

Self-Study
Report
on
Personal Area Networks

Submitted by
H. Srikanth
EE02B068
Guide : Dr. T. G. Venkatesh

Contents

1. Personal Area Networks

- 1.1 Definition
- 1.2 Examples

2. Bluetooth

- 2.1 Introduction
- 2.2 Technical Information
- 2.3 Protocol Stack
- 2.4 Applications

3. IrDA

- 3.1 Introduction
- 3.2 Protocol Stack and Applications

4. WiMax

- 4.1 Introduction
- 4.2 Protocols and Technical Information
- 4.3 Applications

5 WiFi

- 5.1 Introduction
- 5.2 Specification
- 5.3 Security
- 5.4 Advantages and Disadvantages

6. IEEE 802.11

- 6.1 Introduction
- 6.2 Various Technologies
- 6.3 Protocol Stack with a detailed view of the MAC layer
- 6.4 IEEE 802.11 list of Standards

7. Firewire

- 7.1 Introduction
- 7.2 Standards and versions
- 7.3 Security Issues

8. Home RF

References

1. Personal Area Networks

1.1 Definition:

A personal area network (PAN) is a computer network used for Communication among computer devices (including telephones and personal digital assistants) close to one person. The devices may or may not belong to the person in question. The reach of a PAN is typically a few meters. PANs can be used for communication among the personal devices themselves (intrapersonal communication), or for connecting to a higher level network and the Internet (an uplink).

Personal area networks may be wired with computer buses such as USB and Firewire. A wireless personal area network (WPAN) can also be made possible with network technologies such as IrDA and Bluetooth.

1.2 Various Personal Area Networks:

1. Bluetooth
2. IrDA
3. WiMax
4. WiFi
5. IEEE 802.11
6. Firewire
7. Home RF

2. Bluetooth

2.1 Introduction:

Bluetooth is an industrial specification for wireless personal area networks (PANs). Bluetooth provides a way to connect and exchange information between devices like personal digital assistants (PDAs), mobile phones, laptops, PCs, printers and digital cameras via a secure, low cost, globally available short range radio frequency. Bluetooth lets these devices talk to each other when they come in range, even if they are not in the same room, as long as they are within up to 100 metres (328 feet) of each other.

Bluetooth is a wireless radio standard primarily designed for low power consumption, with a short range and with a low cost transceiver microchip in each device. Cell phones with integrated Bluetooth technology have also been sold in large numbers, and are able to connect to computers, PDAs and, specifically, to hands-free devices. The standard also includes support for more powerful, longer range devices suitable for constructing wireless LANs.

2.2 Technical Information:

A Bluetooth device playing the role of the "master" can communicate with up to 7 devices playing the role of the "slave". At any given time, data can be transferred between the master and one slave; but the master switches rapidly from slave to slave in a round-robin fashion. (Simultaneous transmission from the master to multiple slaves is possible, but not used much in practice). These groups of up to 8 devices (1 master and 7 slaves) are called piconets.

The Bluetooth specification also allows connecting two or more piconets together to form a scatternet, with some devices acting as a bridge by simultaneously playing the master role in one piconet and the slave role in another piconet. These devices have yet to come, though are supposed to appear within the next two years.

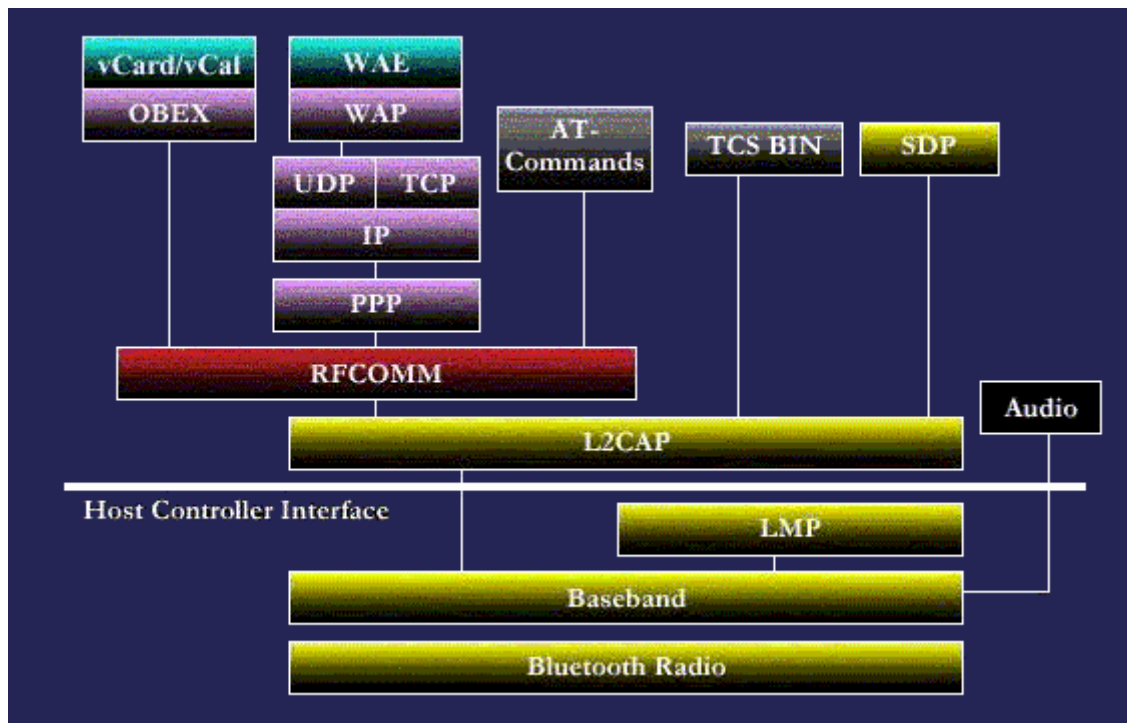
Any device may perform an "inquiry" to find other devices to which to connect, and any device can be configured to respond to such inquiries. However if the device trying to connect knows the address of the device it will always respond.

Every device has a unique 48 bit address. Every device also has a 24bit class identifier. This provides information on what kind of a device it is (Phone, Smartphone, Computer, Headset, etc), which will also be transmitted when other devices perform an inquiry. Devices also have friendly "Bluetooth names" which can be set by the user, and will appear when another user scans for devices and in lists of paired devices.

However since the Bluetooth address is permanent a pairing will be preserved even if the Bluetooth name is changed. Most phones have the Bluetooth name set to the manufacturer and model of the phone. Pairs of devices may establish a trusted relationship by learning (by user input) a shared secret

known as a "passkey". A device that wants to communicate only with a trusted device can cryptographically authenticate the identity of the other device. Trusted devices may also encrypt the data that they exchange over the air so that no one can listen in.

2.3 The Bluetooth Protocol Stack:



Chief Layers and their Functions:

1. **Radio** : The Radio layer defines the requirements for a Bluetooth transceiver operating in the 2.4 GHz ISM band.
2. **Baseband** : The Baseband layer describes the specification of the Bluetooth Link Controller (LC) which carries out the baseband protocols and other low-level link routines.
3. **LMP** : The Link Manager Protocol (LMP) is used by the Link Managers (on either side) for link set-up and control.
4. **HCI** : The Host Controller Interface (HCI) provides a command interface to the Baseband Link Controller and Link Manager, and access to hardware status and control registers.
5. **L2CAP** : Logical Link Control and Adaptation Protocol (L2CAP) supports higher level protocol multiplexing, packet segmentation and reassembly, and the conveying of quality of service information.

6. **RFCOMM** : The RFCOMM protocol provides emulation of serial ports over the L2CAP protocol. The protocol is based on the ETSI standard TS 07.10.
7. **SDP** : The Service Discovery Protocol (SDP) provides a means for applications to discover which services are provided by or available through a Bluetooth device. It also allows applications to determine the characteristics of those available services.

2.3.1 Radio Layer:

The Bluetooth Radio (layer) is the lowest defined layer of the Bluetooth specification. It defines the requirements of the Bluetooth transceiver device operating in the 2.4GHz ISM band.

- Accomplishes spectrum spreading by frequency hopping in 79 hops displaced by 1 MHz, starting at 2.402GHz and finishing at 2.480GHz
- Each device is classified into 3 power classes, Power Class 1, 2 & 3.
Power Class 1 : Long range (~ 100 m), max output power of 20dbm
Power Class 2 : Ordinary range (~ 10 m), max output power of 4dbm
Power Class 3 : Short range (~ 10 cm), max output power of 0 dbm
- Modulation characteristics : The Bluetooth radio module uses GFSK (Gaussian Frequency Shift Keying) where a binary one is represented by a positive frequency deviation and a binary zero by a negative frequency deviation.
- Sensitivity Level: The receiver must have a sensitivity level for which the bit error rate (BER) 0.1% is met. For Bluetooth this means an actual sensitivity level of -70dBm or better.
- Interference Performance: The interference performance on Co-channel and adjacent 1 MHz and 2 MHz are measured with the wanted signal 10 dB over the reference sensitivity level. On all other frequencies the wanted signal shall be 3 dB over the reference sensitivity level.
- Out-of-Band blocking: The Out of band blocking is measured with the wanted signal 3 dB over the reference sensitivity level. The interfering signal shall be a continuous wave signal. The BER shall be less than or equal to 0.1%.

2.3.2 Baseband :

The Baseband is the physical layer of the Bluetooth. It manages physical

channels and links apart from other services like error correction, data whitening, hop selection and Bluetooth security. The Baseband layer lies on top of the Bluetooth radio layer in the bluetooth stack.

- **Physical Channel :**

The channel is represented by a pseudo-random hopping sequence hopping through the 79 or 23 RF channels. Two or more Bluetooth devices using the same channel form a piconet. There is one master and one or more slave(s) in each piconet. The hopping sequence is unique for the piconet and is determined by the Bluetooth device address (BD_ADDR) of the master; the phase in the hopping sequence is determined by the Bluetooth clock of the master. The channel is divided into time slots where each slot corresponds to an RF hop frequency. Consecutive hops correspond to different RF hop frequencies.


- **Logical Channels :**

Bluetooth has five logical channels which can be used to transfer different types of information. LC (Control Channel) and LM (Link Manager) channels are used in the link level while UA, UI and US channels are used to carry asynchronous, isosynchronous and synchronous user information.

- It also takes care of Device Addressing

- **Packet types and Format :**

13 different packet types. Each packet consists of 3 entities, the access code (68/72 bits), the header (54 bits) , and the payload (0-2745 bits).



Scatternet :

Multiple piconets may cover the same area. Since each piconet has a different master, the piconets hop independently, each with their own channel hopping sequence and phase as determined by the respective master. In addition, the packets carried on the channels are preceded by different channel access codes as determined by the master device addresses. As more piconets are added, the probability of collisions increases; a graceful degradation of performance results as is common in frequency-hopping spread spectrum systems.

If multiple piconets cover the same area, a unit can participate in two or more overlaying piconets by applying time multiplexing. To participate on the proper channel, it should use the associated master device address and proper clock offset to obtain the correct phase. A Bluetooth unit can act as a slave in several piconets, but only as a master in a single piconet. A group of piconets in which connections consists between different piconets is called a scatternet.

Sometimes an existing master or slave may wish to swap roles (i.e a master-slave switch) , this can take place in two steps:

1. First a TDD switch of the considered master and slave, followed by a piconet switch of the both participants.
2. Then, if so desired, other slaves of the old piconet can be transferred to the new piconet.

When a unit have acknowledged the reception of the FHS packet, this unit uses the new piconet parameters defined by the new master and the piconet switch is completed.

2.3.3 Link Manager Protocol (LMP) :

The Link Manager carries out link setup, authentication, link configuration and other protocols. It discovers other remote LM's and communicates with them via the Link Manager Protocol (LMP). To perform its service provider role, the LM uses the services of the underlying Link Controller (LC).

The Link Manager Protocol essentially consists of a number of PDU (protocol Data Units), which are sent from one device to another, determined by the AM_ADDR in the packet header. LM PDUs are always sent as single-slot packets and the payload header is therefore one byte.

Functions :

1. Authentication
2. Key pairing
3. Change Link Key
4. Encryption
5. Clock offset Request
6. Slot offset Request
7. Switch of Master-Slave Role
8. Name Request
9. Detach mode : To close connection at any time by the master/slave.
10. Hold mode : No ACL packets will be transferred from the master. The hold mode is typically entered when there is no need to send data for a relatively long period of time.
11. Sniff and Park mode

- 12. Power control
- 13. QoS
- 14. Link supervision
- 15. Connection Establishment
- 16. Error Handling

2.3.4 Host Controller Interface (HCI):

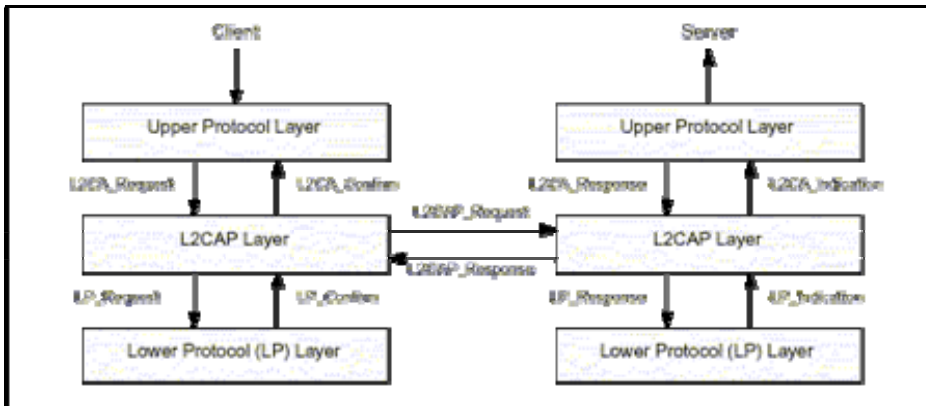
The HCI provides a command interface to the baseband controller and link manager, and access to hardware status and control registers. Essentially this interface provides a uniform method of accessing the Bluetooth baseband capabilities. The HCI exists across 3 sections, the Host - Transport Layer - Host Controller. Each of the sections has a different role to play in the HCI system.

2.3.5 L2CAP:

The Logical Link Control and Adaptation Layer Protocol (L2CAP) is layered over the Baseband Protocol and resides in the data link layer. L2CAP provides connection-oriented and connectionless data services to upper layer protocols with protocol multiplexing capability, segmentation and reassembly operation, and group abstractions. L2CAP permits higher level protocols and applications to transmit and receive L2CAP data packets up to 64 kilobytes in length.

Two link types are supported for the Baseband layer : Synchronous Connection-Oriented (SCO) links and Asynchronous Connection-Less (ACL) links. SCO links support real-time voice traffic using reserved bandwidth. ACL links support best effort traffic. The L2CAP Specification is defined for only ACL links and no support for SCO links is planned.

L2CAP state machine :



The figure above illustrates the events and actions performed by an implementation of the L2CAP layer. Client and Server simply represent the initiator of the request and the acceptor of the request respectively. An application-level Client would both initiate and accept requests.

2.3.6 RFCOMM:

The RFCOMM protocol provides emulation of serial ports over the L2CAP protocol.

Basically two device types exist that RFCOMM must accommodate.

- Type 1 Devices are communication end points such as computers and printers.
- Type 2 Devices are those that are part of the communication segment; e.g. modems.

Control Signals : RFCOMM emulates the 9 circuits of an RS-232 interface.


Functions :

1. Flow control
2. Port emulation and port proxy entities
3. Service registration and discovery
4. Low power modes

2.3.7 Service Discovery Protocol:

SDP is a simple protocol with minimal requirements on the underlying transport. It can function over a reliable packet transport (or even unreliable, if the client implements timeouts and repeats requests as necessary). SDP uses a request/response model where each transaction consists of one request protocol data unit (PDU) and one response PDU. However, the requests may potentially be pipelined and responses may potentially be returned out of order.

SDP uses a request/response model where each transaction consists of one request protocol data unit (PDU) and one response PDU. In the case where SDP is used with the Bluetooth L2CAP transport protocol, only one SDP request PDU per connection to a given SDP server may be outstanding at a given instant. In other words, a client must receive a response to each request before issuing another request on the same L2CAP connection. Limiting SDP to sending one unacknowledged request PDU provides a simple form of flow control.



2.4 Applications:

In order to use Bluetooth, a device must be able to interpret certain Bluetooth profiles. These define the possible applications. Some of these profiles are listed below.

1. Basic Imaging Profile (BIP)

This profile is designed for sending images between devices and includes the ability to resize, and convert images to make them suitable for the receiving device.

2. Basic Printing Profile (BPP)

This allows devices to send text, e-mails, vCards, or other items to printers based on print jobs.

3. Dial-up Networking Profile (DUN)

This profile provides a standard to access the Internet and other dial-up services over Bluetooth.

4. File Transfer Profile (FTP)

Provides access to the file system on another device. This includes support for getting folder listings, changing to different folders, getting files,

putting files and deleting files. It uses OBEX as a transport and is based on GOEP.

5. Personal Area Networking Profile (PAN)

This profile is intended to allow the use of Bluetooth Network Encapsulation Protocol on Layer 3 protocols for transport over a Bluetooth link.

3. IrDA

3.1 Introduction:

The Infrared Data Association (IrDA) defines physical specifications communications protocol standards for the short range exchange of data over infrared light, for uses such as personal area networks (PANs).

IrDA is a very short-range example of free-space optical communication. IrDA interfaces are used in palmtop computers and mobile phones. IrDA specifications include IrPHY, IrLAP, IrLMP, IrCOMM, Tiny TP, IrOBEX, and IrLAN. IrDA has now produced another standard, IrFM, for Infrared financial messaging also known as "Point & Pay". For the devices to communicate via IrDA they must have a direct line of sight.

3.2 Protocol Stack and Applications:

3.2.1 IrPHY

The mandatory IrPHY (Infrared Physical Layer Specification) is the lowest layer of the IrDA specifications. The most important specifications are:

- 1.range (1 m, low-power 0.1 m)
- 2.angle
- 3.speed (2.4 kbit/s to 16 Mbit/s)
- 4.modulation
- 5.Infrared window

IrDA transceivers broadcast infrared pulses in a cone that extends 15 to 30 degrees half angle off center. The IrDA physical specifications require that a minimum irradiance be maintained so that a signal is visible up to a meter away. Similarly, the specifications require that a maximum irradiance not be exceeded so that a receiver is not overwhelmed with brightness when a device comes in close. In practice, there are some devices on the market that do not reach one meter, while other devices may reach up to several meters. There are also devices that do not tolerate extreme closeness. The typical sweet spot for IrDA communications is from 5 cm to 60 cm away from a transceiver, in the center of the cone. IrDA data communications operate in half-duplex mode. The reason is quite simple. While transmitting, a device's receiver is blinded by the light of its own transmitter. Because of this, full duplex communication is not feasible. The two devices that communicate simulate full duplex communication by quickly turning the link around. The primary device controls the timing of the link, but both sides are bound to certain hard constraints and are encouraged to turn the link around as fast as possible. Transmission rates fall into three broad categories: SIR, MIR, and FIR. Serial Infrared (SIR) speeds cover those transmission speeds normally supported by an RS-232 port (9600 bit/s, 19.2

kbit/s, 38.4 kbit/s, 57.6 kbit/s, 115.2 kbit/s). Since the lowest common denominator for all devices is 9600 bit/s, all discovery and negotiation is performed at this baud rate. MIR (Medium Infrared) is not an official term, but is sometimes used to refer to speeds of 0.576 Mbit/s and 1.152 Mbit/s. Fast Infrared (FIR) is deemed an obsolete term by the IrDA physical specification, but is nonetheless in common usage to denote transmission at 4 Mbit/s. "FIR" is sometimes used to refer to all speeds above SIR. However, different encoding approaches are used by MIR and FIR, and different approaches are used to frame MIR and FIR packets. For that reason, these unofficial terms have sprung up to differentiate these two approaches. The future holds faster transmission speeds (currently referred to as Very Fast Infrared, or VFIR) which will support speeds up to 16 Mbit/s. There are (VFIR) infrared transceivers available such as the TFDU8108 operating from 9.6 kbit/s to 16 Mbit/s. UFIR (Ultra Fast Infrared) protocol that will provide up to 100 Mbit/s is also in development.

3.2.2 IrLAP

The mandatory IrLAP (Infrared Link Access Protocol) is the second layer of the IrDA specifications. It lies on top of the IrPHY layer and below the IrLMP layer. It represents the Data Link Layer of the OSI model. The most important specifications are:

1. Access control
2. Discovery of potential communication partners
3. Establishing of a reliable bidirectional connection
4. Negotiation of the Primary/Secondary device roles

On the IrLAP layer the communicating devices are divided into a Primary Device and one or more Secondary Devices. The Primary Device controls the Secondary Devices. Only if the Primary Device requests a Secondary Device to send it is allowed to do so.

3.2.3 IrLMP

The mandatory IrLMP (Infrared Link Management Protocol) is the third layer of the IrDA specifications. It can be broken down into two parts. First, the LM-MUX (Link Management Multiplexer) which lies on top of the IrLAP layer. Its most important achievements are:

1. Provides multiple logical channels
2. Allows change of Primary/Secondary devices

Second, the LM-IAS (Link Management Information Access Service), which provides a list, where service providers can register their services so other devices can access these services via querying the LM-IAS.

3.2.4 Tiny TP

The optional Tiny TP (Tiny Transport Protocol) lies on top of the IrLMP layer. It provides:

1. Transportation of large messages by Segmentation and Reassembly
2. Flow control by giving credits to every logical channel

3.2.5 IrCOMM

The optional IrCOMM (Infrared Communications Protocol) lets the infrared device act like either a serial or parallel port. It lies on top of the IrLMP layer.

3.2.6 IrOBEX

The optional IrOBEX (Infrared Object Exchange) provides the exchange of arbitrary data objects (e.g. vCard, vCalendar or even applications) between infrared devices. It lies on top of the Tiny TP protocol, so Tiny TP is mandatory for IrOBEX to work.

3.2.7 IrLAN

The optional IrLAN (Infrared Local Area Network) provides the possibility to connect an infrared device to a local area network. There are three possible methods:

1. Access Point
2. Peer to Peer
3. Hosted

As IrLAN lies on top of the Tiny TP protocol, the Tiny TP protocol must be implemented for IrLAN to work.

4. WiMax

4.1 Introduction:

WiMAX, an acronym that stands for Worldwide Interoperability for Microwave Access, is a certification mark for products that pass conformity and interoperability tests for the IEEE 802.16 standards. IEEE 802.16 is working group number 16 of IEEE 802, specialising in point-to-multipoint broadband wireless access.

The original WiMAX standard, IEEE 802.16, specifies WiMAX in the 10 to 66 GHz range. 802.16a added support for the 2 to 11 GHz range, of which many parts are already unlicensed internationally and only few still require domestic licenses. Most business interest will probably be in the 802.16a standard, as opposed to the higher frequencies. The WiMAX specification improves upon many of the limitations of the Wi-Fi standard by providing increased bandwidth and stronger encryption. It also aims to provide connectivity to network endpoints without direct line of sight in some circumstances. The details of performance under near-line of sight (NLOS) circumstances, however, are unclear, as they have yet to be demonstrated. It is commonly considered that spectrum under 5-6 GHz is needed to provide reasonable NLOS performance and cost effectiveness for PtM (point to multi-point) deployments.

4.2 Protocols and Technical Information:

4.2.1 Technical Information:

The core components of a WiMAX system are the subscriber station (SS) otherwise known as the CPE and the base station (BS). A BS and one or more SSs can form a cell with a point-to-multipoint (P2MP) structure. On air, the BS controls activity within the cell, including access to the medium by SS, allocations to achieve quality of service (QoS) and admission to the network based on network security mechanisms.

An 802.16-based system often uses fixed antenna at the subscriber station site. The antenna is mounted to the roof or an eave. Provisions such as adaptive-antenna systems (AAS) and sub-channelization are also supported optionally by the standard for enhanced link budget required for in-door installation. IEEE 802.16e sub-committee is currently working on extension to the standard required for mobility and support for the power limited SS terminals.

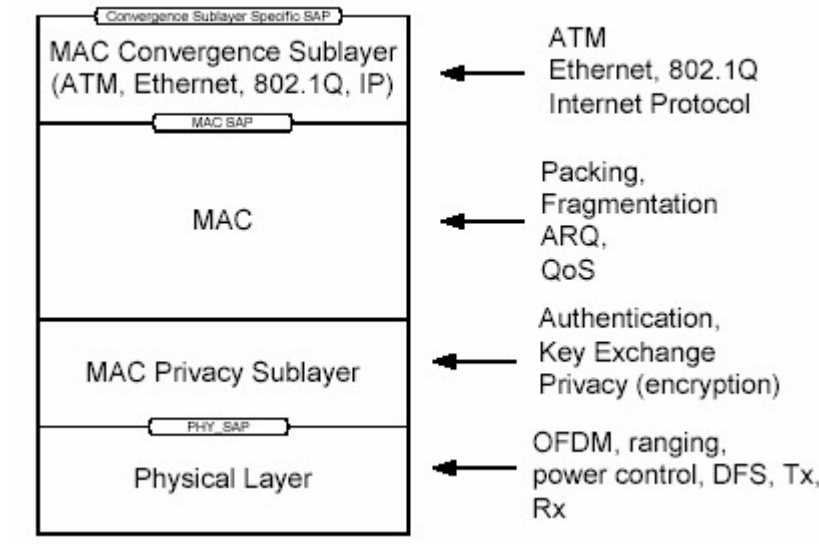
A BS typically uses either sectored/directional or omni-directional antennas. A fixed SS typically uses directional antenna while mobile or portable SS usually uses an omni-directional antenna.

Multiple BSs can be configured to form a cellular wireless network. When orthogonal frequency division multiplexing (OFDM) is used, the cell radius can

ideally reach up to 30 miles, however this requires a favorable channel environment and only the lowest data rate can be achieved. Practical cell sizes usually have a small radius of around 5 miles or less. The 802.16 standard also can be used in a point-to-point (P2P) or mesh topology, using pairs of directional antennas. This can be used to increase the effective range of the system relative to what can be achieved in P2MP mode.

Features :

WiMAX supports both time division duplex (TDD) and frequency division duplex (FDD) modes of operation on air, along with a range of channel bandwidths. The OFDM PHY mode, which is also known WirelessMAN-OFDM, is specified for use between 2 and 11 GHz.

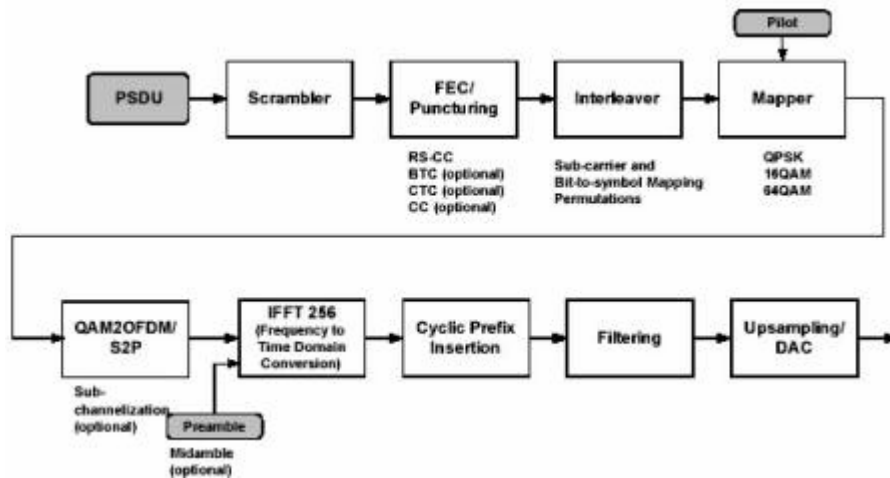


Profiles :

Identifier	Description
ProfM3_PMP	WirelessMAN-OFDM Basic packet PMP MAC profile
ProfM3_Mesh	profM3_Mesh WirelessMAN-OFDM Basic packet Mesh MAC profile
ProfP3_1.75	WirelessMAN-OFDM 1.75 MHz channel basic PHY profile
ProfP3_3.5	WirelessMAN-OFDM 3.5 MHz channel basic PHY profile
ProfP3_7	WirelessMAN-OFDM 7 MHz channel basic PHY profile
ProfP3_3	WirelessMAN-OFDM 3 MHz channel basic PHY profile
ProfP3_5.5	WirelessMAN-OFDM 5.5 MHz channel basic PHY profile
ProfP3_10	WirelessHUMAN-OFDM 10 MHz channel basic PHY profile

4.2.2 Protocols:

4.2.2.1 The OFDM Physical Layer :



Framing and Media Access :

On air transmission time is divided into frames. In the case of an FDD system, there are uplink (SS to BS) and downlink (BS to SS) subframes that are time aligned on separate uplink and downlink channels. In the case of a TDD system, each frame is divided up into a downlink subframe and an uplink subframe.

In both TDD and FDD modes, the length of the frame can vary (under the control of the BS scheduler) per frame. In TDD mode, the division point between uplink and downlink can also vary per frame, allowing asymmetric allocation of on air time between uplink and downlink if required.

The downlink frame format includes a preamble, a DL_MAP, a UL_MAP, and downlink slots. The DL_MAP is a directory of the slot locations within the downlink subframe. The UL_MAP is a directory of slot locations within the uplink subframe. It is through the DL_MAP and UL_MAP frame descriptors that the BS allocates access to the channel for both uplink and downlink.

The SS uses the DL_MAP to identify the location of MPDUs within the frame and listens to each of the MPDUs in turn, receiving those that match a connection ID targeted at that SS.

Uplink framing is more complex, since for best effort delivery and network entry, a contention-based multiple access scheme is required in order to mediate between SSeS that are simultaneously seeking access to the medium. Based on

the QoS service used for a connection, a connection may have either a guaranteed slot, may get access to a guaranteed slot on a per frame basis through polling from the BS or it may have to contend for uplink access on a contention basis in a multiple access (TDMA) slot.

Contention access takes place in slots set aside for the purpose in the uplink, the contention slot for initial ranging slots and the contention slot for bandwidth requests. Each of these slots is divided into minislots. SSES contending for access use a truncated binary exponential backoff algorithm to elect which mini slot to begin its transmission in.

The initial ranging contention slot is used as part of the network entry algorithm. An SS transmits a ranging request (RNG-REQ) packet in the initial ranging contention slot. The RNG-REQ packet has a long preamble, enabling the BS to better identify the timing of the received RNG-REQ packet. If the RNG-REQ is received, the BS responds a RNG-RSP (ranging response) giving timing and power adjustment information to the SS. The SS can then adjust the timing to account for transit delays and path loss of its transmissions such that the timing and power of the signal as received at the base station aligns with transmissions from other SSES.

The bandwidth request contention slot is used by SSES to contend for access to the channel. Bandwidth requests are transmitted into this slot. Once a bandwidth request has been received and granted, the SS may use non-contention slots allocated by the BS.

The BS dictates the length of the contention slots. The optimal length for either of the contention slots might change based on any of a variety of parameters such as the number of SSES, the number and type of QoS connections allocated and current activity levels.

4.2.2.2 The MAC layer :

In an 802.16 system, the MAC communicates using MAC protocol data units (MPDUs) that are carried by the PHY.

The generic MAC header (GMH) contains details of the MPDU. Principally the connection ID (CID) that defines the connection that this packet is servicing, the length of the frame and bits to qualify the presence of the CRC, sub headers and whether or not the payload is encrypted and if so, with which key.

A 32-bit CCITT standard CRC of the entire MPDU may be appended to the frame if required.

The payload can contain either a management message or transport data. Specific connections are set aside as management connections and these carry

management messages and not anything else. All other channels are transport channels that do not carry management messages.

A payload in a transport connection can contain a MAC service data unit (MSDU), fragments of MSDUs, aggregates of MSDUs, aggregates of fragments of MSDUs, bandwidth requests or retransmission requests according to the MAC rules on bandwidth requesting, fragmentation, packing and ARQ.

4.3 Applications:

- WiMax will allow interpenetration for broadband service provision of VoIP, video, and internet access - simultaneously.
- There is also interesting potential for interoperability of WiMax with cellular networks. WiMax antennas can "share" a cell tower without compromising the function of cellular arrays already in place.
- Another application under consideration is gaming. Sony and Microsoft are closely considering the addition of WiMax as a standard feature in their next generation game consoles, already anticipated as a sellout on arrival. This will allow gamers to create ad hoc networks other players with the same gear can connect to, play against each other, communicate, and share data.

5. Wi-Fi

5.1 Introduction:

Wi-Fi (sometimes written Wi-fi, WiFi, Wifi, wifi) is a trademark for set of product compatibility standards for wireless local area networks (WLANs). Wi-Fi, short for "Wireless Fidelity", was intended to allow mobile devices, such as laptop computers and personal digital assistants (PDAs) to connect to local area networks, but is now often used for Internet access and wireless VoIP phones. Desktop computers can use Wi-Fi too, allowing offices to be networked without expensive wiring. Many computers are sold today with Wi-Fi built-in, others require adding a Wi-Fi network card. Other devices, such as digital cameras, are sometimes equipped with Wi-Fi.

A person with a Wi-Fi-enabled device can connect to a local area network when near one of the network's access points. The connection is made by radio signals; there is no need to plug the device into the network. If the local area network is connected to the Internet, the Wi-Fi device can have Internet access as well. The geographical region covered by one or several access points is called a hotspot. The range of an access point varies. The access point built into a typical Wi-Fi home router might have a range of 45 m (150 ft) indoors and 90 m (300 ft) outdoors.

5.2 Specification:

Wi-Fi is based on the IEEE 802.11 specifications. There are currently four deployed 802.11 variations: 802.11a, 802.11b, 802.11g, and 802.11n. The b specification was used in the first Wi-Fi products. The g and n variants are the ones most often sold as of 2005.

<i>Specification</i>	<i>Speed</i>	<i>Frequency Band</i>	<i>Compatibility</i>
802.11a	11 Mbps	2.4 GHz	b
802.11b	54 Mbps	5 GHz	a
802.11g	54 Mbps	2.4 GHz	b,g
802.11n	100 Mbps	2.4 GHz	b,g,n

In most of the world, the frequencies used by Wi-Fi do not require user licenses from local regulators (eg, the Federal Communications Commission in the US). 802.11a equipment, using a higher frequency, has reduced range, all other things being equal.

The most widespread version of Wi-Fi in the US market today (based in IEEE 802.11b/g) operates in the 2,400 MHz to 2,483.50 MHz. It allows to operate in 11 channels (5 MHz each), centered on the following frequencies:

Channel 1 - 2,412 MHz;
Channel 2 - 2,417 MHz;
Channel 3 - 2,422 MHz;
Channel 4 - 2,427 MHz;
Channel 5 - 2,432 MHz;
Channel 6 - 2,437 MHz;
Channel 7 - 2,442 MHz;
Channel 8 - 2,447 MHz;
Channel 9 - 2,452 MHz;
Channel 10 - 2,457 MHz;
Channel 11 - 2,462 MHz

Europe, France and Spain have adopted their own allowed channels set, and Japan has also done so. In all areas, the maximum radio transmitter power and the maximum effective radiated power (essentially the power output at the antenna) are strictly limited. In the US, maximum transmitter power is 1 watt, and maximum effective radiated power is 4 watts; in Europe these limits are somewhat lower. An antenna which concentrates 1 watt of transmitter energy into 1/4 of an 'omnidirectional' sphere will achieve 4 watts of effective power. Most WiFi equipment (eg, PCMCIA or Cardbus cards for laptops, PCI cards for desktop equivalent computers, or standalone units often with other functions included) has transmitter power levels of between 15mw and perhaps 200mw, so antennas with some gain are permissible.

New standards beyond the 802.11 specifications are currently in the works and offer many enhancements, anywhere from longer range to greater transfer speeds. One example is 802.16 WiMAX, with a range of several miles and data rates of up to 70Mbps. 802.16a permits operation between 2 and 11 GHz, so there may eventually be some interoperability between 802.11 units and some 802.16a units.

5.3 Security:

WiFi equipment could be used to steal personal information (passwords, financial information, identity information, and so on) transmitted from Wi-Fi users, if sensible protections are not used.

The first and most commonly used wireless encryption standard, Wired Equivalent Privacy or WEP, has been shown to be easily breakable even when correctly configured. Most wireless products now on the market support the Wi-Fi Protected Access (WPA) encryption protocol, which is considered much stronger, though some older access points have to be replaced to support it. The adoption

of the 802.11i standard (marketed as WPA2) makes available a rather better security scheme — when properly configured. As of mid-2005, both Microsoft Windows XP and Mac OS X support WPA2, but on newer equipment only. While waiting for better standards to be available, many enterprises have chosen to deploy additional layers of encryption (such as VPNs) to protect against interception.

Some report that interference of a closed or encrypted access point with other open access points on the same or a neighboring channel can prevent access to the open access points by others in the area. This can pose a problem in high-density areas such as large apartment buildings where many residents are operating Wi-Fi access points.

5.4 Advantages and Disadvantages:

5.4.1 Wi-Fi vs. Cellular:

Some argue that Wi-Fi and related consumer technologies will replace cellular telephone networks such as 3G and GSM. The current generation of Wi-Fi still lacks roaming and authentication features (see 802.1x, SIM cards and RADIUS) and the limited range of Wi-Fi as well as the narrowness of the available spectrum are holding back its proliferation as 3G replacement.

However, the bandwidth and overall capabilities of Wi-Fi are already exceeding those once promised by 3G cellular telephone standards which lead to the use of the term 4G being used for Wi-Fi.

Many vendors are now selling mobile Internet products that link Wi-Fi and cellular radio system in a more or less transparent way to take advantage of the benefits of both systems. Future wireless systems are expected to routinely switch between a variety of radio systems.

The main difference between cellular and Wi-Fi is that the cellular system uses the licensed spectrum, and Wi-Fi is implemented in unlicensed bands. The economic basis for its implementation is therefore completely different. The success of Wi-Fi has made many people look to the unlicensed spectrum as the future of wireless access, rather than the spectrum licensed and controlled by large corporations.

5.4.2 Advantages:

- Unlike packet radio systems, Wi-Fi uses unlicensed radio spectrum and does not require regulatory approval for individual deployers.
- Allows LANs to be deployed without cabling, potentially reducing the costs of network deployment and expansion.

- Wi-Fi products are widely available in the market. Different brands of access points and client network interfaces are interoperable at a basic level of service.
- Competition amongst vendors has lowered prices considerably since their inception.
- Many Wi-Fi networks support roaming, in which a mobile client station such as a laptop computer can move from one access point to another as the user moves around a building or area.
- Many access points and network interfaces support various degrees of encryption to protect traffic from interception.
- Wi-Fi is a global set of standards. Unlike cellular carriers, the same Wi-Fi client works in different countries around the world.

5.4.3 Disadvantages:

- Use of the 2.4 GHz Wi-Fi band does not require a license in most of the world provided that one stays below the local regulatory limits and provided one accepts interference from other sources, including interference which causes your devices to no longer function.
- Legislation/regulation is not consistent worldwide; most of Europe allows for an additional 2 channels over those allowed for b and g; Japan has one more on top of that - and some countries, like Spain, prohibit use of the lower-numbered channels. Furthermore some countries, such as Italy, used to require a 'general authorization' for any WiFi used outside the owned premises; or required something akin to operator registration.
- The 802.11b and 802.11g flavors of Wi-Fi use the 2.4 GHz spectrum, which is crowded with other equipment such as Bluetooth devices, microwave ovens, cordless phones (900 MHz or 5.8 GHz are, therefore, alternative phone frequencies one can use to avoid interference if one has a Wi-Fi network), or video sender devices, among many others. This may cause a degradation in performance. Other devices which use these microwave frequencies can also cause degradation in performance.
- Power consumption is fairly high compared to other standards, making battery life and heat a concern.
- Free access points (or improperly configured access points) may be used by a hacker to anonymously initiate an attack that would be difficult to track beyond the owner of the access point.

6. IEEE 802.11

6.1 Introduction:

IEEE 802.11 or Wi-Fi denotes a set of Wireless LAN standards developed by working group 11 of the IEEE LAN/MAN Standards Committee (IEEE 802). The term is also used to refer to the original 802.11, which is now sometimes called "802.11legacy."

The 802.11 family currently includes six over-the-air modulation techniques that all use the same protocol, the most popular (and prolific) techniques are those defined by the a, b, and g amendments to the original standard; security was originally included, and was later enhanced via the 802.11i amendment. Other standards in the family (c–f, h–j, n) are service enhancement and extensions, or corrections to previous specifications. 802.11b was the first widely accepted wireless networking standard, followed (somewhat counterintuitively) by 802.11a and 802.11g.

802.11b and 802.11g standards use the unlicensed 2.4 gigahertz (GHz) band. The 802.11a standard uses the 5 GHz band. Operating in an unregulated frequency band, 802.11b and 802.11g equipment can incur interference from microwave ovens, cordless phones, and other appliances using the same 2.4 GHz band.

6.2 Various Technologies:

The original version of the standard IEEE 802.11 released in 1997 specifies two raw data rates of 1 and 2 megabits per second (Mbit/s) to be transmitted via infrared (IR) signals or in the Industrial Scientific Medical frequency band at 2.4 GHz. IR remains a part of the standard but has no actual implementations.

The original standard also defines Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) as the media access method. A significant percentage of the available raw channel capacity is sacrificed (via the CSMA/CA mechanisms) in order to improve the reliability of data transmissions under diverse and adverse environmental conditions.

6.2.1 802.11b:

802.11b has a maximum raw data rate of 11 Mbit/s and uses the same CSMA/CA media access method defined in the original standard. Due to the CSMA/CA protocol overhead, in practice the maximum 802.11b throughput that an application can achieve is about 5.9 Mbit/s over TCP and 7.1 Mbit/s over UDP.

802.11b products appeared on the market very quickly, since 802.11b is a direct extension of the DSSS (Direct-sequence spread spectrum) modulation technique defined in the original standard. Hence, chipsets and products were easily upgraded to support the 802.11b enhancements. The dramatic increase in throughput of 802.11b (compared to the original standard) along with substantial price reductions led to the rapid acceptance of 802.11b as the definitive wireless LAN technology.

802.11b is usually used in a point-to-multipoint configuration, wherein an access point communicates via an omni-directional antenna with one or more clients that are located in a coverage area around the access point. With high-gain external antennas, the protocol can also be used in fixed point-to-point arrangements, typically at ranges up to eight kilometers (km) although some report success at ranges up to 80–120 km where line of sight can be established. This is usually done in place of costly leased lines or very cumbersome microwave communications equipment.

802.11b cards can operate at 11 Mbit/s, but will scale back to 5.5, then 2, then 1 Mbit/s (a.k.a Adaptive Rate Selection), if signal quality becomes an issue. Since the lower data rates use less complex and more redundant methods of encoding the data, they are less susceptible to corruption due to interference and signal attenuation. Extensions have been made to the 802.11b protocol (e.g., channel bonding and burst transmission techniques) in order to increase speed to 22, 33, and 44 Mbit/s, but the extensions are proprietary and have not been endorsed by the IEEE. Many companies call enhanced versions "802.11b+". These extensions have been largely obviated by the development of 802.11g, which has data rates up to 54 Mbit/s and is backwards-compatible with 802.11b.

6.2.2 802.11a:

The 802.11a standard uses the same core protocol as the original standard, operates in 5 GHz band, and uses a 52-subcarrier orthogonal frequency-division multiplexing (OFDM) with a maximum raw data rate of 54 Mbit/s, which yields realistic net achievable throughput in the mid-20 Mbit/s. The data rate is reduced to 48, 36, 24, 18, 12, 9 then 6 Mbit/s if required. 802.11a has 12 non-overlapping channels, 8 dedicated to indoor and 4 to point to point. It is not interoperable with 802.11b, except if using equipment that implements both standards.

Since the 2.4 GHz band is heavily used, using the 5 GHz band gives 802.11a the advantage of less interference. However, this high carrier frequency also brings disadvantages. It restricts the use of 802.11a to almost line of sight, necessitating the use of more access points; it also means that 802.11a cannot penetrate as far as 802.11b since it is absorbed more readily, other things (such as power) being equal.

Of the 52 OFDM subcarriers, 48 are for data and 4 are pilot subcarriers with a carrier separation of 0.3125 MHz (20 MHz/64). Each of these subcarriers can be a BPSK, QPSK, 16-QAM or 64-QAM. The total bandwidth is 20 MHz with an occupied bandwidth of 16.6 MHz. Symbol duration is 4 microseconds with a guard interval of 0.8 microseconds. The actual generation and decoding of orthogonal components is done in baseband using DSP which is then upconverted to 5 GHz at the transmitter. Each of the subcarriers could be represented as a complex number. The time domain signal is generated by taking an Inverse Fast Fourier transform (IFFT). Correspondingly the receiver downconverts, samples at 20 MHz and does an FFT to retrieve the original coefficients. The advantages of using OFDM include reduced multipath effects in reception and increased spectral efficiency.

802.11a products started shipping in 2001, lagging 802.11b products due to the slow availability of the 5 GHz components needed to implement products. 802.11a was not widely adopted overall because 802.11b was already widely adopted, because of 802.11a's disadvantages, because of poor initial product implementations, making its range even shorter, and because of regulations. Manufacturers of 802.11a equipment responded to the lack of market success by improving the implementations (current-generation 802.11a technology has range characteristics much closer to those of 802.11b), and by making technology that can use more than one 802.11 standard. There are dual-band, or dual-mode or tri-mode cards that can automatically handle 802.11a and b, or a, b and g, as available. Similarly, there are mobile adapters and access points which can support all these standards simultaneously.

6.2.3 802.11g:

This flavor works in the 2.4 GHz band (like 802.11b) but operates at a maximum raw data rate of 54 Mbit/s, or about 24.7 Mbit/s net throughput like 802.11a. It is fully backwards compatible with b and uses the same frequencies. Details of making b and g work well together occupied much of the lingering technical process. In older networks, however, the presence of an 802.11b participant significantly reduces the speed of an 802.11g network.

While 802.11g held the promise of higher throughput, actual results were mitigated by a number of factors: conflict with 802.11b-only devices (see above), exposure to the same interference sources as 802.11b, limited channelization (only 3 fully non-overlapping channels like 802.11b) and the fact that the higher data rates of 802.11g are often more susceptible to interference than 802.11b, causing the 802.11g device to reduce the data rate to effectively the same rates used by 802.11b. The move to dual-mode/tri-mode products also carries with it economies of scale (e.g. single chip manufacturing). The use of dual-band/tri-mode products ensures the best possible throughput in any given environment.

A new proprietary feature called Super G is now integrated in certain access points. These can boost network speeds up to 108 Mbit/s by using channel bonding. This feature may interfere with other networks and may not support all b and g client cards. In addition, packet bursting techniques are also available in some chipsets and products which will also considerably increase speeds. Again, they may not be compatible with some equipment.

6.2.4 802.11n:

In January 2004 IEEE announced that it had formed a new 802.11 Task Group (TGn) to develop a new amendment to the 802.11 standard for local-area wireless networks. The real data throughput is estimated to reach a theoretical 540 Mbit/s (which may require an even higher raw data rate at the physical layer), and should be up to 10 times faster than 802.11a or 802.11g, and near 40 times faster than 802.11b. It is projected that 802.11n will also offer a better operating distance than current networks.

There are two competing proposals of the 802.11n standard, expected to be ratified: WWiSE (World-Wide Spectrum Efficiency), backed by companies including Broadcom, and TGn Sync backed by Intel and Philips.

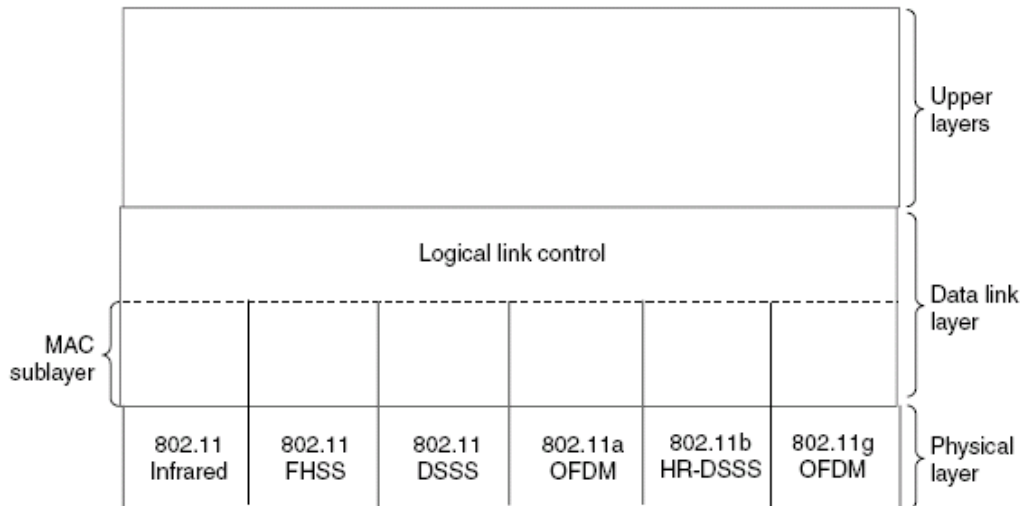
802.11n builds upon previous 802.11 standards by adding MIMO (multiple-input multiple-output) and orthogonal frequency-division multiplexing (OFDM). MIMO uses multiple transmitter and receiver antennas to allow for increased data throughput through spatial multiplexing and increased range by exploiting the spatial diversity, perhaps through coding schemes like Alamouti coding.

6.3 The 802.11 Protocol Stack:

The protocols used by all the 802 variants, including Ethernet, have a certain commonality of structure. The physical layer corresponds to the OSI physical layer fairly well, but the data link layer in all the 802 protocols is split into two or more sublayers. In 802.11, the MAC (Medium Access Control) sublayer determines how the channel is allocated, that is, who gets to transmit next. Above it is the LLC (Logical Link Control) sublayer, whose job it is to hide the differences between the different 802 variants and make them indistinguishable as far as the network layer is concerned.

The 1997 802.11 standard specifies three transmission techniques allowed in the physical layer. The infrared method uses much the same technology as television remote controls do. The other two use short-range radio, using techniques called FHSS and DSSS. Both of these use a part of the spectrum that does not require licensing (the 2.4-GHz ISM band). Radio-controlled garage door openers also use this piece of the spectrum, so your notebook computer may find itself in

competition with your garage door. Cordless telephones and microwave ovens also use this band. All of these techniques operate at 1 or 2 Mbps and at low enough power that they do not conflict too much. In 1999, two new techniques were introduced to achieve higher bandwidth. These are called OFDM and HRDSSS. They operate at up to 54 Mbps and 11 Mbps, respectively. In 2001, a second OFDM modulation was introduced, but in a different frequency band from the first one. Now we will examine each of them briefly.



6.3.1 The 802.11 Physical Layer:

Each of the five permitted transmission techniques makes it possible to send a MAC frame from one station to another. They differ, however, in the technology used and speeds achievable. The infrared option uses diffused (i.e., not line of sight) transmission at 0.85 or 0.95 microns. Two speeds are permitted: 1 Mbps and 2 Mbps. At 1 Mbps, an encoding scheme is used in which a group of 4 bits is encoded as a 16-bit codeword containing fifteen 0s and a single 1, using what is called **Gray code**. This code has the property that a small error in time synchronization leads to only a single bit error in the output. At 2 Mbps, the encoding takes 2 bits and produces a 4-bit codeword, also with only a single 1, that is one of 0001, 0010, 0100, or 1000. Infrared signals cannot penetrate walls, so cells in different rooms are well isolated from each other. Nevertheless, due to the low bandwidth (and the fact that sunlight swamps infrared signals), this is not a popular option.

FHSS (Frequency Hopping Spread Spectrum) uses 79 channels, each 1-MHz wide, starting at the low end of the 2.4-GHz ISM band. A pseudorandom **294 THE MEDIUM ACCESS CONTROL SUBLAYER CHAP. 4** number generator is used to produce the sequence of frequencies hopped to. As long as all stations use the same seed to the pseudorandom number generator and stay synchronized in time, they will hop to the same frequencies simultaneously. The

amount of time spent at each frequency, the **dwell time**, is an adjustable parameter, but must be less than 400 msec. FHSS' randomization provides a fair way to allocate spectrum in the unregulated ISM band. It also provides a modicum of security since an intruder who does not know the hopping sequence or dwell time cannot eavesdrop on transmissions. Over longer distances, multipath fading can be an issue, and FHSS offers good resistance to it. It is also relatively insensitive to radio interference, which makes it popular for building-to-building links. Its main disadvantage is its low bandwidth.

The third modulation method, **DSSS (Direct Sequence Spread Spectrum)**, is also restricted to 1 or 2 Mbps. Each bit is transmitted as 11 chips, using what is called a **Barker sequence**. It uses phase shift modulation at 1 Mbaud, transmitting 1 bit per baud when operating at 1 Mbps and 2 bits per baud when operating at 2 Mbps.

The first of the high-speed wireless LANs, **802.11a**, uses **OFDM (Orthogonal Frequency Division Multiplexing)** to deliver up to 54 Mbps in the wider 5-GHz ISM band. As the term FDM suggests, different frequencies are used—52 of them, 48 for data and 4 for synchronization—not unlike ADSL. Since transmissions are present on multiple frequencies at the same time, this technique is considered a form of spread spectrum, but different from both CDMA and FHSS. Splitting the signal into many narrow bands has some key advantages over using a single wide band, including better immunity to narrowband interference and the possibility of using noncontiguous bands. A complex encoding system is used, based on phase-shift modulation for speeds up to 18 Mbps and on QAM above that. At 54 Mbps, 216 data bits are encoded into 288-bit symbols. The technique has a good spectrum efficiency in terms of bits/Hz and good immunity to multipath fading.

HR-DSSS (High Rate Direct Sequence Spread Spectrum), another spread spectrum technique, uses 11 million chips/sec to achieve 11 Mbps in the 2.4-GHz band. It is called **802.11b** but is not a follow-up to 802.11a. Data rates supported by 802.11b are 1, 2, 5.5, and 11 Mbps. The two slow rates run at 1 Mbaud, with 1 and 2 bits per baud, respectively, using phase shift modulation (for compatibility with DSSS). The two faster rates run at 1.375 Mbaud, with 4 and 8 bits per baud, respectively, using **Walsh/Hadamard** codes. The data rate may be dynamically adapted during operation to achieve the optimum speed possible under current conditions of load and noise. In practice, the operating speed of 802.11b is nearly always 11 Mbps. Although 802.11b is slower than 802.11a, its range is about 7 times greater, which is more important in many situations.

An enhanced version of 802.11b, **802.11g**, was approved by IEEE in November 2001 after much politicking about whose patented technology it would use. It uses the OFDM modulation method of 802.11a but operates in the narrow 2.4-GHz ISM band along with 802.11b. In theory it can operate at up to 54 MBps. It is not yet clear whether this speed will be realized in practice.

6.3.2 The 802.11 MAC Sublayer Protocol:


The 802.11 MAC sublayer protocol is quite different from that of Ethernet due to the inherent complexity of the wireless environment compared to that of a wired system. With Ethernet, a station just waits until the ether goes silent and starts transmitting. If it does not receive a noise burst back within the first 64 bytes, the frame has almost assuredly been delivered correctly. With wireless, this situation does not hold.

To start with, there is the hidden station problem. Since not all stations are within radio range of each other, transmissions going on in one part of a cell may not be received elsewhere in the same cell. In this example, station *C* is transmitting to station *B*. If *A* senses the channel, it will not hear anything and falsely conclude that it may now start transmitting to *B*.

In addition, there is the inverse problem, the exposed station problem. Here *B* wants to send to *C* so it listens to the channel. When it hears a transmission, it falsely concludes that it may not send to *C*, even though *A* may be transmitting to *D* (not shown). In addition, most radios are half duplex, meaning that they cannot transmit and listen for noise bursts at the same time on a single frequency. As a result of these problems, 802.11 does not use CSMA/CD, as Ethernet does.

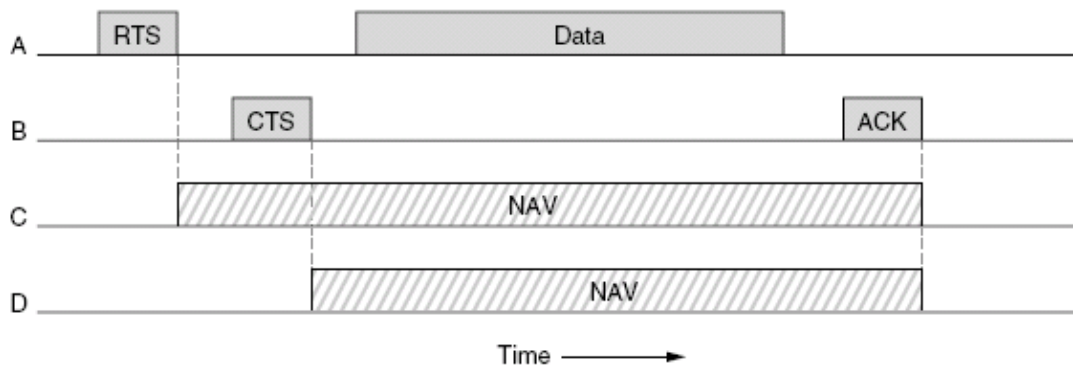
To deal with this problem, 802.11 supports two modes of operation. The first, called **DCF (Distributed Coordination Function)**, does not use any kind of central control (in that respect, similar to Ethernet). The other, called **PCF (Point Coordination Function)**, uses the base station to control all activity in its cell. All implementations must support DCF but PCF is optional. We will now discuss these two modes in turn.

When DCF is employed, 802.11 uses a protocol called **CSMA/CA (CSMA**



with Collision Avoidance). In this protocol, both physical channel sensing and virtual channel sensing are used. Two methods of operation are supported by CSMA/CA. In the first method, when a station wants to transmit, it senses the channel. If it is idle, it just starts transmitting. It does not sense the channel while transmitting but emits its entire frame, which may well be destroyed at the receiver due to interference there. If the channel is busy, the sender defers until it goes idle and then starts transmitting. If a collision occurs, the colliding stations wait a random time, using the Ethernet binary exponential backoff algorithm, and then try again later.

The other mode of CSMA/CA operation is based on MACAW and uses virtual channel sensing. In this example, *A* wants to send to *B*. *C* is a station within range of *A* (and possibly within range of *B*, but that does not matter). *D* is a station within range of *B* but not within range of *A*.



The protocol starts when *A* decides it wants to send data to *B*. It begins by sending an RTS frame to *B* to request permission to send it a frame. When *B* receives this request, it may decide to grant permission, in which case it sends a CTS frame back. Upon receipt of the CTS, *A* now sends its frame and starts an ACK timer. Upon correct receipt of the data frame, *B* responds with an ACK frame, terminating the exchange. If *A*'s ACK timer expires before the ACK gets back to it, the whole protocol is run again.

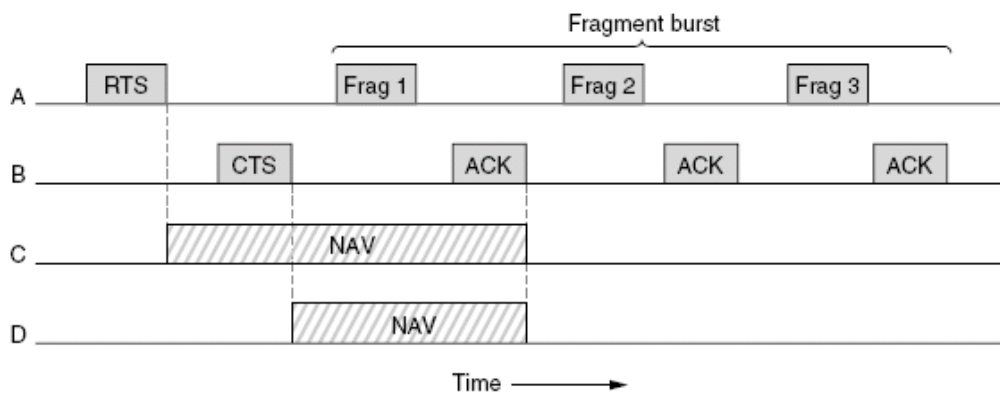
Now let us consider this exchange from the viewpoints of *C* and *D*. *C* is within range of *A*, so it may receive the RTS frame. If it does, it realizes that someone is going to send data soon, so for the good of all it desists from transmitting anything until the exchange is completed. From the information provided in the RTS request, it can estimate how long the sequence will take, including the final ACK, so it asserts a kind of virtual channel busy for itself, indicated by **NAV (Network Allocation Vector)**. *D* does not hear the RTS, but it does hear the CTS, so it also asserts the NAV signal for itself. Note that the NAV signals are not transmitted; they are just internal reminders to keep quiet for a certain period of time.

In contrast to wired networks, wireless networks are noisy and unreliable, in no small part due to microwave ovens, which also use the unlicensed ISM

bands. As a consequence, the probability of a frame making it through successfully decreases with frame length. If the probability of any bit being in error is p , then the probability of an n -bit frame being received entirely correctly is $(1 - p)^n$. For example, for $p = 10^{-4}$, the probability of receiving a full Ethernet frame (12,144 bits) correctly is less than 30%. If $p = 10^{-5}$, about one frame in 9 will be damaged. Even if $p = 10^{-6}$, over 1% of the frames will be damaged, which amounts to almost a dozen per second, and more if frames shorter than the maximum are used. In summary, if a frame is too long, it has very little chance of getting through undamaged and will probably have to be retransmitted.

To deal with the problem of noisy channels, 802.11 allows frames to be fragmented into smaller pieces, each with its own checksum. The fragments are individually numbered and acknowledged using a stop-and-wait protocol (i.e., the sender may not transmit fragment $k + 1$ until it has received the acknowledgment for fragment k). Once the channel has been acquired using RTS and CTS, multiple fragments can be sent in a row. sequence of fragments is called a **fragment burst**.

Fragmentation increases the throughput by restricting retransmissions to the bad fragments rather than the entire frame. The fragment size is not fixed by the standard but is a parameter of each cell and can be adjusted by the base station. The NAV mechanism keeps other stations quiet only until the next acknowledgement, but another mechanism (described below) is used to allow a whole fragment burst to be sent without interference. All of the above discussion applies to the 802.11 DCF mode. In this mode, there is no central control, and stations compete for air time, just as they do with Ethernet. The other allowed mode is PCF, in which the base station polls the other stations, asking them if they have any frames to send. Since transmission order is completely controlled by the base station in PCF mode, no collisions ever occur. The standard prescribes the mechanism for polling, but not the polling frequency, polling order, or even whether all stations need to get equal service.



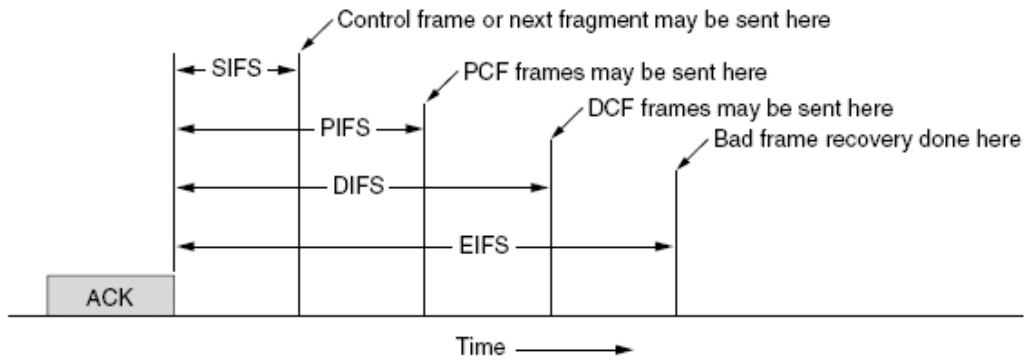
The basic mechanism is for the base station to broadcast a **beacon frame** periodically (10 to 100 times per second). The beacon frame contains system parameters, such as hopping sequences and dwell times (for FHSS), clock

synchronization, etc. It also invites new stations to sign up for polling service. Once a station has signed up for polling service at a certain rate, it is effectively guaranteed a certain fraction of the bandwidth, thus making it possible to give quality-of-service guarantees. Battery life is always an issue with mobile wireless devices, so 802.11 pays attention to the issue of power management. In particular, the base station can direct a mobile station to go into sleep state until explicitly awakened by the base station or the user. Having told a station to go to sleep, however, means that the base station has the responsibility for buffering any frames directed at it while the mobile station is asleep. These can be collected later.

PCF and DCF can coexist within one cell. At first it might seem impossible to have central control and distributed control operating at the same time, but 802.11 provides a way to achieve this goal. It works by carefully defining the inter-frame time interval. After a frame has been sent, a certain amount of dead time is required before any station may send a frame. Four different intervals are defined, each for a specific purpose. The four intervals are depicted in the figure.

The shortest interval is **SIFS (Short InterFrame Spacing)**. It is used to allow the parties in a single dialog the chance to go first. This includes letting the receiver send a CTS to respond to an RTS, letting the receiver send an ACK for a fragment or full data frame, and letting the sender of a fragment burst transmit the next fragment without having to send an RTS again. There is always exactly one station that is entitled to respond after a SIFS interval. If it fails to make use of its chance and a time **PIFS (PCF InterFrame Spacing)** elapses, the base station may send a beacon frame or poll frame. This mechanism allows a station sending a data frame or fragment sequence to finish its frame without anyone else getting in the way, but gives the base station a chance to grab the channel when the previous sender is done without having to compete with eager users.

If the base station has nothing to say and a time **DIFS (DCF InterFrame Spacing)** elapses, any station may attempt to acquire the channel to send a new frame. The usual contention rules apply, and binary exponential backoff may be needed if a collision occurs. The last time interval, **EIFS (Extended InterFrame Spacing)**, is used only by a station that has just received a bad or unknown frame to report the bad frame. The idea of giving this event the lowest priority is that since the receiver may have no idea of what is going on, it should wait a substantial time to avoid interfering with an ongoing dialog between two stations.



6.4 IEEE 802.11 list of Standards:

- IEEE 802.11 - The original 1 Mbit/s and 2 Mbit/s, 2.4 GHz RF and IR standard
- IEEE 802.11a - 54 Mbit/s, 5 GHz standard (1999, shipping products in 2001)
- IEEE 802.11b - Enhancements to 802.11 to support 5.5 and 11 Mbit/s (1999)
- IEEE 802.11d - International (country-to-country) roaming extensions
- IEEE 802.11e - Enhancements: QoS, including packet bursting
- IEEE 802.11F - Inter-Access Point Protocol (IAPP)
- IEEE 802.11g - 54 Mbit/s, 2.4 GHz standard (backwards compatible with b) (2003)
- IEEE 802.11h - 5 GHz spectrum, Dynamic Channel/Frequency Selection (DCS/DFS) and Transmit Power Control (TPC) for European compatibility
- IEEE 802.11i (ratified 24 June 2004) - Enhanced security
- IEEE 802.11j - Extensions for Japan
- IEEE 802.11k - Radio resource measurement enhancements
- IEEE 802.11n - Higher throughput improvements
- IEEE 802.11p - WAVE - Wireless Access for the Vehicular Environment (such as ambulances and passenger cars)
- IEEE 802.11r - Fast roaming
- IEEE 802.11s - Wireless mesh networking
- IEEE 802.11T - Wireless Performance Prediction (WPP) - test methods and metrics
- IEEE 802.11u - Interworking with non-802 networks (e.g., cellular)
- IEEE 802.11v - Wireless network management
- IEEE 802.11w - Protected Management Frames

7. Firewire

7.1 Introduction:

FireWire (also known as i.Link or IEEE 1394) is a personal computer and digital video serial bus interface standard offering high-speed communications and isochronous real-time data services. FireWire can be considered a successor technology to the obsolescent SCSI Parallel Interface. It is hub able to 63 ports (as opposed to USB, which can be hubbed up to 128), FireWire hubs are typically more costly than their USB equivalents.

Almost all modern digital camcorders have included this connection since 1995. All Macintosh computers currently produced have built-in FireWire ports, as do all Sony PCs and many PCs intended for home or professional audio/video use. FireWire is also used on the Apple iPod music player, permitting new tracks to be uploaded in a few seconds and also for the battery to be recharged concurrently with one cable.

7.2 Standards and Versions:

FireWire was developed primarily by Apple Computer in the 1990s, after work defining a slower version of the interface by the IEEE 1394 working committee in the 1980s. IEEE proposed the standard as a serial replacement for the SCSI bus. Apple's development was completed in 1995. It is defined in IEEE standard 1394 which is currently a composite of three documents: the original IEEE Std. 1394-1995, the IEEE Std. 1394a-2000 amendment, and the IEEE Std. 1394b-2002 amendment. Sony's implementation of the system is known as i.Link, and uses only the four signal pins, discarding the two pins that provide power to the device in favor of a separate power connector on Sony's i.Link products.

The system is commonly used for connection of data storage devices and digital video cameras, but is also popular in industrial systems for machine vision and professional audio systems. It is used instead of the more common USB due to its faster effective speed, higher power distribution capabilities, and because it does not need a computer host. Perhaps more importantly, since the advent of USB2, FireWire makes full use of all SCSI capabilities and supports delivery of data with known, low lag (isochronous), important for audio or video editing.

However, the small royalty that Apple Computer and other patent holders have initially demanded from users of FireWire (\$0.25 per end-user system) and the more expensive hardware needed to implement it (\$1–\$2) has prevented FireWire from displacing USB in low-end mass-market computer peripherals where cost of product is a major constraint.

FireWire can connect together up to 63 peripherals in an acyclic network structure (hubs, as opposed to SCSI's linear structure). It allows peer-to-peer device communication, such as communication between a scanner and a printer, to take place without using system memory or the CPU. FireWire also supports multiple hosts per bus, and IP networks can be formed through software between FireWire-linked computers (as opposed to USB, where a specialist chipset is needed for a USB-USB link between two computers, effectively resulting in the need for a special, expensive cable, whereas FireWire requires only a cable with the correct number of pins on either end (normally 6)). It is designed to support plug-and-play and hot swapping. Its six-wire cable is not only more convenient than SCSI cables but can supply up to 45 watts of power per port, allowing moderate-consumption devices to operate without a separate power cord. The Sony-inspired i.Link usually omits the power part of the cable/connector system and only uses a 4-pin connector.

FireWire 400 can transfer data between devices at 100, 200, or 400 Mbit/s data rates (actually 98.304, 196.608, or 393.216 Mbit/s, but commonly referred to as S100, S200, and S400). Although USB2 claims to be capable of higher speeds (480mb/s), FireWire is, in practice, faster. Cable length is limited to 4.5 metres but up to 16 cables can be daisy chained yielding a total length of 72 meters under the specification. FireWire 800 (Apple's name for the 9-pin "S800 bilingual" version of the IEEE1394b standard) was introduced commercially by Apple in 2003, allows an increase to 786.432 Mbit/s with backwards compatibility to the slower rates and 6-pin connectors of FireWire 400.

The full IEEE 1394b specification supports optical connections up to 100 metres in length and data rates all the way to 3.2 Gbit/s. Standard category-5 unshielded twisted pair supports 100 metres at S100, and the new p1394c technology goes all the way to S800. The original 1394 and 1394a standards used data/strobe (D/S) encoding (called legacy mode) on the signal wires, while 1394b adds a data encoding scheme called 8B10B (also referred to as beta mode). With this new technology, FireWire, which was arguably already slightly faster, is now substantially faster than Hi-Speed USB.

FireWire devices implement the ISO/IEC 13213 "configuration ROM" model for device configuration and identification, to provide plug-and-play capability. All FireWire devices are identified by an IEEE EUI-64 unique identifier (an extension of the 48-bit Ethernet MAC address format) in addition to well-known codes indicating the type of device and protocols it supports.

7.3 Security Issues:

Devices on a FireWire bus can communicate by direct memory access, where a device can use hardware to map internal memory to FireWire's "Physical Memory Space". The SBP (serial bus protocol) used by FireWire disk drives use this capability to minimize interrupts and buffer copies. In SBP, the initiator

(controlling device) sends a request by remotely writing a command into a specified area of the target's FireWire address space. This command usually includes buffer addresses in the initiator's FireWire "Physical Address Space", which the target is supposed to use for moving I/O data to and from the initiator.

On many implementations, particularly those like PCs and Macintoshes using the popular OHCI interface, the mapping between the FireWire "Physical Memory Space" and device physical memory is done in hardware, without operating-system intervention. While this enables extremely high-speed and low-latency communication between data sources and sinks without unnecessary copying (such as between a video camera and a software video recording application, or between a disk drive and the application buffers), this can also be a security risk if untrustworthy devices are attached to the bus. For this reason, high-security installations will typically either purchase newer machines that map a virtual memory space to the FireWire "Physical Memory Space" (such as a G5 Macintosh, or any Sun workstation), disable the OHCI hardware mapping between FireWire and device memory, physically disable the entire FireWire interface, or do not have FireWire at all.

This feature can also be used to debug a machine whose operating system has crashed, and in some systems for remote-console operations.

8. Home RF

- Home RF is a PAN standard
- It is a [wireless](#) networking specification that uses direct sequence spread spectrum (DSSS) in the 2.4 [GHz](#) frequency band
- Can achieve a maximum of 10 Mbit/s throughput and its nodes can travel within a 50 meter range of an access point while remaining connected to the PAN
- Both traditional telephone signals and data signals can be exchanged over the same wireless network
- But Home RF is obsolete and no group is developing the standard further.

References

1. Computer Networks, Tanenbaum
2. A Technical Tutorial on IEEE 802.11 by Breeze.com
3. Wireless communications tutorials from palowireless.com
4. Online encyclopedias (Wikipedia)