN systems were not allowed to interfere with the primary ISM band users, which forced them to use a low power profile and an inefficient signaling scheme. Moreover, the interference from primary users within this frequency band was quite high. As a result, these initial wireless LANs had very poor performance in terms of data rates and coverage. This poor performance – coupled with concerns about security, lack of standardization, and high cost

gradually evolved to support data services such as email, Internet access, and short messaging. Unfortunately, the great market potential for cellular phones led to a proliferation of second-generation cellular standards: three different standards in the United States alone, other standards in Europe and Japan, and all incompatible. The fact that different cities have different incompatible standards makes roaming throughout the United States and the world with only one cellular phone standard impossible. Moreover, some countries have initiated service for third-generation systems, for which there are also multiple incompatible standards. As a result of this proliferation of standards, many cellular phones today are multimode: they incorporate multiple digital standards to facilitate nationwide and worldwide roaming and possibly the first-generation analog standard as well, since only this standard provides universal coverage throughout the United States.

Satellite systems are typically characterized by the height of the satellite orbit: low-earth orbit (LEOs at roughly 2000 km altitude), medium-earth orbit (MEOs, 9000 km), or geosynchronous orbit (GEOs, 40,000 km). The geosynchronous orbits are seen as stationary from the earth, whereas satellites with other orbits have their coverage area change over time. The concept of using geosynchronous satellites for communications was first suggested by the science-fiction writer Arthur C. Clarke in 1945. However, the first deployed satellites – the Soviet Union's *Sputnik* in 1957 and the NASA/Bell Laboratories' *Echo-I* in 1960 – were not geosynchronous owing to the difficulty of lifting a satellite into such a high orbit. The first GEO satellite was launched by Hughes and NASA in 1963; GEOs then dominated both commercial and government satellite systems for several decades.

Geosynchronous satellites have large coverage areas, so fewer satellites (and dollars) are necessary to provide wide area or global coverage. However, it takes a great deal of power to reach the satellite, and the propagation delay is typically too large for delay-constrained applications like voice. These disadvantages caused a shift in the 1990s toward lower-orbit satellites [2; 3]. The goal was to provide voice and data service competitive with cellular systems. However, the satellite mobile terminals were much bigger, consumed much more power, and cost much more than contemporary cellular phones, which limited their appeal. The most compelling feature of these systems is their ubiquitous worldwide coverage, especially in remote areas or third-world countries with no landline or cellular system infrastructure. Unfortunately, such places do not typically have large demand or the resources to pay for satellite service either. As cellular systems became more widespread, they took away most revenue that LEO systems might have generated in populated areas. With no real market left, most LEO satellite systems went out of business.

A natural area for satellite systems is broadcast entertainment. Direct broadcast satellites operate in the 12-GHz frequency band. These systems offer hundreds of TV channels and are major competitors to cable. Satellite-delivered digital radio has also become popular. These systems, operating in both Europe and the United States, offer digital audio broadcasts at near-CD quality.

1.2 Wireless Vision

The vision of wireless communications supporting information exchange between people or devices is the communications frontier of the next few decades, and much of it already

exists in the world. The world is top, laptop, from the electronic devices connecting homes, cars, emergency, people, continents, who miss, will enable the world, application, stress and, icals and, process a, control l, detection, counterter, devices, systems, trial auto.

The v, what, ex, topic into, clude vo, mation s, control, data syst, building, commun, the indu, being off, cations h, (around, of arou, ticeable, rates (l, must be, have hig.

¹ Wired t, voice o

cess, and short mes-
led to a proliferation
e United States alone,
ct that different cities
United States and the
, some countries have
so multiple incompati-
llular phones today are
ionwide and worldwide
ince only this standard

satellite orbit: low-earth
s, 9000 km), or geosyn-
seen as stationary from
a change over time. The
as first suggested by the
deployed satellites – the
Echo-I in 1960 – were not
ch a high orbit. The first
hen dominated both com-

satellites (and dollars) are
akes a great deal of power
arge for delay-constrained
1990s toward lower-orbit
competitive with cellular
ger, consumed much more
which limited their appeal.
s worldwide coverage, es-
ne or cellular system infra-
ge demand or the resources
more widespread, they took
populated areas. With no real

t. Direct broadcast satellites
undreds of TV channels and
io has also become popular.
offer digital audio broadcasts

on exchange between people
ades, and much of it already

exists in some form. This vision will allow multimedia communication from anywhere in the world using a small handheld device or laptop. Wireless networks will connect palm-top, laptop, and desktop computers anywhere within an office building or campus, as well as from the corner cafe. In the home these networks will enable a new class of intelligent electronic devices that can interact with each other and with the Internet in addition to providing connectivity between computers, phones, and security/monitoring systems. Such “smart” homes can also help the elderly and disabled with assisted living, patient monitoring, and emergency response. Wireless entertainment will permeate the home and any place that people congregate. Video teleconferencing will take place between buildings that are blocks or continents apart, and these conferences can include travelers as well – from the salesperson who missed his plane connection to the CEO off sailing in the Caribbean. Wireless video will enable remote classrooms, remote training facilities, and remote hospitals anywhere in the world. Wireless sensors have an enormous range of both commercial and military applications. Commercial applications include monitoring of fire hazards, toxic waste sites, stress and strain in buildings and bridges, carbon dioxide movement, and the spread of chemicals and gasses at a disaster site. These wireless sensors self-configure into a network to process and interpret sensor measurements and then convey this information to a centralized control location. Military applications include identification and tracking of enemy targets, detection of chemical and biological attacks, support of unmanned robotic vehicles, and counterterrorism. Finally, wireless networks enable distributed control systems with remote devices, sensors, and actuators linked together via wireless communication channels. Such systems in turn enable automated highways, mobile robots, and easily reconfigurable industrial automation.

The various applications described here are all components of the wireless vision. So then what, exactly, is wireless communications? There are many ways to segment this complex topic into different applications, systems, or coverage regions [4]. Wireless applications include voice, Internet access, Web browsing, paging and short messaging, subscriber information services, file transfer, video teleconferencing, entertainment, sensing, and distributed control. Systems include cellular telephone systems, wireless LANs, wide area wireless data systems, satellite systems, and ad hoc wireless networks. Coverage regions include in-building, campus, city, regional, and global. The question of how best to characterize wireless communications along these various segments has resulted in considerable fragmentation in the industry, as evidenced by the many different wireless products, standards, and services being offered or proposed. One reason for this fragmentation is that different wireless applications have different requirements. Voice systems have relatively low data-rate requirements (around 20 kbps) and can tolerate a fairly high probability of bit error (bit error rates, or BERs, of around 10^{-3}), but the total delay must be less than about 100 ms or else it becomes noticeable to the end user.¹ On the other hand, data systems typically require much higher data rates (1–100 Mbps) and very small BERs (a BER of 10^{-8} or less, and all bits received in error must be retransmitted) but do not have a fixed delay requirement. Real-time video systems have high data-rate requirements coupled with the same delay constraints as voice systems,

¹ Wired telephones have a delay constraint of ~30 ms. Cellular phones relax this constraint to ~100 ms, and voice over the Internet relaxes the constraint even further.

spectrum sharing are very complex, and the decision of which technique is best for a given system and operating environment is never straightforward.

Efficient cellular system designs are *interference limited* – that is, the interference dominates the noise floor, since otherwise more users could be added to the system. As a result, any technique to reduce interference in cellular systems leads directly to an increase in system capacity and performance. Some methods for interference reduction in use today or proposed for future systems include cell sectorization, directional and smart antennas, multi-user detection, and dynamic resource allocation. Details of these techniques will be given in Chapter 15.

The first-generation (1G) cellular systems in the United States, called the Advance Mobile Phone Service (AMPS), used FDMA with 30-kHz FM-modulated voice channels. The FCC initially allocated 40 MHz of spectrum to this system, which was increased to 50 MHz shortly after service introduction to support more users. This total bandwidth was divided into two 25-MHz bands, one for mobile-to-base station channels and the other for base station-to-mobile channels. The FCC divided these channels into two sets that were assigned to two different service providers in each city to encourage competition. A similar system, the Total Access Communication System (TACS), emerged in Europe. AMPS was deployed worldwide in the 1980s and remains the only cellular service in some areas, including rural parts of the United States.

Many of the first-generation cellular systems in Europe were incompatible, and the Europeans quickly converged on a uniform standard for second-generation (2G) digital systems called GSM.² The GSM standard uses a combination of TDMA and slow frequency hopping with frequency-shift keying for the voice modulation. In contrast, the standards activities in the United States surrounding the second generation of digital cellular provoked a raging debate on spectrum-sharing techniques, resulting in several incompatible standards [9; 10; 11]. In particular, there are two standards in the 900-MHz cellular frequency band: IS-136,³ which uses a combination of TDMA and FDMA and phase-shift keyed modulation; and IS-95, which uses direct-sequence CDMA with phase-shift keyed modulation and coding [12; 13]. The spectrum for digital cellular in the 2-GHz PCS (personal communication system) frequency band was auctioned off, so service providers could use any standard for their purchased spectrum. The end result has been three different digital cellular standards for this frequency band: IS-136, IS-95, and the European GSM standard. The digital cellular standard in Japan is similar to IS-136 but in a different frequency band, and the GSM system in Europe is at a different frequency than the GSM systems in the United States. This proliferation of incompatible standards in the United States and internationally makes it impossible to roam between systems nationwide or globally without a multimode phone and/or multiple phones (and phone numbers).

All of the second-generation digital cellular standards have been enhanced to support high-rate packet data services [14]. GSM systems provide data rates of up to 140 kbps by aggregating all timeslots together for a single user. This enhancement is called GPRS. A

² The acronym GSM originally stood for Groupe Spéciale Mobile, the name of the European charter establishing the GSM standard. As GSM systems proliferated around the world, the underlying acronym meaning was changed to Global Systems for Mobile Communications.

³ IS-136 is the evolution of the older IS-54 standard and subsumes it.

18-GHz band, had data rates on the order of 6 Mbps. However, performance of Altair was hampered by the high cost of components and the increased path loss at 18 GHz, and Altair was discontinued within a few years of its release.

The second-generation wireless LANs in the United States operate with 83.5 MHz of spectrum in the 2.4-GHz ISM band. A wireless LAN standard for this frequency band, the IEEE 802.11b standard, was developed to avoid some of the problems with the proprietary first-generation systems. The standard specifies direct-sequence spread spectrum with data rates of around 1.6 Mbps (raw data rates of 11 Mbps) and a range of approximately 100 m. The network architecture can be either star or peer-to-peer, although the peer-to-peer feature is rarely used. Many companies developed products based on the 802.11b standard, and after slow initial growth the popularity of 802.11b wireless LANs has expanded considerably. Many laptops come with integrated 802.11b wireless LAN cards. Companies and universities have installed 802.11b base stations throughout their locations, and many coffee houses, airports, and hotels offer wireless access, often for free, to increase their appeal.

Two additional standards in the 802.11 family were developed to provide higher data rates than 802.11b. The IEEE 802.11a wireless LAN standard operates with 300 MHz of spectrum in the 5-GHz U-NII band. The 802.11a standard is based on multicarrier modulation and provides 54-Mbps data rates at a range of about 30 m. Because 802.11a has much more bandwidth and consequently many more channels than 802.11b, it can support more users at higher data rates. There was some initial concern that 802.11a systems would be significantly more expensive than 802.11b systems, but in fact they quickly became quite competitive in price. The other standard, 802.11g, has the same design and data rates as 802.11a, but it operates in the 2.4-GHz band with a range of about 50 m. Many wireless LAN cards and access points support all three standards to avoid incompatibilities.

In Europe, wireless LAN development revolves around the HIPERLAN (high-performance radio LAN) standards. The HIPERLAN/2 standard is similar to the IEEE 802.11a wireless LAN standard. In particular, it has a similar link layer design and also operates in a 5-GHz frequency band similar to the U-NII band. Hence it has the same maximum data rate of 54 Mbps and the same range of approximately 30 m as 802.11a. It differs from 802.11a in its access protocol and its built-in Quality-of-Service (QoS) support.

1.4.4 Wide Area Wireless Data Services

Wide area wireless data services provide wireless data to high-mobility users over a large coverage area. In these systems a given geographical region is serviced by base stations mounted on towers, rooftops, or mountains. The base stations can be connected to a backbone wired network or form a multihop ad hoc wireless network.

Initial wide area wireless data services had very low data rates, below 10 kbps, which gradually increased to 20 kbps. There were two main players providing this service: Motient and Bell South Mobile Data (formerly RAM Mobile Data). Metricom provided a similar service using architecture that consisted of a large network of small, inexpensive base stations with small coverage areas. The increased efficiency of the small coverage areas allowed for higher data rates in Metricom, 76 kbps, than in the other wide area wireless data systems. However, the high infrastructure cost for Metricom eventually forced it into bankruptcy, and the system was shut down. Some of the infrastructure was bought and is operating in a few areas as Ricochet.

The
layed o
of the a
tal cell
is availa
tems als
these ne
althoug

1.4.5 B

Broadb
cess poi
tive vid
high-sp
high-sp
bands v
systems
for metr
LMDS
tition an
a televis
[5, Cha
TV char
with ex
called F

Wi
[16; 17]
operati
and bet
will be
meters.
to incor
cess an
cellular
The abi
forman

1.4.6 I

Paging
tions or
with ter
city or
internat
tions ar
high co

formance of Altair was
ss at 18 GHz, and Altair
erate with 83.5 MHz of
this frequency band, the
ems with the proprietary
read spectrum with data
of approximately 100 m.
ugh the peer-to-peer fea-
the 802.11b standard, and
as expanded considerably.
Companies and universi-
, and many coffee houses,
se their appeal.
to provide higher data rates
es with 300 MHz of spec-
on multicarrier modulation
use 802.11a has much more
it can support more users at
stems would be significantly
became quite competitive in
a rates as 802.11a, but it op-
any wireless LAN cards and
ties.
e HIPERLAN (high-perfor-
similar to the IEEE 802.11a
design and also operates in a
the same maximum data rate
11a. It differs from 802.11a in
support.

gh-mobility users over a large
on is serviced by base stations
ns can be connected to a back-
work.
ata rates, below 10 kbps, which
s providing this service: Motient
Metricom provided a similar ser-
small, inexpensive base stations
small coverage areas allowed for
wide area wireless data systems.
ully forced it into bankruptcy, and
bought and is operating in a few

The cellular digital packet data (CDPD) system is a wide area wireless data service overlaid on the analog cellular telephone network. CDPD shares the FDMA voice channels of the analog systems, since many of these channels are idle owing to the growth of digital cellular. The CDPD service provides packet data transmission at rates of 19.2 kbps and is available throughout the United States. However, since newer generations of cellular systems also provide data services and at higher data rates, CDPD is mostly being replaced by these newer services. Thus, wide area wireless data services have not been very successful, although emerging systems that offer broadband access may have more appeal.

1.4.5 Broadband Wireless Access

Broadband wireless access provides high-rate wireless communications between a fixed access point and multiple terminals. These systems were initially proposed to support interactive video service to the home, but the application emphasis then shifted to providing both high-speed data access (tens of Mbps) to the Internet and the World Wide Web as well as high-speed data networks for homes and businesses. In the United States, two frequency bands were set aside for these systems: part of the 28-GHz spectrum for local distribution systems (local multipoint distribution service, LMDS) and a band in the 2-GHz spectrum for metropolitan distribution service (multichannel multipoint distribution services, MMDS). LMDS represents a quick means for new service providers to enter the already stiff competition among wireless and wireline broadband service providers [5, Chap. 2.3]. MMDS is a television and telecommunication delivery system with transmission ranges of 30–50 km [5, Chap. 11.11]. MMDS has the capability of delivering more than a hundred digital video TV channels along with telephony and access to the Internet. MMDS will compete mainly with existing cable and satellite systems. Europe is developing a standard similar to MMDS called Hiperaccess.

WiMax is an emerging broadband wireless technology based on the IEEE 802.16 standard [16; 17]. The core 802.16 specification is a standard for broadband wireless access systems operating at radio frequencies between 2 GHz and 11 GHz for non-line-of-sight operation, and between 10 GHz and 66 GHz for line-of-sight operation. Data rates of around 40 Mbps will be available for fixed users and 15 Mbps for mobile users, with a range of several kilometers. Many manufacturers of laptops and PDAs (personal digital assistants) are planning to incorporate WiMax once it becomes available to satisfy demand for constant Internet access and email exchange from any location. WiMax will compete with wireless LANs, 3G cellular services, and possibly wireline services like cable and DSL (digital subscriber line). The ability of WiMax to challenge or supplant these systems will depend on its relative performance and cost, which remain to be seen.

1.4.6 Paging Systems

Paging systems broadcast a short paging message simultaneously from many tall base stations or satellites transmitting at very high power (hundreds of watts to kilowatts). Systems with terrestrial transmitters are typically localized to a particular geographic area, such as a city or metropolitan region, while geosynchronous satellite transmitters provide national or international coverage. In both types of systems, no location management or routing functions are needed because the paging message is broadcast over the entire coverage area. The high complexity and power of the paging transmitters allows low-complexity, low-power,

pocket paging receivers capable of long usage times from small and lightweight batteries. In addition, the high transmit power allows paging signals to easily penetrate building walls. Paging service also costs less than cellular service, both for the initial device and for the monthly usage charge, although this price advantage has declined considerably in recent years as cellular prices dropped. The low cost, small and lightweight handsets, long battery life, and ability of paging devices to work almost anywhere indoors or outdoors are the main reasons for their appeal.

Early radio paging systems were analog 1-bit messages signaling a user that someone was trying to reach him or her. These systems required callback over a landline telephone to obtain the phone number of the paging party. The system evolved to allow a short digital message, including a phone number and brief text, to be sent to the paged as well. Radio paging systems were initially extremely successful, with a peak of 50 million subscribers in the United States alone. However, their popularity began to wane with the widespread penetration and competitive cost of cellular telephone systems. Eventually the competition from cellular phones forced paging systems to provide new capabilities. Some implemented "answer-back" capability (i.e., two-way communication). This required a major change in design of the pager because now it needed to transmit signals in addition to receiving them, and the transmission distance to a satellite or base station can be very large. Paging companies also teamed up with palmtop computer makers to incorporate paging functions into these devices [18]. Despite these developments, the market for paging devices has shrunk considerably, although there is still a niche market among doctors and other professionals who must be reachable anywhere.

1.4.7 Satellite Networks

Commercial satellite systems are another major component of the wireless communications infrastructure [2; 3]. Geosynchronous systems include Inmarsat and OmniTRACS. The former is geared mainly for analog voice transmission from remote locations. For example, it is commonly used by journalists to provide live reporting from war zones. The first-generation Inmarsat-A system was designed for large (1-m parabolic dish antenna) and rather expensive terminals. Newer generations of Inmarsats use digital techniques to enable smaller, less expensive terminals, about the size of a briefcase. Qualcomm's OmniTRACS provides two-way communications as well as location positioning. The system is used primarily for alphanumeric messaging and location tracking of trucking fleets. There are several major difficulties in providing voice and data services over geosynchronous satellites. It takes a great deal of power to reach these satellites, so handsets are typically large and bulky. In addition, there is a large round-trip propagation delay; this delay is quite noticeable in two-way voice communication. Geosynchronous satellites also have fairly low data rates of less than 10 kbps. For these reasons, lower-orbit LEO satellites were thought to be a better match for voice and data communications.

LEO systems require approximately 30–80 satellites to provide global coverage, and plans for deploying such constellations were widespread in the late 1990s. One of the most ambitious of these systems, the Iridium constellation, was launched at that time. However, the cost to build, launch, and maintain these satellites is much higher than costs for terrestrial base stations. Although these LEO systems can certainly complement terrestrial systems in low-population areas and are also appealing to travelers desiring just one handset and phone

number for
bitious pla
eventually
launched
provides vo
of the Iridi

The mo
geographic
and satellit
broadcastin
since they
over, the co
service qui

1.4.8 Low

As radios
into more t
works, and
Bluetooth

Bluetoc
rudimentar
incorporati
place of a
ers, portab
short-range
phone to a
and this ran
system ope
without any
nel at 723.2
is a reverse
synchronou
Connection
also be use
1 Mbps. R
frequency
synchronou

Bluetoc
Typically,
any given t
the bandwi

⁴ The Bluetoc
united Der
for its nam

lightweight batteries. rate building walls. device and for the siderably in recent andsets, long battery outdoors are the main

a user that someone a landline telephone to allow a short digi- pagee as well. Radio 0 million subscribers e with the widespread tually the competition es. Some implemented ied a major change in tion to receiving them. ry large. Paging com- e paging functions into ing devices has shrunk and other professionals

wireless communications d OmniTRACS. The for- ations. For example, it is ones. The first-generation antenna) and rather expen- niques to enable smaller, 's OmniTRACS provides stem is used primarily for s. There are several major onous satellites. It takes a ally large and bulky. In ad- quite noticeable in two-way y low data rates of less than ght to be a better match for

provide global coverage, and late 1990s. One of the most ched at that time. However, gher than costs for terrestrial lement terrestrial systems in g just one handset and phone

number for global roaming, the growth and diminished cost of cellular prevented many ambitious plans for widespread LEO voice and data systems from materializing. Iridium was eventually forced into bankruptcy and disbanded, and most of the other systems were never launched. An exception to these failures was the Globalstar LEO system, which currently provides voice and data services over a wide coverage area at data rates under 10 kbps. Some of the Iridium satellites are still operational as well.

The most appealing use for a satellite system is the broadcast of video and audio over large geographic regions. Approximately 1 in 8 U.S. homes have direct broadcast satellite service, and satellite radio is also emerging as a popular service. Similar audio and video satellite broadcasting services are widespread in Europe. Satellites are best tailored for broadcasting, since they cover a wide area and are not compromised by an initial propagation delay. Moreover, the cost of the system can be amortized over many years and many users, making the service quite competitive in cost with that of terrestrial entertainment broadcasting systems.

1.4.8 Low-Cost, Low-Power Radios: Bluetooth and ZigBee

As radios decrease their cost and power consumption, it becomes feasible to embed them into more types of electronic devices, which can be used to create smart homes, sensor networks, and other compelling applications. Two radios have emerged to support this trend: Bluetooth and ZigBee.

Bluetooth⁴ radios provide short-range connections between wireless devices along with rudimentary networking capabilities. The Bluetooth standard is based on a tiny microchip incorporating a radio transceiver that is built into digital devices. The transceiver takes the place of a connecting cable for devices such as cell phones, laptop and palmtop computers, portable printers and projectors, and network access points. Bluetooth is mainly for short-range communications – for example, from a laptop to a nearby printer or from a cell phone to a wireless headset. Its normal range of operation is 10 m (at 1-mW transmit power), and this range can be increased to 100 m by increasing the transmit power to 100 mW. The system operates in the unlicensed 2.4-GHz frequency band, so it can be used worldwide without any licensing issues. The Bluetooth standard provides one asynchronous data channel at 723.2 kbps. In this mode, also known as Asynchronous Connection-Less (ACL), there is a reverse channel with a data rate of 57.6 kbps. The specification also allows up to three synchronous channels each at a rate of 64 kbps. This mode, also known as Synchronous Connection Oriented (SCO), is mainly used for voice applications such as headsets but can also be used for data. These different modes result in an aggregate bit rate of approximately 1 Mbps. Routing of the asynchronous data is done via a packet switching protocol based on frequency hopping at 1600 hops per second. There is also a circuit switching protocol for the synchronous data.

Bluetooth uses frequency hopping for multiple access with a carrier spacing of 1 MHz. Typically, up to eighty different frequencies are used for a total bandwidth of 80 MHz. At any given time, the bandwidth available is 1 MHz, with a maximum of eight devices sharing the bandwidth. Different logical channels (different hopping sequences) can simultaneously

⁴ The Bluetooth standard is named after Harald I Bluetooth, the king of Denmark between 940 and 985 A.D. who united Denmark and Norway. Bluetooth proposes to unite devices via radio connections, hence the inspiration for its name.

share the same 80-MHz bandwidth. Collisions will occur when devices in different piconets that are on different logical channels happen to use the same hop frequency at the same time. As the number of piconets in an area increases, the number of collisions increases and performance degrades.

The Bluetooth standard was developed jointly by 3 Com, Ericsson, Intel, IBM, Lucent, Microsoft, Motorola, Nokia, and Toshiba. The standard has now been adopted by over 1,300 manufacturers, and many consumer electronic products incorporate Bluetooth. These include wireless headsets for cell phones, wireless USB or RS232 connectors, wireless PCMCIA cards, and wireless settop boxes.

The ZigBee⁵ radio specification is designed for lower cost and power consumption than Bluetooth [19]. Its specification is based on the IEEE 802.15.4 standard. The radio operates in the same ISM band as Bluetooth and is capable of connecting 255 devices per network. The specification supports data rates of up to 250 kbps at a range of up to 30 m. These data rates are slower than Bluetooth, but in exchange the radio consumes significantly less power with a larger transmission range. The goal of ZigBee is to provide radio operation for months or years without recharging, thereby targeting applications such as sensor networks and inventory tags.

1.4.9 Ultrawideband Radios

Ultrawideband (UWB) radios are extremely wideband radios with very high potential data rates [20; 21]. The concept of ultrawideband communications actually originated with Marconi's spark-gap transmitter, which occupied a very wide bandwidth. However, since only a single low-rate user could occupy the spectrum, wideband communications was abandoned in favor of more efficient communication techniques. Renewed interest in wideband communications was spurred by the FCC's decision in 2002 to allow operation of UWB devices underlaid beneath existing users in the 3.1–10.6-GHz range. The underlay in theory interferes with all systems in that frequency range, including critical safety and military systems, unlicensed systems such as 802.11 wireless LANs and Bluetooth, and cellular systems where operators paid billions of dollars for dedicated spectrum use. The FCC's ruling was quite controversial given the vested interest in interference-free spectrum of these users. To minimize the impact of UWB on primary band users, the FCC put in place severe transmit power restrictions. This requires UWB devices to be within close proximity of their intended receiver.

Ultrawideband radios come with unique advantages that have long been appreciated by the radar and communications communities. Their wideband nature provides precise ranging capabilities. Moreover, the available UWB bandwidth has the potential for extremely high data rates. Finally, power restrictions dictate that the devices be small and with low power consumption.

Initial UWB systems used ultrashort pulses with simple amplitude or position modulation. Multipath can significantly degrade performance of such systems, and proposals to mitigate the effects of multipath include equalization and multicarrier modulation. Precise

⁵ ZigBee takes its name from the dance that honey bees use to communicate information about newly found food sources to other members of the colony.

and rapid sy
challenges
cially and in

1.5 The V

1.5.1 Met

Most countr
radio spectr
Commission
for military
Telecommu
munications
commercial
the FCC allo
to specific g
the 800-MH
operators in
tory bodies
blocks are n
argue that th
allocate the
ment beside
hurts techn
companies c
tain spectru
in high initi
eral compan
auctions.

In additi
that are free
may correspo
pose of thes
Many extren
LANs, Blue
can be killed
close proxim

Underlay
operates as a
ondary users
This is usual
is an exampl
Such underla
ing how inter

es in different piconets
ency at the same time.
ons increases and per-

n, Intel, IBM, Lucent,
a adopted by over 1,300
uetooth. These include
ors, wireless PCMCIA

ower consumption than
ard. The radio operates
55 devices per network.
up to 30 m. These data
significantly less power
dio operation for months
sensor networks and in-

a very high potential data
ally originated with Mar-
dth. However, since only
mmunications was aban-
ewed interest in wideband
ow operation of UWB de-
ge. The underlay in theory
cal safety and military sys-
etooth, and cellular systems
use. The FCC's ruling was
pectrum of these users. To
out in place severe transmit
proximity of their intended

ve long been appreciated by
ure provides precise ranging
potential for extremely high
e small and with low power

plitude or position modula-
ch systems, and proposals to
ticarrier modulation. Precise

information about newly found food

and rapid synchronization is also a big challenge for these systems. Although many technical challenges remain, the appeal of UWB technology has sparked great interest both commercially and in the research community to address these issues.

1.5 The Wireless Spectrum

1.5.1 Methods for Spectrum Allocation

Most countries have government agencies responsible for allocating and controlling use of the radio spectrum. In the United States, spectrum is allocated by the Federal Communications Commission (FCC) for commercial use and by the Office of Spectral Management (OSM) for military use. Commercial spectral allocation is governed in Europe by the European Telecommunications Standards Institute (ETSI) and globally by the International Telecommunications Union (ITU). Governments decide how much spectrum to allocate between commercial and military use, and this decision is dynamic depending on need. Historically the FCC allocated spectral blocks for specific uses and assigned licenses to use these blocks to specific groups or companies. For example, in the 1980s the FCC allocated frequencies in the 800-MHz band for analog cellular phone service and provided spectral licenses to two operators in each geographical area based on a number of criteria. The FCC and regulatory bodies in other countries still allocate spectral blocks for specific purposes, but these blocks are now commonly assigned through spectral auctions to the highest bidder. Some argue that this market-based method is the fairest and most efficient way for governments to allocate the limited spectral resource and that it provides significant revenue to the government besides; others believe that this mechanism stifles innovation, limits competition, and hurts technology adoption. Specifically, the high cost of spectrum dictates that only large companies or conglomerates can purchase it. Moreover, the large investment required to obtain spectrum can delay the ability to invest in infrastructure for system rollout and results in high initial prices for the end user. The 3G spectral auctions in Europe, in which several companies ultimately defaulted, have provided ammunition to the opponents of spectral auctions.

In addition to spectral auctions, spectrum can be set aside in specific frequency bands that are free to use without a license according to a specific set of etiquette rules. The rules may correspond to a specific communications standard, power levels, and so forth. The purpose of these *unlicensed bands* is to encourage innovation and low-cost implementation. Many extremely successful wireless systems operate in unlicensed bands, including wireless LANs, Bluetooth, and cordless phones. A major difficulty of unlicensed bands is that they can be killed by their own success. If many unlicensed devices in the same band are used in close proximity then they interfere with each other, which can make the band unusable.

Underlay systems are another alternative for allocating spectrum. An underlay system operates as a secondary user in a frequency band with other primary users. Operation of secondary users is typically restricted so that primary users experience minimal interference. This is usually accomplished by restricting the power per hertz of the secondary users. UWB is an example of an underlay system, as are unlicensed systems in the ISM frequency bands. Such underlay systems can be extremely controversial given the complexity of characterizing how interference affects the primary users. Yet the trend toward spectrum allocation for

Diversity

In Chapter 6 we saw that both Rayleigh fading and log-normal shadowing exact a large power penalty on the performance of modulation over wireless channels. One of the best techniques to mitigate the effects of fading is diversity combining of independently fading signal paths. Diversity combining exploits the fact that independent signal paths have a low probability of experiencing deep fades simultaneously. Thus, the idea behind diversity is to send the same data over independent fading paths. These independent paths are combined in such a way that the fading of the resultant signal is reduced. For example, consider a system with two antennas at either the transmitter or receiver that experience independent fading. If the antennas are spaced sufficiently far apart, it is unlikely that they both experience deep fades at the same time. By selecting the antenna with the strongest signal, a technique known as *selection combining*, we obtain a much better signal than if we had just one antenna. This chapter focuses on common methods used at the transmitter and receiver to achieve diversity. Other diversity techniques that have potential benefits beyond these schemes in terms of performance or complexity are discussed in [1, Chap. 9.10].

Diversity techniques that mitigate the effect of multipath fading are called *microdiversity*, and that is the focus of this chapter. Diversity to mitigate the effects of shadowing from buildings and objects is called *macrodiversity*. Macrodiversity is generally implemented by combining signals received by several base stations or access points, which requires coordination among these different stations or points. Such coordination is implemented as part of the networking protocols in infrastructure-based wireless networks. We will therefore defer discussion of macrodiversity until Chapter 15, where we examine the design of such networks.

7.1 Realization of Independent Fading Paths

There are many ways of achieving independent fading paths in a wireless system. One method is to use multiple transmit or receive antennas, also called an antenna array, where the elements of the array are separated in distance. This type of diversity is referred to as *space diversity*. Note that with receiver space diversity, independent fading paths are realized without an increase in transmit signal power or bandwidth. Moreover, coherent combining of the diversity signals increases the signal-to-noise power ratio at the receiver over the SNR that

would be obtained with just a single receive antenna. This SNR increase, called array gain, can also be obtained with transmitter space diversity by appropriately weighting the antenna transmit powers relative to the channel gains. In addition to array gain, space diversity also provides diversity gain, defined as the change in slope of the error probability resulting from the diversity combining. We will describe both the array gain and diversity gain for specific diversity-combining techniques in subsequent sections.

The maximum diversity gain for either transmitter or receiver space diversity typically requires that the separation between antennas be such that the fading amplitudes corresponding to each antenna are approximately independent. For example, from equation (3.26) we know that, in a uniform scattering environment with omnidirectional transmit and receive antennas, the minimum antenna separation required for independent fading on each antenna is approximately one half-wavelength ($.38\lambda$, to be exact). If the transmit or receive antennas are directional (which is common at the base station if the system has cell sectorization), then the multipath is confined to a small angle relative to the LOS ray, which means that a larger antenna separation is required to obtain independent fading samples [2].

A second method of achieving diversity is by using either two transmit antennas or two receive antennas with different polarization (e.g., vertically and horizontally polarized waves). The two transmitted waves follow the same path. However, since the multiple random reflections distribute the power nearly equally relative to both polarizations, the average receive power corresponding to either polarized antenna is approximately the same. Since the scattering angle relative to each polarization is random, it is highly improbable that signals received on the two differently polarized antennas would be simultaneously in a deep fade. There are two disadvantages of polarization diversity. First, you can have at most two diversity branches, corresponding to the two types of polarization. The second disadvantage is that polarization diversity loses effectively half the power (3 dB) because the transmit or receive power is divided between the two differently polarized antennas.

Directional antennas provide angle (or directional) diversity by restricting the receive antenna beamwidth to a given angle. In the extreme, if the angle is very small then at most one of the multipath rays will fall within the receive beamwidth, so there is no multipath fading from multiple rays. However, this diversity technique requires either a sufficient number of directional antennas to span all possible directions of arrival or a single antenna whose directivity can be steered to the arrival angle of one of the multipath components (preferably the strongest one). Note also that with this technique the SNR may decrease owing to the loss of multipath components that fall outside the receive antenna beamwidth – unless the directional gain of the antenna is sufficiently large to compensate for this lost power. *Smart antennas* are antenna arrays with adjustable phase at each antenna element: such arrays form directional antennas that can be steered to the incoming angle of the strongest multipath component [3].

Frequency diversity is achieved by transmitting the same narrowband signal at different carrier frequencies, where the carriers are separated by the coherence bandwidth of the channel. This technique requires additional transmit power to send the signal over multiple frequency bands. Spread-spectrum techniques, discussed in Chapter 13, are sometimes described as providing frequency diversity in that the channel gain varies across the bandwidth of the transmitted signal. However, this is not equivalent to sending the same information




Figure 16.1: Ad hoc network.

16.1 Applications

This section describes some of the most prevalent applications for ad hoc wireless networks. The self-configuring nature and lack of infrastructure inherent to these networks make them highly appealing for many applications, even if it results in a significant performance penalty. The lack of infrastructure is highly desirable for low-cost commercial systems, since it obviates the need for a large investment to get the network up and running, and deployment costs may then scale with network success. Lack of infrastructure is also highly desirable for military systems, where communication networks must be configured quickly as the need arises, often in remote areas. Other advantages of ad hoc wireless networks include ease of network reconfiguration and reduced maintenance costs. However, these advantages must be balanced against any performance penalty resulting from the multihop routing and distributed control inherent to these networks.

We will focus on the following applications: data networks, home networks, device networks, sensor networks, and distributed control systems. This list is by no means comprehensive; in fact, the success of ad hoc wireless networks hinges on making them sufficiently flexible so that there can be accidental successes. Therein lies the design dilemma for these networks. If the network is designed for maximum flexibility to support many applications (a one-size-fits-all network) then it will be difficult to tailor the network to different application requirements. This will likely result in poor performance for some applications, especially

those w
network
what th
work sh
perform
sign can
constrai

16.1.1

Ad-hoc
palmtop
works go
WANs (f
less LAN
ever, ad
networks
are dispe
installati
cate throu
can excha

Wireles
cover a la
effective n
changes w
users furth
mobile int
ad hoc wi
less MAN
[9]. Altho
politan are
bankruptcy

Wireles
area is too
into remote
that must b
are unfeasi

16.1.2 Ho

Home netw
less phones
entertainme
rooms that s
"aware hom
those with d

those with high rate requirements or stringent delay constraints. On the other hand, if the network is optimized for a few specific applications then designers must predict in advance what these “killer applications” will be – a risky proposition. Ideally an ad hoc wireless network should be sufficiently flexible to support many different applications while adapting its performance to the applications in operation at any given time. An adaptive cross-layer design can provide this flexibility along with the ability to tailor protocol design to different constraints in the nodes.

16.1.1 Data Networks

Ad-hoc wireless data networks primarily support data exchange between laptop computers, palmtops, personal digital assistants (PDAs), and other information devices. These data networks generally fall into three categories based on their coverage area: LANs, MANs, and WANs (for “local”, “metropolitan”, and “wide” area networks). Infrastructure-based wireless LANs are already quite prevalent, and they deliver good performance at low cost. However, ad hoc wireless data networks have some advantages over these infrastructure-based networks. First, with multihop routing only one access point is needed to connect nodes that are dispersed throughout the area to the backbone wired infrastructure: this reduces cost and installation requirements. In addition, it can sometimes be inefficient for nodes to communicate through an access point or base station. For example, PDAs that are next to each other can exchange information directly rather than routing through an intermediate node.

Wireless MANs typically require multihop routing, often over many hops, because they cover a larger area. The challenge in these networks is to support high data rates in a cost-effective manner and over multiple hops, where the link quality of each hop is different and changes with time. The lack of centralized network control and potential for high-mobility users further complicate this objective. Military programs such as DARPA’s GLOMO (global mobile information systems) have invested much time and money in building high-speed ad hoc wireless MANs that support multimedia, with limited success [7; 8]. Ad hoc wireless MANs have also permeated the commercial sector, with Metricom the best example [9]. Although Metricom did deliver fairly high data rates throughout several major metropolitan areas, significant demand never materialized and so the company eventually filed for bankruptcy.

Wireless WANs are needed for applications where network infrastructure to cover a wide area is too costly or impractical to deploy. For example, sensor networks may be dropped into remote areas where network infrastructure cannot be developed. In addition, networks that must be built up and torn down quickly (e.g., for military applications or disaster relief) are unfeasible without an ad hoc approach.

16.1.2 Home Networks

Home networks are envisioned to support communication between PCs, laptops, PDAs, cordless phones, smart appliances, security and monitoring systems, consumer electronics, and entertainment systems anywhere in and around the home. Such networks could enable smart rooms that sense people and movement and adjust light and heating accordingly, as well as “aware homes” that feature network sensors and computers for assisted living of seniors and those with disabilities. Home networks also encompass video or sensor monitoring systems

the higher layers. We now give some examples of this interaction for physical layer design choices related to packet error rate, multiple antennas, and power control.

In most wireless ad hoc networks, bits are packetized for transmission as described in Section 14.3. The design choices at the physical layer along with the channel and interference conditions determine the link packet error rate (PER). Many access layer protocols retransmit packets received in error, so PER based on the physical layer design affects the retransmission requirements at the access layer. Similarly, as described in Section 10.8, multiple antennas give rise to a multiplexing/diversity/directionality trade-off: the antennas can be used to increase the data rate on the link, to provide diversity in fading so that average BER is reduced, or to provide directionality to reduce fading and the interference a signal causes to other signals. Diversity gain will reduce PER, leading to fewer retransmissions. Multiplexing will increase the link rate, which reduces congestion and delay on the link and benefits all multihop routes using that link. Directionality reduces interference to other links, thereby improving their performance. Thus, the best use of multiple antennas in an ad hoc wireless network clearly transcends just the physical layer; in fact, it simultaneously affects the physical, access, network, and transport layers.

The transmit power of a node at the physical layer also has a broad impact across many layers of the protocol stack. Increasing transmit power at the physical layer reduces PER, thereby affecting the retransmissions required at the access layer. In fact any two nodes in the network can communicate directly with sufficiently high transmit power, so this power drives link connectivity. However, a high transmit power from one node in the network can cause significant interference to other nodes, thereby degrading their performance and possibly breaking their connectivity on some links. In particular, link performance in an ad hoc wireless network is driven by SINR, so the transmit power of all nodes has an impact on the performance of all links in the network. Broadly speaking, the transmit power coupled with adaptive modulation and coding for a given node defines its "local neighborhood" – the collection of nodes that it can reach in a single hop – and thus defines the context in which access, routing, and other higher-layer protocols operate. Therefore, the transmit power of all nodes in the network must be optimized with respect to all layers that it affects. As such, it is a prime motivator for a cross-layer design.

16.3.2 Access Layer Design

The access layer controls how different users share the available spectrum and ensures successful reception of packets transmitted over this shared spectrum. Allocation of signaling dimensions to different users is accomplished through either multiple access or random access, and a detailed discussion of these access techniques can be found in Sections 14.2 and 14.3. Multiple access divides the signaling dimensions into dedicated channels via orthogonal or nonorthogonal channelization methods. The most common of these methods are TDMA, FDMA, and CDMA. The access layer must also provide control functionality in order to assign channels to users and to deny access to users when they cannot be accommodated by the system. In random access, channels are assigned to active users dynamically, and in multihop networks these protocols must contend with hidden and exposed terminals. The most common random access methods are different forms of ALOHA, carrier-sense

multi
incorp
As
other
func
all lin
power
multi
follow
transm

where g
of the k
the rece
 $\rho \approx 1/6$
link requ
quireme
form as
powers a

is the vec
matrix w

where k_j
in Section
of $\mathbf{F}$ is les
SINR requ
tion [33; 3
on all links
on the k th

² As noted pr
with K no
SINR on th

physical layer design
control.

mission as described in the channel and interference access layer protocols. The layer design affects the delay in Section 10.8, multiple-off: the antennas can be fading so that average interference a signal fewer retransmissions. delay on the link and interference to other links, antennas in an ad hoc simultaneously affects

ad impact across many
cal layer reduces PER,
fact any two nodes in
t power, so this power
ode in the network can
performance and pos-
formance in an ad hoc
odes has an impact on
transmit power coupled
al neighborhood”—the
es the context in which
the transmit power of
that it affects. As such,

trum and ensures suc-
Allocation of signaling
e access or random ac-
ound in Sections 14.2
icated channels via or-
on of these methods are
ontrol functionality in
y cannot be accommo-
ive users dynamically,
and exposed terminals.
LOHA, carrier-sense

multiple access (CSMA), and scheduling (see Section 14.3). These random access methods incorporate channel assignment and access denial into their protocols.

As discussed in the prior section, transmit power associated with a single node affects all other nodes. Thus, power control across all nodes in the network is part of the access layer functionality. The main role of power control is to ensure that SINR targets can be met on all links in the network. This is often unfeasible, as discussed in more detail below. The power control algorithms described in Sections 14.4 and 15.5.3 for meeting SINR targets in multiple access and cellular systems, respectively, can be extended to ad hoc networks as follows. Consider an ad hoc wireless network with K nodes and N links between different transmitter–receiver pairs of these nodes.² The SINR on link k is given by

$$\gamma_k = \frac{g_{kk} P_k}{n_k + \rho \sum_{i \neq k} g_{ki} P_i}, \quad k, j \in \{1, 2, \dots, N\}, \quad (16.1)$$

where $g_{kj} > 0$ is the channel power gain from the transmitter of the j th link to the receiver of the k th link, P_k is the power of the transmitter on the k th link, n_k is the noise power of the receiver on the k th link, and ρ is the interference reduction due to signal processing (i.e., $\rho \approx 1/G$ for CDMA with processing gain G and $\rho = 1$ in TDMA). Suppose that the k th link requires an SINR of γ_k^* as determined, for example, by the connectivity and data-rate requirements for that link. Then the SINR constraints for all links can be represented in matrix form as $(\mathbf{I} - \mathbf{F})\mathbf{P} \geq \mathbf{u}$ with $\mathbf{P} > 0$, where $\mathbf{P} = (P_1, P_2, \dots, P_N)^T$ is the vector of transmit powers associated with the transmitters on the N links,

$$\mathbf{u} = \left(\frac{\gamma_1^* n_1}{g_{11}}, \frac{\gamma_2^* n_2}{g_{22}}, \dots, \frac{\gamma_N^* n_N}{g_{NN}} \right)^T \quad (16.2)$$

is the vector of noise powers scaled by the SINR constraints and channel gains, and $\mathbf{F}$ is a matrix with

$$F_{kj} = \begin{cases} 0 & k = j, \\ \gamma_k^* g_{kj} \rho / g_{kk} & k \neq j, \end{cases} \quad (16.3)$$

where $k, j \in \{1, 2, \dots, N\}$. As in the uplink and cellular power control problems described in Section 14.4 and Section 15.5.3, if the Perron–Frobenius (maximum modulus) eigenvalue of $\mathbf{F}$ is less than unity then there exists a vector $\mathbf{P}^* > 0$ (i.e., $P_k > 0$ for all k) such that the SINR requirements of all links are satisfied, with $\mathbf{P}^* = (\mathbf{I} - \mathbf{F})^{-1} \mathbf{u}$ the Pareto optimal solution [33; 34; 35]. Thus $\mathbf{P}^*$ meets the SINR requirements with the minimum transmit power on all links. Moreover, a distributed iterative power control algorithm, where the transmitter on the k th link updates its transmit power at time $i + 1$ as

$$P_k(i+1) = \frac{\gamma_k^*}{\gamma_k(i)} P_k(i), \quad (16.4)$$

² As noted previously, all nodes can communicate with all other nodes, so there are $K(K-1)$ links for a network with K nodes. However, we assume that only $N \leq K(K-1)$ of these are used, so we need only consider the SINR on these N links.

can be shown [34; 35] to converge to the optimal solution P^* . This is a simple distributed algorithm for power control in an ad hoc wireless network, since it requires only that the SINR of the k th link be known to the transmitter on that link. Then, if this SINR is below its target the transmitter increases its power, and if it is above this target the transmitter decreases the power. It is quite remarkable that such a simple distributed algorithm converges to a globally optimal power control. However, when the channel gains are not static, SINR constraints can no longer be met with certainty, and it is much more difficult to develop distributed power control algorithms that meet a desired performance target [36]. In particular, the algorithm described by (16.4) can exhibit large fluctuations in link SINR when the channel gains vary over time. More importantly, it is often impossible to meet the SINR constraints of all nodes simultaneously even when the link gains are static, owing to the large number of interferers and the range of channel gains associated with all signals in the network. When the SINR constraints cannot be met, the distributed power control algorithm will diverge such that all nodes transmit at their maximum power and still cannot meet their SINR constraints. This is obviously an undesirable operational state, especially for energy-constrained nodes.

In [37], the distributed power control algorithm assuming static link gains is extended to include distributed admission control. The admission control provides protection for existing link SINR targets when a new user enters the system. In this scheme the active links have a slightly higher SINR target than needed. This buffer is used so that, when a new user who attempts to access the system transmits at a low power level, he will not cause the active links to fall below their minimum SINR targets. The new user gradually ramps up his power and checks if he gets closer to his SINR target as a result. If the new user can be accommodated by the system without violating the SINR constraints of existing links, then the distributed algorithm with this gradual ramp-up will eventually converge to a new P^* that satisfies the SINR constraints of the new and existing links. However, if the new user cannot be accommodated, his gradual ramp-up will not approach his required SINR and he will eventually be forced to leave the system. Note that the criteria for denial of access is difficult to optimize in time-varying channels with distributed control [36]. These ideas are combined with transmission scheduling in [38; 39] to improve power efficiency and reduce interference. The power control algorithm can also be modified to take into account delay constraints [23; 40]. However, delay constraints are associated with the full multihop route of a packet, so power control should be coordinated with network layer protocols to ensure delay constraints are met on the end-to-end route.

The access layer is also responsible for retransmissions of packets received in error over the wireless link, often referred to as the automatic repeat request (ARQ) protocol. Specifically, data packets typically have an error detection code that is used by the receiver to determine if one or more bits in the packet were corrupted and cannot be corrected. For such packets, the receiver will usually discard the corrupted packet and inform the transmitter via a feedback channel that the packet must be retransmitted. However, rather than discarding the packet, the access layer can save it and use a form of diversity to combine the corrupted packet with the retransmitted packet for a higher probability of correct packet reception. Alternatively, rather than retransmitting the original packet in its entirety, the transmitter can just send some additional coded bits to provide a stronger error correction capability for the receiver to correct for the corrupted bits in the original packet. This technique is called

Table D.5: 802.11 wireless LAN link layer standards

Parameter	802.11	802.11a	802.11b	802.11g
Bandwidth (MHz)	83.5	300	83.5	83.5
Frequency range (GHz)	2.4–2.4835	5.15–5.25 (lower) 5.25–5.35 (middle) 5.725–5.825 (upper)	2.4–2.4835	2.4–2.4835
Number of channels	3	12 (4 per subband)	3	3
Modulation	BPSK, QPSK DSSS, FHSS	BPSK, QPSK, MQAM OFDM	BPSK, QPSK DSSS	BPSK, QPSK, MQAM OFDM
Coding	Undefined	Convolutional (rate 1/2, 2/3, 3/4)	Barker, CCK	Convolutional (rate 1/2, 2/3, 3/4)
Maximum data rate (Mbps)	2	54	11	54
Range (m)	Undefined	27–30 (lower band)	75–100	30
Random access	CSMA/CA	CSMA/CA	CSMA/CA	CSMA/CA

peak rates of about 300 kbps and actual rates of about 144 kbps. There are two evolutions of this core technology to provide high data rates (referred to as HDR service) above 1 Mbps; these evolutions are referred to as cdma2000 1XEV. The first phase of evolution, cdma2000 1XEV-DO (data only), enhances the cdmaOne system using a separate 1.25-MHz dedicated high-speed data channel that supports downlink data rates up to 3 Mbps and uplink data rates up to 1.8 Mbps for an average combined rate of 2.4 Mbps. The second phase of the evolution, cdma2000 1XEV-DV (data and voice), is projected to support up to 4.8-Mbps data rates as well as legacy 1X voice users, 1XRTT data users, and 1XEV-DO data users – and all within the same radio channel. Another proposed enhancement to cdma2000 is to aggregate three 1.25-MHz channels into one 3.75-MHz channel. This aggregation is referred to as cdma2000 3X, and its exact specifications are still under development.

W-CDMA is the primary competing 3G standard to cdma2000. It has been selected as the 3G successor to GSM and in this context is referred to as the Universal Mobile Telecommunications System (UMTS). W-CDMA is also used in the Japanese FOMA and J-Phone 3G systems. These different systems share the W-CDMA link layer protocol (air interface) but have different protocols for other aspects of the system such as routing and speech compression. W-CDMA supports peak rates of up to 2.4 Mbps, with typical rates anticipated in the 384-kbps range. W-CDMA uses 5-MHz channels, in contrast to the 1.25-MHz channels of cdma2000. An enhancement to W-CDMA called High Speed Data Packet Access (HSDPA) provides data rates of around 9 Mbps, and this may be the precursor to fourth-generation systems.

D.2 Wireless Local Area Networks

Wireless local area networks (WLANs) are built around the family of IEEE 802.11 standards. The main characteristics of this standards family are summarized in Table D.5. The baseline 802.11 standard, released in 1997, occupies 83.5 MHz of bandwidth in the unlicensed 2.4-GHz frequency band. It specifies PSK modulation with FHSS or DSSS. Data rates of up to 2 Mbps are supported, with CSMA/CA used for random access. The baseline standard

was expanded in 1999 to create the 802.11b standard, operating in the same 2.4-GHz band using only DSSS. This standard uses variable-rate modulation and coding, with BPSK or QPSK for modulation and channel coding via either Barker sequences or complementary code keying (CCK). This leads to a maximum channel rate of 11 Mbps and a maximum user data rate of about 1.6 Mbps. The transmission range is approximately 100 m. The network architecture in 802.11b is specified as either star or peer-to-peer, although the peer-to-peer feature is not typically used. This standard has been widely deployed, with manufacturers integrating 802.11b wireless LAN cards into many laptop computers.

The 802.11a standard was finalized in 1999 as an extension to 802.11 to improve on the 802.11b data rates. The 802.11a standard occupies 300 MHz of spectrum in the 5-GHz U-NII band. In fact, the 300 MHz of bandwidth is segmented into three 100-MHz subbands: a lower band (5.15–5.25 GHz), a middle band (5.25–5.35 GHz), and an upper band (5.725–5.825 GHz). Channels are spaced 20 MHz apart except on the outer edges of the lower and middle bands, where they are spaced 30 MHz apart. Three maximum transmit power levels are specified: 40 mW for the lower band, 200 mW for the middle band, and 800 mW for the upper band. These restrictions imply that the lower band is mostly suitable for indoor applications only, the middle band for indoor and outdoor, and the high band for outdoor. Variable-rate modulation and coding is used on each channel: the modulation varies over BPSK, QPSK, 16-QAM, and 64-QAM, and the convolutional code rate varies over 1/2, 2/3, and 3/4. This leads to a maximum data rate per channel of 54 Mbps. For indoor systems, the 5-GHz carrier coupled with the power restriction in the lower band reduces the range of 802.11a relative to 802.11b and also makes it more difficult for the signal to penetrate walls and other obstructions. The 802.11a standard uses orthogonal frequency division multiplexing (OFDM) multiple access instead of FHSS or DSSS, and in that sense it diverges from the original 802.11 standard.

The 802.11g standard, finalized in 2003, attempts to combine the best of 802.11a and 802.11b, with data rates of up to 54 Mbps in the 2.4-GHz band for greater range. The standard is backward compatible with 802.11b, so 802.11g access points will work with 802.11b wireless network adapters and vice versa. However, 802.11g uses the OFDM, modulation, and coding schemes of 802.11a. Access points and wireless LAN cards are available with all three standards to avoid incompatibilities. The 802.11a/b/g family of standards are collectively referred to as *Wi-Fi*, for wireless fidelity. Extending these standards to frequency allocations in countries other than the United States falls under the 802.11d standard. There are several other standards in the 802.11 family that are under development; these are summarized in Table D.6.

A potential competitor to the 802.11 standards as well as cellular systems is the emerging IEEE 802.16 standard called WiMAX. This standard promises broadband wireless access with data rates on the order of 40 Mbps for fixed users and 15 Mbps for mobile users, with a range of several kilometers. Details of the specification are still being worked out.

D.3 Wireless Short-Distance Networking Standards

This last section summarizes the main characteristics of ZigBee, Bluetooth, and UWB (ultra-wideband), which have emerged to support a wide range of short-distance wireless network

Table D.6: IEEE 802.11 ongoing standards work

Standard	Scope
802.11e	Provides quality of service (QoS) at the MAC layer
802.11f	Roaming protocol across multivendor access points
802.11h	Adds frequency and power management features to 802.11a to make it more compatible with European operation
802.11i	Enhances security and authentication mechanisms
802.11j	Modifies 802.11a link layer to meet Japanese requirements
802.11k	Provides an interface to higher layers for radio and network measurements that can be used for radio resource management
802.11m	Maintenance of 802.11 standard (technical/editorial corrections)
802.11n	MIMO link enhancements to enable higher throughput

Table D.7: Short-range wireless network standards

Parameter	ZigBee (802.15.4)	Bluetooth (802.15.1)	UWB (802.15.3 proposal)
Frequency range (GHz)	2.4–2.4835	2.4–2.4835	3.1–10.6
Bandwidth (MHz)	83.5	83.5	7500
Modulation	BPSK, OQPSK DSSS	GFSK FHSS	BPSK, QPSK OFDM or DSSS
Maximum data rate (Mbps)	.25	1	100
Range (m)	30	10 (100)	10
Power consumption (mW)	5–20	1 (100)	80–150 mW
Access	CSMA/CA (optional TD)	TD	Undefined
Networking	Mesh/Star/Tree	Subnet clusters (8 nodes)	Undefined

applications. These specifications are designed to be compliant with the IEEE 802.15 standards, a family of IEEE standards for short-distance wireless networking called Wireless Personal Area Networks (WPANs). Bluetooth operates in the 2.4-GHz unlicensed band; ZigBee operates in the same band as well as in the 800-MHz and 900-MHz unlicensed bands; and UWB operates across a broad range of frequencies in an underlay to existing systems. ZigBee and Bluetooth include link, MAC, and higher-layer protocol specifications, whereas UWB specifies just the link layer protocol. Table D.7 summarizes the main characteristics of ZigBee (2.4-GHz band only), Bluetooth, and UWB.

ZigBee consists of link and MAC layer protocols that are compliant with the IEEE 802.15.4 standard, as well as higher-layer protocols for ad hoc networking (mesh, star, or tree topologies), power management, and security. ZigBee supports data rates of up to 250 kbps with PSK modulation and DSSS. ZigBee generally targets applications requiring relatively low data rates, low duty cycles, and large networks. Power efficiency is key, with the goal of nodes operating for months or years on a single battery charge.

In contrast to ZigBee, Bluetooth provides up to 1-Mbps data rate, including three guaranteed low-latency voice channels, using FHSS and FSK modulation with Gaussian pulse

REFERENCES

- [1] T. S. Eng
- [2] J. D. mov
- [3] D. F. IEE
- [4] I. Po
- [5] S. H. NJ, 2
- [6] W. S. NJ, 2

shaping (Gaussian FSK or GFSK). Bluetooth normally transmits at a power of 1 mW with a transmission range of about 10 m, although this can be extended to 100 m by increasing the transmit power to 100 mW. Networks are formed in subnet clusters (piconets) of up to eight nodes, with one node acting as a master and the rest as slaves. Time division is used for channel access, with the master node coordinating the frequency-hopping sequence and synchronization with the slave nodes. Extended networks, or *scatternets*, can be formed when one node is part of multiple piconets. However, forming large networks through this approach is difficult owing to the synchronization requirements of FHSS. Portions of the Bluetooth standard were formally adopted by the IEEE as its 802.15.1 standard.

Ultrawideband has significantly higher data rates – up to 100 Mbps – than either ZigBee or Bluetooth. It also occupies significantly more bandwidth and has stringent power restrictions (to prevent it from interfering with primary band users). Thus, it is suitable only for short-range indoor applications. The UWB protocol defines only a link layer technology, so it requires a compatible MAC protocol as well as higher-layer protocols to become part of a wireless network standard. The modulation is BPSK or QPSK, with competing camps recommending either OFDM or DSSS overlayed on the data modulation. Ultrawideband is likely to become the link layer technology for the IEEE 802.15.3 standard, a family of standards for wireless networks supporting imaging and multimedia applications.

REFERENCES

- [1] T. S. Rappaport, *Wireless Communications – Principles and Practice*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ, 2001.
- [2] J. D. Vriendt, P. Lainé, C. Lerouge, and X. Xu, "Mobile network evolution: A revolution on the move," *IEEE Commun. Mag.*, pp. 104–11, April 2002.
- [3] D. Porcino and W. Hirt, "Ultra-wideband radio technology: Potential and challenges ahead," *IEEE Commun. Mag.*, pp. 66–74, July 2003.
- [4] I. Poole, "What exactly is ... ZigBee?" *IEEE Commun. Eng.*, pp. 44–5, August/September 2004.
- [5] S. Haykin and M. Moher, *Modern Wireless Communications*, Prentice-Hall, Englewood Cliffs, NJ, 2005.
- [6] W. Stallings, *Wireless Communications and Networks*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ, 2005.

Bibliography

- Abou-Faycal, I. C., M. D. Trott, and S. Shamai, "The capacity of discrete-time memoryless Rayleigh fading channels," *IEEE Trans. Inform. Theory*, pp. 1290–1301, May 2001.
- Abramson, N., "The ALOHA system – Another alternative for computer communications," *Proc. Amer. Federation Inform. Proc. Soc. Fall Joint Comput. Conf.*, pp. 281–5, November 1970.
- Abramson, N., "Wide-band random-access for the last mile," *IEEE Pers. Commun. Mag.*, pp. 29–33, December 1996.
- Abreshamkar, F., and Z. Siveski, "PCS global mobile satellites," *IEEE Commun. Mag.*, pp. 132–6, September 1996.
- Abu-Dayya, A., and N. Beaulieu, "Analysis of switched diversity systems on generalized-fading channels," *IEEE Trans. Commun.*, pp. 2959–66, November 1994.
- Abu-Dayya, A., and N. Beaulieu, "Switched diversity on microcellular Ricean channels," *IEEE Trans. Veh. Tech.*, pp. 970–6, November 1994.
- Agrawal, P., "Energy efficient protocols for wireless systems," *Proc. Internat. Sympos. Pers., Indoor, Mobile Radio Commun.*, pp. 564–9, September 1998.
- Ahlsweide, R., N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, pp. 1204–16, July 2000.
- Akerberg, D., "Properties of a TDMA picocellular office communication system," *Proc. IEEE Globecom Conf.*, pp. 1343–9, December 1988.
- Akin, H. C., and K. M. Wasserman, "Resource allocation and scheduling in uplink for multimedia CDMA wireless systems," *IEEE/Sarnoff Sympos. Adv. Wired Wireless Commun.*, pp. 185–8, April 2004.
- Alamouti, S., "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, pp. 1451–8, October 1998.
- Alamouti, S. M., and S. Kallel, "Adaptive trellis-coded multiple-phased-shift keying for Rayleigh fading channels," *IEEE Trans. Commun.*, pp. 2305–14, June 1994.
- Alasti, M., K. Sayrafian-Pour, A. Ephremides, and N. Farvardin, "Multiple description coding in networks with congestion problem," *IEEE Trans. Inform. Theory*, pp. 891–902, March 2001.
- Al-Dhahir, N., C. Fragouli, A. Stamoulis, W. Younis, and R. Calderbank, "Space-time processing for broadband wireless access," *IEEE Commun. Mag.*, pp. 136–42, September 2002.
- Algans, A., K. I. Pedersen, and P. E. Mogensen, "Experimental analysis of the joint statistical properties of azimuth spread, delay spread, and shadow fading," *IEEE J. Sel. Areas Commun.*, pp. 523–31, April 2002.
- Alouini, M.-S., and A. J. Goldsmith, "Area spectral efficiency of cellular mobile radio systems," *IEEE Trans. Veh. Tech.*, pp. 1047–66, July 1999.
- Alouini, M.-S., and A. J. Goldsmith, "Capacity of Rayleigh fading channels under different adaptive transmission and diversity combining techniques," *IEEE Trans. Veh. Tech.*, pp. 1165–81, July 1999.