

BSc (Hons) Computer Science

University of Portsmouth

Third Year

Advanced Networks

M21279

Semester 1

Hugh Baldwin

hugh.baldwin@myport.ac.uk

Contents

1	Lecture - Introduction	2
2	Lecture - Signal Encoding	3
3	Seminar - Encoding Exercises	6
4	Lecture - Signal Encoding II	7
5	Seminar - Encoding Exercises II	10
6	Lecture - Security	11
7	Lecture - Local Area Networks	14
8	Lecture - Wireless Local Area Networks	16
9	Lecture - Bluetooth Networks	19
10	Lecture - Cellular Networks	22
11	Lecture - Software-Defined Networks (SDN)	24
12	Lecture - Satellite Networks	26

Lecture - Introduction

11:00

30/09/24

Asim Ali

Admin

- Lectures given by Dr Asim Ali
- Office - BK 2.20
 - Monday 1300-1400
 - Thursday 1100-1200
- 2 hour lecture every week (PK 2.01)
- 1 hour seminar every week (AG 1.03)

The module information is available at the URL below:

<https://course-module-catalog.port.ac.uk/#/moduleDetail/M21279/2024%2F25>

Assessments

This module is assessed by one exam and one piece of coursework as below

- 60%, 90 minute Exam covering LO1,2,3&4 - TB1 assessment period (January)
- 40%, Group coursework covering LO5 - 2-4 members *or* individual (but not recommended)
 - 5 pages maximum
 - Template provided
 - Previous samples provided

Lecture - Signal Encoding

11:30

30/09/24

Asim Ali

Signals

(In this context) Signals are typically electromagnetic waves that carry information using some method of encoding. This means that the 'shape' of the signal is changed to represent information. There are three main types of signal–

- Analogue Signals, which vary continuously and smoothly over time
- Digital Signals, which have only 2 levels and change near instantly
- Discrete Signals, which have 2 or more levels and change near instantly

When information is transmitted, it is often converted between many different signals before it reaches the destination. For example, if you have ADSL internet, your computer sends the data as a digital signal to the modem, which is then converted to an analogue signal before it is transmitted across a POTS network (Plain Old Telephone System), then the receiver's modem converts it back to a digital signal before it is sent on to the receiver.

Digital Signal Encoding

There are many methods of encoding a digital signal, but some of the most common are–

- Non-return to Zero Level (NRZ-L)
 - 0 – High voltage level
 - 1 – Low voltage level
- Non-return to Zero Inverted (NRZ-I)
 - 0 – No transition (high-low or low-high voltage level) on clock pulse
 - 1 – Transition on clock pulse
- Bipolar-AMI
 - 0 – Zero voltage
 - 1 – Alternating positive or negative voltage level
- MLT-3 (Multi-Level Transmit 3)
 - 0 – Remain at the same voltage level
 - 1 – Transition to the next voltage level
 - Uses 3 voltage levels, named +1, 0 & –1 but can be any arbitrary voltages
- Manchester
 - 0 – Transition from high-low voltage level in the middle of the clock interval
 - 1 – Transition from low-high voltage level
- Differential Manchester

- 0 – Transition (high-low or low-high voltage level) on clock pulse
- 1 – No transition on clock pulse
- Always transitions in the middle of the clock interval

In this context, the high and low voltage levels can be whatever you like, but are typically either a positive and negative voltage of the same magnitude, or are a positive voltage and zero. The most common voltages used for signalling are 3.3, 5 and 12 volts. In the case of Bipolar-AMI, the 3 voltage levels can be arbitrary, and don't have to be, for example +5, 0 and -5 volts.

Diagrams

Below is the binary string 01100110 encoded using all of the above encoding schemes. This assumes that the previous state before the transmission starts is always low, and that each vertical dotted line represents a clock pulse.




Figure 2.1: Non-return to Zero Level (NRZ-L)




Figure 2.2: Non-return to Zero Inverted (NRZ-I)




Figure 2.3: Bipolar-AMI




Figure 2.4: Multi-Level Transmit 3 (MLT-3)




Figure 2.5: Manchester




Figure 2.6: Differential Manchester

Carrier Waves and Modulation

A carrier wave is a continuous waveform that, on it's own, carries no information. This carrier wave is then modified by another signal to convey information. This modification can be of it's amplitude, frequency, phase, or a combination of all 3. This is known as modulation, hence the names AM (Amplitude Modulation) and FM (Frequency Modulation) for the two types of analogue radio.

Digital-to-Analogue Encoding

There are several methods of encoding a digital signal on an analogue transmission medium. In the case of sending digital data over a POTS network, the signal is encoded onto a carrier wave in the range of 300-3400Hz using a MoDem (Modulator-Demodulator). The main methods of encoding are-

- Amplitude-Shift keying (ASK)
- Frequency-Shift keying (FSK)
 - Binary FSK (BFSK)
 - Multiple FSK (MFSK)
- Phase-Shift keying (PSK)
 - Binary PSK – 1 = A sine wave, 0 = The same sine wave shifted by 180°
 - Differential PSK (DPSK) – 1 = The previous sine wave shifted by 180° , 0 = The previous sine wave
 - Multiple-level PSK
- Quadratic AM (Combination of ASK and PSK)

The sender and receiver may use different frequencies as to allow full-duplex data transmission on the same physical medium or radio channel. Full-duplex meaning full-speed transmission in both directions simultaneously.

Diagrams




Figure 2.7: Amplitude-Shift Keying




Figure 2.8: Binary Frequency-Shift Keying




Figure 2.9: Binary Phase-Shift Keying




Figure 2.10: Differential Phase-Shift Keying

Seminar - Encoding Exercises

12:00

03/10/24

Asim Ali

Figure 3.1: Encode the binary data 11000011 with the MLT-3 encoding scheme




Figure 3.2: Encode the binary data 11000011 with the Non-Differential Manchester encoding scheme





Figure 3.3: Encode the binary data 11000011 with the Differential Manchester encoding scheme



Lecture - Signal Encoding II

11:00

07/10/24

Asim Ali

Transmission Media

There are many different mediums which can be used to transmit information. In the case of networks, there are two main types, guided and unguided media.

- Guided Media
 - The signal is contained within a medium, such as
 - Copper cable (twisted pair, coaxial, etc)
 - Fibre optics
- Unguided Media
 - The signal is sent out without containment, through the air
 - This does not mean that the signal is not directional, microwave networks are often highly directional to reduce signal drop-off and interference
 - Microwave, Radio, Cellular and Satellite

Signal Properties

- **Data Rate** – The speed in bits per second (bps or b/s) at which data can be communicated
- **Error** – The reception of a 1 when 0 was transmitted, or vice versa
- **Error Rate** – The rate at which errors occur, usually as a ratio or percentage
- **Frequency Bandwidth** – The difference between the upper and lower frequency in a continuous frequency band
- **Channel Capacity** – The maximum rate at which data can be transmitted through a channel
- **Signal-to-Noise Ratio** – The ratio of the signal power to noise power in decibels (dB)

Interference and Noise

Definition

Interference is when two signals are added together. There are two main types of interference – constructive and destructive. When two signals combine to create a signal with a higher amplitude, this is constructive interference. If the amplitude is reduced, this is destructive interference.

Definition

Noise is any other signal which interferes with the desired signal along the medium of transmission. Typically, any noise in a signal will make it more difficult to decode, since the signal may be reduced in amplitude or shifted in phase.




Figure 4.1: Constructive interference

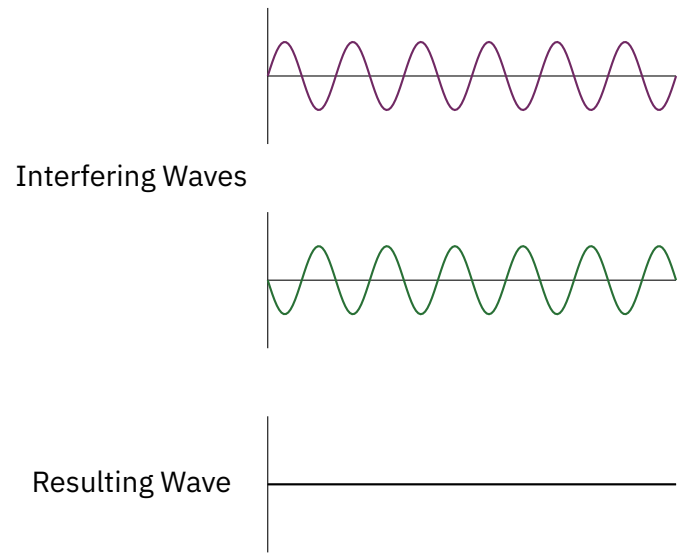


Figure 4.2: Destructive interference

Securing a Signal

When a signal is sent through any media, if a third party has access to the same media, they can intercept, block and even fake signals. This is because there is no inherent security in an electromagnetic wave, so we need to find ways of securing it.

Spread Spectrum

If we were to transmit a signal using a wide range of the spectrum, rather than a single frequency, it makes it harder for a third party to entirely block the signal. This method is often used in military applications, as well as wired and wireless networks. Three common techniques are–

- Frequency Hopping Spread Spectrum (FHSS)
 - Data is sent using one of the previously discussed analogue transmission techniques, but the carrier frequency is changed rapidly across a wide spectrum, using a code or pattern known by both the transmitter and receiver
 - This makes it harder to intercept or block transmissions, as long as the pattern of frequency hopping is not known
 - It is also very resilient to narrowband interference (interference over a small band of frequencies)
- Direct Sequence Spread Spectrum (DSSS)
 - Each bit of the data is represented by multiple bits of the transmitted signal, using a spreading code
 - The bits in the PN sequence are known as chips, and the sequence of them the chip sequence
 - One technique is to combine the data with the spreading code bit stream using an exclusive-or
- Code Division Multiple Access (CDMA)

Multiplexing

Definition

Multiplexing is sending multiple data streams or signals over a single transmission medium.

Frequency Division

With frequency division multiplexing, you have a transmission medium which can send data over a wide spectrum of frequencies, and multiple signals you need to send at the same time. If you modulate the different signals using a different carrier frequency for each, and leave a suitable gap between them, you can then combine the waves together to create a single composite wave which encodes the data of all the streams. The receiver then needs to know only the frequency of the signal they need to receive, and pass the composite wave through a filter for that specific frequency to recover the signal intended for them.

Time Division

With time division multiplexing, each of the signals gets a certain share of time in which it can be transmitted. This often means that the data needs to be buffered on both ends, so may be more costly to implement. This also greatly limits the bandwidth available to each signal, as it would be split between all of the signals on the same channel.

Code Division Multiple Access (CDMA)

CDMA is a multiplexing technique which can be used with a spread-spectrum signal. Given a data stream of bit rate R , we assign each bit a unique user code of n according to a Walsh matrix. If a user, k , sends a 1, the transmitter sends the chip code ck , and if they send a 0, the transmitter sends the inverse of that chip code, $c'k$. The chip codes of all users will add up to a bipolar signal, D . The receiver decodes the signal for a specific user with a function, by taking the cartesian product of D and the specific chip code ck , $D \times ck$. If $D \times ck = n$ then a 1 is received, and if $D \times ck = -n$ then a 0 is received.

Example

We have 4 users, with chip codes as follows

- A: $(-1, -1, -1, -1)$
- B: $(-1, +1, -1, +1)$
- C: $(-1, -1, +1, +1)$
- D: $(-1, +1, +1, -1)$

If the users were to send 1,1,1,0 respectively, then we would add up the signals such that $A + B + C + D' = (-1, -1, -1, -1) + (-1, +1, -1, +1) + (-1, -1, +1, +1) + (+1, -1, -1, +1) = (-2, -2, -2, +2)$

Then, the receivers filter out each individual signal by multiplying the received signal by the chip code of the signal they're interested in. To get back the signal for A, you would do the following $(-2, -2, -2, +2) \times (-1, -1, -1, -1) = (-2 \times -1) + (-2 \times -1) + (-2 \times -1) + (+2 \times -1) = 2 + 2 + 2 - 2 = 4$, which shows that A was sending a 1.

Seminar - Encoding Exercises II

12:00

10/10/24

Asim Ali




Figure 5.1: BFSK representation of the binary number 11000011




Figure 5.2: BPSK representation of the binary number 11000011




Figure 5.3: DPSK representation of the binary number 11000011

INPUT	1				0				1				1			
PN Stream	0	1	1	0	1	0	0	1	0	1	1	0	1	0	1	1
Transmitted	1	0	0	1	1	0	0	1	1	0	0	1	0	1	0	0
Received	1	0	0	1	1	0	0	1	1	0	0	1	0	1	0	0
PN Stream	0	1	1	0	1	0	0	1	0	1	1	0	1	0	1	1
OUTPUT	1				0				1				1			

Figure 5.4: DSSS encoding and decoding of the binary number 1001

Lecture - Security

11:00

14/10/24

Asim Ali

Security Factors

There are three main security factors we need to consider–

- Confidentiality
 - Only authorised parties should be able to read the information
- Integrity
 - Only authorised parties should be able to modify, create and delete information
- Availability
 - Authorised parties should always have access to the information

Types of Attackers

- Passive Attackers
 - The attacker only listens to the messages, so they can only read the information
- Active Attackers
 - The attacker listens, modifies and sends messages, so they can read the information, modify it or fabricate it entirely

To overcome both of these types of attacker, you can encrypt the data, meaning that it is unreadable by any third parties, and if someone attempts to modify or spoof a message, it will not be encrypted with the same key and so the receiver will know it was not sent by the correct person.

Symmetric (Shared Key) Encryption

Symmetric key encryption requires the sender and receiver to already have a shared key, which will then be used to both encrypt and decrypt the messages.

DES (Data Encryption Standard)

Plaintext is split into 64-bit blocks. The key is 56-bits long and 16 subkeys are generated for 16 rounds of cryptography. The subkeys are then used in reverse order to decrypt the data. Because this quickly became insecure, Triple-DES was created as a new standard, which is as literal as the name suggests. You simply encrypt the data three times with three different keys, then decrypt 3 times in reverse order.

AES (Advanced Encryption Standard)

Plaintext is split into 128-bit blocks. The key is 128,192 or 256-bits long, giving the algorithm the name AES-128/192/256. AES is a much more advanced algorithm than DES, and it was created with the express purpose of replacing DES

Confidentiality

Because the messages are encrypted, they cannot be read or examined by a third party. This does rely upon the keys being shared ahead of time, in a secure way such that no third parties know what the key is.

Asymmetric (Public Key) Encryption

In asymmetric encryption, the sender and receiver have two different but related keys. The sender uses the public key of the receiver to encrypt the message, and the receiver uses their own private key to decrypt the message.

RSA (Rivest-Shamir-Adleman)

1. Select two large primes p and q (the larger the better)
2. $n = pq$ and $z = (p - 1)(q - 1)$
3. Select a number relatively prime (such that they have no common divisors) to z and call it d
4. Find e such that $e * d = 1 \pmod{z}$
5. The public key is then (e, n) and the private key (d, n)
6. To encrypt a message, convert the plaintext M into an integer m such that $0 \leq m < n$, using a known and reversible padding scheme. Then compute the ciphertext c such as $c = m^e \pmod{n}$
7. To then decrypt the message, you use the private key d by computing $c^d = (m^e)^d = m \pmod{n}$. Then given m , you recover the message M by reversing the padding scheme

Figure 6.1: Encrypt then decrypt the string 'PORT' with the public key (3, 33) and the private key (7, 33)

Char	Numeric	P^3	$P^3 \text{ Mod } 33$	C^7	$C^7 \text{ mod } 33$	Char
	P	16	4096	4	16384	P
	O	15	3375	9	4782969	O
	R	18	5832	24	4586471424	R
	T	20	8000	14	105413504	T
Sender				Receiver		

Confidentiality

Because the messages are encrypted, they cannot be read or examined by a third party. Since the public key can only be used to encrypt messages (theoretically it can be used to determine the private key but this is effectively impossible with current computational power), it can be sent out on the internet for anyone to use to encrypt messages for the receiver. As long as the private key is kept secret by the receiver, the messages are perfectly confidential.

Ensuring Integrity

To make sure that a message was sent by the correct person, and without being modified we can use either signatures, message digests or both.

Digital Signatures

When making a digital signature, the inverse of public-key cryptography is done. This means that only the owner of the signature knows the private key, but everyone knows the public key. The private key is used to cryptographically sign a file or message, which anyone can use the public key to verify, but only the owner of the keypair can use their key to sign. This means that if a message is signed, it must've come from the correct person.

Message Digests

A message digest (or hash function) takes a message M and produces a small 'fingerprint' known as the message digest $H(M)$. A secure hash function H encrypts a small block of the message which is produced as a function of the message, known as the authenticator. The properties of H are such that $H(M) \neq H(M')$ and given $H(M)$ is computationally impossible to find M . The message digest is encrypted with the sender's private key, and serves as a signature that verifies the content and order of the message. Two popular digest functions are MD5 and SHA-1.

Authentication

Authentication verifies that a message has come from a verified source. This can be achieved using conventional encryption techniques as discussed previously. This does not protect against so-called playback attacks however. If a third party records an encrypted message, they can re-send it again at a later time, and since the message is correctly encrypted, the receiver will still believe it to be a valid message.

This can be avoided by adding a number to the cipher that relates to the time the message was sent, either directly or by using it as the seed for a random number generator. When the message is received, if the time is too far in the past, the message will either be ignored, or requested for re-sending to ensure the correct message was sent.

Lecture - Local Area Networks

11:00

04/11/24

Asim Ali

Data Link Layer

Within the Data Link layer of the OSI reference model, there are two control systems– the Logical Link Control and Media Access Control. The logical link control is responsible for interfacing with the layers above it, and performing flow and error control.

Media Access Control is responsible for more of the addressing and actual communications, and handles assembling and disassembling the transmission frame, which consists of

- Receiver's address
- Error detection data, e.g. checksum or crc

Media Access Protocol

The original ethernet protocol was designed for use in bus or star networks, and so includes error and collision handling. Since only one station is able to transmit at a time in a bus or (hub) star network, there needs to be a way to determine if a collision has occurred and how it should be resolved. This is made up of a set of rules–

- If there is already a station transmitting, do not transmit.
- If another station transmits while you are transmitting, jam the medium and wait for a random length of time, using the back-off algorithm. (Select a random length of time from 0 to $2^n - 1$ slots, where n is the number of collisions).
- If there have been an excessive number of collisions, cancel the transmission and send an error up through the layers.
- After waiting, attempt to transmit again.

The length of one transmission slot depends upon the speed of the medium, and the minimum size of packets. If the slot size is 512bits, then at 10Mbps, the slot time is $\frac{512}{10000000}$ or 0.0000512 seconds (or 51.2 μ sec)

To determine if a collision has occurred, each station is connected to the medium with two wires, one for transmitting and one for receiving. During transmission, if the signal sent out on the transmission line is not the same as what is received, then another station is interfering with the transmission, and a collision has occurred.

Switching

Store-and-Forward

When the switch receives a packet, it waits for the entire packet to arrive and stores it in memory before forwarding it to the intended recipient. This mode of operation works for both half and full-duplex transmission.

The latency in this mode can be calculated by sender-switch transmission time (packet size / transmission speed) + switch latency + switch-recipient transmission time + propagation delay of the medium (how long the data takes to get from one end of the medium to the other).

Virtual-cut Through

When the switch receives a packet, it only buffers enough of the packet to allow it time to process the MAC header, and determine where the packet should be sent. It then just forwards the buffer and the data as it flows into the switch.

The latency here can be calculated by sender-switch transmission time + switch latency + propagation delay of the medium.

Ethernet Frame Format

Each frame sent on an ethernet network contains

- Preamble– A series of 56 1s and 0s, used to sync the clock of the receiver (56 bits)
- Frame Delimiter– A byte which signals the start of the frame, specifically 10101011 (8 bits)
- Source– Address of the sending station (48 bits)
- Destination– Address of the recipient station (48 bits)
- Length– The length of the data in bytes (16 bits)
- Data– Up to 12000 bits (1500 bytes) of data (0 – 12000 bits)
- Pad– Padding added to ensure the data is at least 46 bytes long (0 – 368 bits)
- CRC– Cyclic Redundancy Check information, used to detect errors in the frame, so it can be re-sent on error (32 bits)

This gives us a minimum frame length of $48 + 48 + 16 + 0 + 368 + 32 = 512$ bits

Lecture - Wireless Local Area Networks

11:00

11/11/24

Asim Ali

A single cell network is one which has a single Access Point (or Control Module), using a single frequency, but may have any number of clients (User Modules). A multiple-cell network may have multiple access points using different frequencies, each with multiple clients, but they must be connected to the same wired network for back-haul.

Each cell acts like a star network, since client can only communicate with the access point, and not directly with any of the other clients on the same network. In a multiple-cell network, they also cannot communicate directly between cells without a wired back-haul, which is usually part of an existing wired network.

Before a client can communicate in a cell, it needs to associate with the access point. In a single cell network, this is simple since there is only ever a single access point that needs to be associated with. In a multi-cell network, the client must disassociate from one access point, and associate to the new one. This may be done manually, if each access point is broadcasting a different SSID, or it may be done automatically if it's a cohesive network where the access points communicate to hand-off clients.

There are also ad-hoc WLAN networks, in which the devices are all simultaneously clients and access points, and communicate directly with each other.

IEEE 802.11

Terms

- Associated Stations– Devices connected to a WLAN
- Access Point– Provides access to the network for associated devices
- Basic Service Set (BSS)– A set of stations controlled by one coordinator
- Extended Service Set (ESS)– A set of one or more connected BSSs
- Distribution System (DS)– A system used to connect a set of BSSs and LANs to create an ESS
- Frame– A unit of data in the MAC protocol

Services

The MAC Service Data Unit is effectively the mechanism by which packets are prepared and sent over the medium.

Service	Provider	Used for
Association	Distribution System	MSDU
Authentication	Stations	Access & Security
De-authentication	Stations	Access & Security
Disassociation	Distribution System	MSDU
Distribution	Distribution System	MSDU
Integration	Distribution System	MSDU
MAC Service Data Unit (MSDU)	Stations	
Privacy	Stations	Access & Security
Re-association	Distribution System	MSDU

Access Control (CSMA/CA)

A very similar protocol to CSMA/CD, but designed specifically for wireless networks rather than wired bus networks. The difference is that the nodes might not know that they are causing a collision, since they may be in range of the access point, but not the other transmitting node.

There is an optional centralised controller known as the Point Coordination Function (PCF), which provides collision avoidance by providing permission to each node before it is allowed to send a message. There is also a distributed version known as the distributed coordination function (DCF), which is used in ad-hoc networks or in networks where throughput is more important than latency.

Point Coordination Function

The controller (usually the access point) offers a transmission slot to each known client, which the client then replies to, saying that it either does or does not wish to transmit. If the client does not wish to transmit, the slot is offered to the next client, and so on. It uses a round-robin method to give each client an opportunity to transmit, and guarantees there are no collisions. This does however waste a lot of capacity, since nodes often do not wish to transmit, and if there are some devices with lots of data, and some with no data, they are still given the same share of the throughput.

The access point also sends out a 'beacon frame' 10 – 100 times per second, which advertises the network to any stations within the coverage area. The frame also contains information such as hopping frequencies, dwell times, clock synchronisation and more.

Distributed Coordination Function

DCF uses CSMA/CA to provide collision avoidance, as well as various levels of Inter-Frame Spacing (IFS). There are two methods which a network can use to confirm that a packet has been received correctly.

- Frame Exchange Protocol
 - Sender transmits frame immediately
 - Receiver response with an acknowledgement
 - If sender doesn't receive acknowledgement, it retransmits the frame
- Four Frame Exchange
 - Sender sends request-to-send (RTS)
 - Receiver responds with a clear-to-send (CTS)
 - Sender transmits frame
 - Receiver responds with an acknowledgement

If any station receives an RTS or CTS message not addressed to them, they immediately put themselves into a Non-Active Mode (NAV) until the medium is again idle. That way, even if the station is not in range of the requesting station, it should always be in range of the clearing station, as it will most likely be the access point. If the sending station does not receive an acknowledgement within a short time, the transmission is cancelled and it must send a new RTS before it can transmit.

The Distributed Inter-Frame Spacing (DIFS) is the length of time that a node must wait between transmissions, and the length of time that is multiplied to get the exponential back-off delays.

Priority Traffic

Some traffic needs to be given priority over others, such as acknowledgements, CTS messages and poll responses. These types of traffic use the Short Inter-Frame Spacing (SIFS) so they are more likely to transmit immediately. The access point also has priority over all other stations, and uses the Point-Coordinator

Inter-Frame Spacing (PIFS), which is shorter than the DIFS, but longer than the SIFS.

Super Frames

The bandwidth of the frequency is split into so-called **Super Frames**, which are a certain length of time.

Each 'Super Frame' has two sections. The first part of the frame uses PCF to poll each station for time-sensitive frames which need to be sent before anything else. For the second part, the point-coordinator goes idle and allows time for DCF frames to be sent asynchronously. This allows both time-sensitive and throughput-sensitive traffic on the same network.

If the end of one transmission overlaps the start of a super frame, the length of the super frame is reduced and the PCF section is delayed as to not collide with the previous transmission. This typically means that the DCF time in this super frame is reduced.

Lecture - Bluetooth Networks

11:00

18/11/24

Asim Ali

The original purpose was to create a universal interface for using small, local ad-hoc networks. It was also intended as a replacement for irDA— an infrared communication protocol.

The design specifications required a maximum range of 10 – 100m, low power consumption, and a license free frequency to keep costs low. The specific frequency it uses is roughly 2.45GHz, which is within the license free frequency band also occupied by 2.4GHz Wi-Fi.

Bluetooth Protocols

Ordered from lowest to highest in the protocol stack, these are some of the main protocols used to control and communicate between bluetooth devices.

Radio

Responsible for frequency hopping, modulation and demodulation, and managing transmission power.

Baseband

Responsible for establishing connections, addressing different devices, formatting packets, timing messages and controlling transmission power.

Link Manager Protocol

The Link Manager Protocol (LMP) is responsible for initiating and maintaining connections between devices, authentication and encryption, and is in the lower layers of the bluetooth protocol.

Logical Link Control and Adaption

Adapts existing protocols to work over a baseband network, such as TCP/IP and telephony. It offers connectionless and connection-oriented services, similar to TCP and UDP.

Service Discovery Protocol

Provides device information and service advertisement, which can be queried by other devices to initiate establishing a connection.

Piconet

A **Piconet** is a collection of bluetooth devices connected in an ad-hoc manner. They form a star network with one device acting as the master, and up to 255 devices connected to it as clients. All devices in a piconet are synchronised to the same hopping sequence, and each piconet has its own hopping pattern to avoid interfering with each other. Before a device is able to join a network, it has to synchronise with the hopping pattern.

While up to 255 devices can be in a piconet, only 7 devices can be actively communicating with the master at any time. All other devices are placed into a parked state, where they master is still aware they exist and may want to connect, but they cannot actively communicate.

Active clients can be in one of three states– Active (Transmits and receives), Sniffing (Reduced number of slots at a reduced power) or Holding (Further reduced power, only supporting SCO links).

There is also a standby state, in which the device is in range of the piconet, but is not participating.

Radio Specifications

Include table from slide 9

Baseband Specification

Bluetooth baseband uses Frequency Hopping Spread Spectrum (FHSS) to make it less likely for interference to occur, and to make it harder to eavesdrop on communications. The hopping frequency for baseband is 1600 hops/sec, with a dwell time of 625 µsec, which is also the slot time. There are 79 carrier frequencies within the band, each of which is hopped to in a pre-determined but pseudorandom sequence, which further minimises both the likelihood and impact of interference.

Bluetooth devices use Time Division Duplex, where the same frequency is used for sending and receiving data, but not at the same time. Each slot is labelled from 0 to $2^{27} - 1$, and loops. The master uses even slots, and clients share the odd slots using Time Division Multiple Access (TDMA). In general, Piconets use FH-TDD-TDMA (Frequency Hopping Time Division Duplex Time Division Multiple Access).

Scatternet

Multiple piconets may exist in one area, and may form a scatternet. A bridge device connects two piconets together, and may be a client in both networks, or a master in one and client in another. The bridge acts as a relay between the two networks, and forwards messages between them. Clients transmit only with the permission of the master, using Time Division Multiplexing to access the medium.

Links Between Master and Clients

Bluetooth Packet Format

- Access Code– 72 bits used for timing synchronisation, offset compensation, paging and inquiry
- Header– 54 bits used to identify packet type and contain protocol control information
- Payload– 0 – 2745 bits of actual data

Synchronous Connection Oriented (SCO)

A fixed number of slots are assigned between point-to-point and master-client communication. The master maintains connections with all devices using reserved slots at a regular interval. Slots must be reserved, and the smallest unit of reservation is 2 consecutive slots– one in each direction. The payload of each slot may be 80, 160 or 240 bits. The most reliable variant is 80 bits, at 800 slots/sec as the packets are replicated 3 times. The master can support three simultaneous links, if each of them uses the 80 bit variant. SCO communications are never re-transmitted, and error correction techniques are used at each end instead. Each variant uses a different type of Forward Error Correction (FEC), 80 using 1/3 FEC, 160 using 2/3 FEC and 240 using no FEC, where the fraction determines how much of the bandwidth is used for actual data. This is typically used for realtime data, video and audio.

Asynchronous Connectionless (ACL)

There can be only a single ACL link at a time, and transmits in slots not already reserved for SCO communications. ACL links use 1, 3 or 5 slot packets, which are assigned by the master, depending upon the requirements of devices on the network. The first 259 μ sec of the slot are used as settling time, then 126 bits of header, and finally 240, 1490 or 2740 bits of data, depending on the number of slots used. There is no bandwidth guaranteed, and retransmission may be required. A client may only reply to the master with an ACL slot, and may never initiate communication without the master having assigned slots. After every ACL slot, a single SCO slot is reserved for the client to reply to the master.

Lecture - Cellular Networks

11:00

25/11/24

Asim Ali

In a cellular network, each cell contains a base station, a number of relay stations and a number of client devices, called user equipment. Each cell has a set bandwidth and number of available channels, each of which can be used to communicate with one device at a time. For example, if a base station has 10 channels, it can communicate directly with 10 devices at once. In this context, a base station is very similar to that in a WLAN network, as it has a controller which handles communication within the cell's channels and a set of transceivers, one for each channel.

A relay station is a smaller station which has equipment to relay communications from devices to the base station. These are usually used near the edges of the cell, or in places with particularly poor reception, such as big concrete warehouses. This is especially needed in cells that span a larger area, since they can be anywhere from a few hundred to 30,000 metres in size.

A cellular network is made up of many cells in close proximity, which form an area of continuous coverage. Cells typically have omnidirectional transceivers, meaning they cover a circular area around them, and usually overlap with other cells around the edges. The space each cell actually covers is usually a hexagon, which is the most effective way to pack cells together, and the overlaps are ignored.

Cell Geometry

The area of each cell is $2.6 \times R \times R$ or $2.6R^2$, and the distance between the centre of each cell is $d = \sqrt{3}R$. In this case, R is the radius and d the diameter of the circular coverage range of the cell, and the given area is the area of the theoretical hexagonal area it covers.

Frequency Reuse

It is not possible to use the same frequency in directly adjacent cells, since the signals would interfere with each other. After a certain distance, D , it becomes possible to reuse frequencies since the cells do not have overlapping coverage areas. This distance can be calculated by $\frac{D}{d} = \sqrt{N}$ or replacing d , $\frac{D}{R} = \sqrt{3N}$ where N is the number of cells in a repeating pattern.

Two cells using the same frequency cause co-channel interference (CCI). The frequency reuse distance can be evaluated using the following method with two variables, i and j .

- Move i steps from the reference cell in any direction
- Turn counter-clockwise by 60 degrees
- Move j steps in that direction
- The number of cells in a cluster, $N = i^2 + j^2 + ij$

Calculating Network Capacity

In a network with 32 cells, each of which have a radius of 1.6km, a total frequency bandwidth of 336 channels, and reuse factor of 7, how would you find the total coverage area and total channel capacity?

- Area of each cell is $2.6 \times 1.6 \times 1.6 = 6.7km^2$
- Area of all cells is $6.7 \times 32 = 212.9km^2$
- Number of channels per cell is $\frac{336}{7} = 48$

- Total number of channels is $48 \times 32 = 1536$

Increasing Cell Capacity

There are several ways to increase the capacity of an individual cell within a cellular network. One option is to simply add more channels, but this is not always possible. To increase capacity without increasing the number of channels, you can use

- **Frequency Borrowing**– Cells with lower usage can have their channels ‘borrowed’ by other cells which are more congested.
- **Cell Splitting**– Congested cells can be split into smaller cells, reducing the radius of a cell by a factor of p , and increasing the number of base stations by p^2 , which increases the number of channels in the same area.
- **Cell Sectioning**– Congested cells can be divided into a number of wedge-shaped sectors, each of which have their own set of channels transmitted using directional antennae such that they don’t overlap.
- **MicroCells**– Antennae are moved to buildings, lamp posts, etc, in cities and have their transmission power greatly reduced, such that they only cover a small area, but with far more cells to increase the capacity.

Cellular Network Operations

Out of every 66 total channels, 11 are control channels and 55 are traffic channels. The client device listens to the control channels, and selects a base station to connect to, based on the relative strength of the control channels it can receive. Base stations are connected to a Mobile Telecommunications Switching Office (MTSO), which connects the cellular network to the main telephone system. The MTSO also assigns channels, performs hand-offs and monitors calls to meter bills. Each MTSO is usually responsible for more than one cell, and often one cluster of cells as defined by the frequency reuse pattern.

GSM Networks

diagram of GSM network?

Channels

Each GSM provider is typically assigned two 25MHz chunks of the spectrum, one for uplink channels and one for downlink channels. Within this, each channel is 200KHz, meaning that there can be up to 125 channels for uplink and downlink. This means that there are 125 full-duplex channels, each of which can be split up using 8-way TDMA, to get a total user count of 1000 in each cell.

2nd Generation, CDMA

Up to 35 users per channel, using CDMA with a 64×64 Walsh Matrix.

Lecture - Software-Defined Networks (SDN)

11:00

02/12/24

Asim Ali

VLANS

VLANS are virtual networks made up of one or more physical LANs, combined together using switches to create a virtual LAN. This is useful for separating different types of devices on the same network, and making sure they don't conflict with each other.

Management, Control and Data Planes

The Data plane is responsible for transmitting and receiving packets, as well as forwarding them from one device to another. It is also responsible for enforcing the routing rules decided upon by the Control plane. Uses common protocols such as TCP/IP and UDP. All functions relating to data reside within the Data plane.

The Control plane is responsible for routing packets between networks, naming devices, policy declaration (such as bandwidth and speeds, encryption and authentication), security and routing protocols such as OSPF, BGP and MPLS. It also decides upon the route packets should take from the source to the destination, which is enforced by the Data plane.

The Management plane is responsible for configuration and monitoring of network devices. Uses protocols such as SSH, Telnet and SNMP.

SDN Architecture

The network is logically divided into three planes– the Data, Control and Application planes. The data plane contains the physical switches, routers and devices in the network. The control plane is made up of software which manages routing, traffic engineering and mobility, typically running on a physical machine somewhere on the network. The application plane contains the SDN applications, business applications and any cloud orchestration software.

The OpenFlow API is used to communicate between the SDN software and physical network devices, which performs the configuration and monitoring of devices, to form the network how the control plane decides. Because the control layer is decoupled from the physical devices, it has only a single interface, allowing administrators to easily and dynamically adjust the network to meet the needs of applications.

OpenFlow Switches

The first packet of each new device is encapsulated and forwarded to the SDN controller, so routing decisions can be made. The switch forwards all incoming packets through the appropriate port using a flow table. Each flow can temporarily or permanently drop packets, as decided by the controller.

Message Types

- Controller-to-Switch– Sent from the controller to one or all switches to configure the switch, or retrieve information from it
- Asynchronous (Switch-to-Controller)– Sent from the switch to the controller to report packet information, state changes, error rates, and so on

- Symmetric (Switch-to-Controller or Controller-to-Switch)– Messages sent without solicitation from either side, typically used for health check messages or to check that channels are available

Controller Actions

- Reactive Mode
 - The controller is queried when a packet arrives that has no flow entry or an expired timer
 - The controller creates and installs a rule to the flow table of switches on the path with rules for that specific packet
 - The first switch then forwards the packet according to the new rule(s)
- Proactive Mode
 - The controller creates flow rules for all possible types of traffic
 - These rules are installed statically ahead of time
 - The controller is not queried since all incoming packets should find a flow entry
 - This means that packets are switched at line-rate with no delay, but implementing the tables requires lots of expensive memory for moderately large tables

Lecture - Satellite Networks

11:00

09/12/24

Asim Ali

Most communication satellites effectively act as a reflector which sends signals between two stations on the ground. They typically have several transponders, each of which is able to listen to several portions of the spectrum, and amplify the incoming (uplink) signal, change its frequency and then broadcast it to the receiving station (downlink). Each ground station has its own transmitters and receivers, with antennae pointed towards the satellite. They allow communication between two points on the ground which do not have line-of-sight, typically due to their distance and the curvature of the Earth.

Satellites can either be point-to-point, or have a broadcast link, where there is one transmitter, and the signal is amplified and broadcasted to multiple receivers on the ground.

Satellite Orbits

There are several important factors which determine which orbit a satellite will use, such as where it needs to have line-of-sight, and when. All orbits have an apogee (the furthest point from Earth) and a perigee (closest point to Earth). If the orbit is elliptical, then it is usually off-centre, and so the satellite gets much closer and further away along its orbit. There are also circular orbits, which are more useful for Geostationary orbits, or for satellites which need to be a constant distance from the ground.

Orbital Periods

The orbital period is the time it takes for the satellite to make one complete orbit around the planet, and it can be calculated as

$$T = \frac{(R + A)^{1.5}}{100}$$

where R is the radius of the earth (6378Km), A is the altitude of the satellite, and T is the orbital period in seconds.

Different Orbits

There are three main types of orbits which are used for communication satellites–

- Low-Earth Orbit (LEO)– 500 to 1500Km altitude, moving very quickly. Much lower latency than other orbits
- Medium-Earth Orbit (MEO)– 5000 to 18000Km altitude, moving slower than LEO. Higher latency than LEO but lower than geostationary
- Geostationary-Earth Orbit (GEO)– Roughly 35786Km altitude, moving much slower than MEO. Much higher latency than LEO or MEO, but always above the same point on Earth, so useful for static communication links

There are also two layers of the Van Allen belt, which are layers of charged particles, which very quickly destroy any satellites which enter them. They are from 1500 to 5000 and 15000 to 20000Km in altitude.

Using a higher orbit increases the coverage area of each satellite, which means that less are needed for full coverage. However, it also greatly increases latency and so is not suitable for realtime communications.

Geostationary Satellites

Since the orbital period of a geostationary satellite is 24 hours, they always stay above the same point on the equator, which means that antennae don't need to constantly adjust to follow the satellite as they move across the sky. They provide great coverage, theoretically covering 120 degrees of the surface, everywhere except the poles, meaning that only 3 satellites are needed for total coverage. There are at most 180 satellites in a perfect geostationary orbit at any time, as they need 2 degrees of rotational separation to ensure they don't collide. The allocation of 'slots' is managed by the ITU.

LEO Satellites

Typically have a circular or slightly elliptical orbit, with an altitude lower than 2000Km. The orbital period is typically anywhere from 1.5 to 2 hours, and they typically have a coverage area with a diameter of roughly 8000Km. They are typically only visible for up to 20 minutes at any given point, and so the system must be able to aim the antennae to maintain line-of-sight. There is also a large amount of doppler shift in the signal due to the very high relative speed of the satellites. Because they orbit so low, they experience atmospheric drag and must boost their altitude often to avoid orbital deterioration and then falling back to Earth and burning up.

MEO Satellites

Can have either circular or elliptical orbit, with an altitude from 5000 to 18000Km. The orbital period is typically roughly 6 hours. They do not experience atmospheric drag, and so typically have a longer lifespan than LEO satellites. GPS satellites are typically MEO, and orbit at around 18000Km. Their orbits are usually inclined with respect to the equator, which means that they cover a larger portion of the surface over time.

Frequency Bands

add table here

Transmission Performance

There are several factors which determine the performance of a satellite communication link–

- The distance between the station antennae and the satellite antennae, due to signal dissipation and latency
- The ground distance between the station antennae and the main focus point of the satellite
- Atmospheric attenuation due to varying weather and atmospheric pressure

Satellite Capacity Allocation

Since there is a limited number of satellites in each orbit, and each satellite has a limited number of transponders, to serve a large number of users it is necessary to use multiplexing to increase capacity. Methods discussed before are used.

Frequency Division Multiple Access (FDMA)

The bandwidth assigned to the satellite is divided into channels, for example a 500MHz bandwidth divided into 24 40MHz channels, with a guard band of 4MHz built in. Each of the channels could be used for any purpose, and they can all transmit simultaneously without colliding. There are only 12 40MHz channels in a 500MHz band, but satellites often use horizontal and vertical polarisation to double the bandwidth available. Within each channel, there can be multiple stations, each of which uses a different amount of bandwidth and different frequencies.

Time Division Multiple Access (TDMA)

There are a set number of slots in each second, which are divided up so that each station has the same share of time. The uplink is divided up and each of the stations transmits at a different time, only to the satellite. The downlink is divided up and the satellite broadcasts all of the data to all of the stations, and they ignore any data destined for the other stations. There is also a guard time between each slot, to make sure that there are no collisions. At the start of each frame, there is a so-called 'reference burst', which are used to denote the start of a new frame, and to synchronise the clocks in the devices.